Assessing the performance of the SpeedInfo sensor

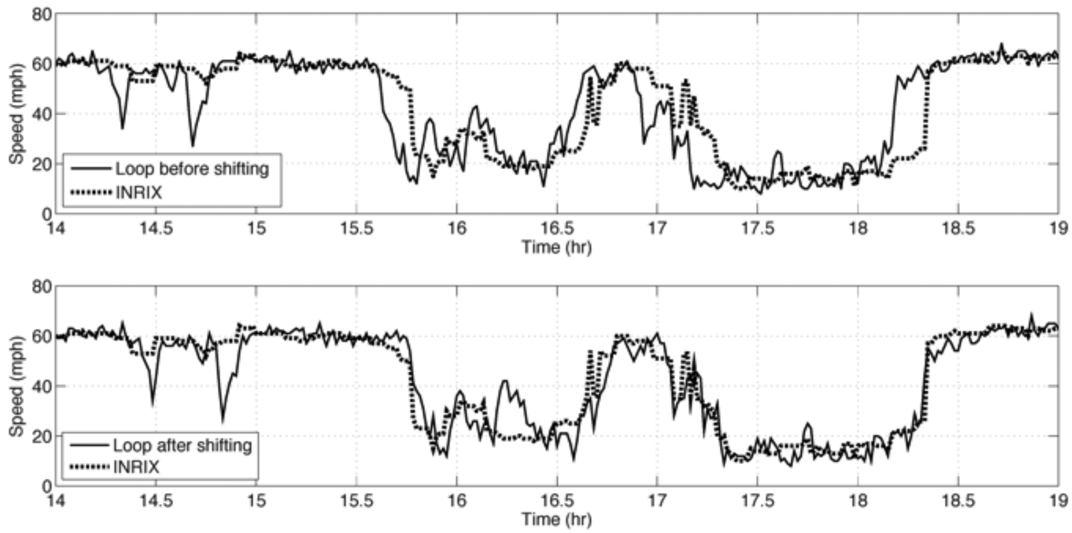

Prepared by:
Benjamin Coifman
Seoungbum Kim

Prepared for:
The Ohio Department of Transportation,
Office of Statewide Planning & Research

State Job Number 134723

December 2013

Final Report

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
FHWA/OH-2013/21		

4. Title and Subtitle	5. Report Date
Assessing the performance of the SpeedInfo sensor	**December 2013**
	6. Performing Organization Code
7. Author(s)	8. Performing Organization Report No.
Benjamin Coifman, Seoungbum Kim	
9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
The Ohio State University **Hitchcock Hall 470** **2070 Neil Ave.** **Columbus, OH 43210**	
	11. Contract or Grant No.
	State Job Number 134723
12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Ohio Department of Transportation **1980 West Broad Street** **Columbus, Ohio 43223**	
	14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

Until recently freeway traffic operations data were collected in house by the Ohio Department of Transportation (ODOT) using loop detectors. In recent years several private companies have emerged with a viable opportunity to outsource traffic data collection. While these new business models and ease of deployment are compelling, there are few published studies explicitly evaluating their performance. ODOT has contracted with two of these data providers: SpeedInfo and INRIX. This study exploited a unique juncture as ODOT transitioned to the new third party data collection: the two new systems and the legacy loop detectors were concurrently operational in Columbus, OH for approximately 6 months. Using the archived data from this period, we evaluated SpeedInfo and INRIX performance over a 14 mi. This long period of time ensured that we would have a better chance of observing intermittent features that might go unobserved in a short-term study.

Both SpeedInfo and INRIX traffic data generally performed within specifications, exhibiting small biases that could likely be fixed with fine-tuning. Both sensors also exhibited a reporting lag, on the order of 60 sec for SpeedInfo and 360 seconds for INRIX (comparable to conventional 60 sec and 5 min aggregation periods used in loop detectors). Although the lag and bias are within specifications, some applications might be sensitive to them, e.g., traffic responsive ramp metering. Each system also exhibited unique behavior not found in the other.

SpeedInfo is sensitive to precipitation: sometimes going offline for tens of minutes. When mounted on the shoulder, SpeedInfo showed slightly degraded performance on the far side traffic. INRIX reported speeds every minute, but after excluding repeated values, the actual effective reporting period was more like 3-5 min, with occasional periods of repeated measurements lasting in excess of 10 min (already reflected in the lag above). INRIX also provides two measures of confidence, neither of which reflects the measurement repetition.

17. Keywords	18. Distribution Statement
Highway traffic monitoring, SpeedInfo, INRIX, Loop detectors	**No restrictions. This document is available to the public through the National Technical Information Service, Springfield, Virginia 22161**

19. Security Classification (of this report)	20. Security Classification (of this page)	21. No. of Pages	22. Price
Unclassified	**Unclassified**	**59**	

Form DOT F 1700.7 (8-72) **Reproduction of completed pages authorized**

Assessing the performance of the SpeedInfo sensor

Prepared by:

Benjamin Coifman, PhD

Associate Professor
The Ohio State University
Joint appointment with the Department of Civil, Environmental, and Geodetic Engineering, and the Department of Electrical and Computer Engineering
Hitchcock Hall 470
2070 Neil Ave, Columbus, OH 43210
Phone: (614) 292-4282
E-mail: Coifman.1@OSU.edu

Seoungbum Kim, PhD

Research Associate
The Ohio State University
Department of Civil, Environmental, and Geodetic Engineering
E-mail: kim.1936@osu.edu

December 2013

Prepared in cooperation with the Ohio Department of Transportation and the U.S. Department of Transportation, Federal Highway Administration

We are particularly grateful for the assistance and input from George Saylor, Bryan Comer, Merih Ocbazghi, Jason Yeray, and John MacAdam at the Ohio Department of Transportation, and Zhuojun Jiang at the Mid-Ohio Regional Planning Commission (MORPC).

Table of Contents

List of Figures

List of Tables

1. Introduction

Until recently freeway traffic operations data were collected in house by the Ohio Department of Transportation (ODOT) using loop detectors that provide speed, flow and occupancy by lane. Ultimately, most real-time management applications only use the speed measurements (e.g., traveler information and congestion monitoring). In recent years a new company, SpeedInfo, emerged with a different paradigm for traffic operations data collection. The SpeedInfo approach is revolutionary in two ways. First, at the sensor level, they employ a single unit that can be strapped to an existing pole and configured in less than half an hour. The unit contains a Doppler radar to measure traffic speeds, a wireless modem, a solar panel, a battery, and an onboard processor. In principle they should provide very accurate speeds because of the radar sensor. Typically the units report speed in two directions (approaching and receding) every minute. The SpeedInfo sensor is attractive because it is completely wireless and out of the right of way. Second, the SpeedInfo business model is innovative. Instead of selling hardware, the company sells data, while the company typically owns and maintains the sensors. SpeedInfo collects the traffic data and covers the communications costs, they then sell the speed data. In effect SpeedInfo offers a viable opportunity to outsource traffic data collection.

The SpeedInfo costs are compelling compared to traditional traffic detection and ODOT has deployed SpeedInfo statewide, with most urban freeways covered at a density of one bidirectional station per mile and rural freeways at a lower density. The measurements are publicly available on a real-time map at Buckeye Traffic [1] aggregated into three speed bins, and are available both real-time and archived at a higher resolution in-house at ODOT.

While the SpeedInfo business model and ease of deployment is frequently cited in the literature, there are few published studies explicitly evaluating the performance SpeedInfo sensor. We were only able to find a few documents that either directly or indirectly provide insight into the performance of SpeedInfo sensors. A PhD dissertation used SpeedInfo data to validate link travel times from a GPS equipped probe vehicle [2], but by extension, the GPS also provides some degree of validation for SpeedInfo. The study compared the concurrent SpeedInfo: 4 min moving average spot speeds at 13 locations against the instantaneous GPS spot speeds from 31 probe vehicle passes of each location. The author excluded 23% of the SpeedInfo locations (3/13) because of large differences relative to the GPS speeds, the average difference at these locations ranged between 6 and 20 mph. A scatter plot of GPS speeds versus SpeedInfo speeds from the remaining data (Fig 7.5 in the dissertation) shows strong correlation between the two sources during free flow periods but almost no correlation when either speed is below 50 mph. The overall poor correlation at 23% of the locations and poor performance during congestion at the remaining 77% of the locations merit further study, the discrepancies may simply be due to comparing an instantaneous speed in one lane against a four-minute average across all lanes. Finally, the SpeedInfo sensors were eliminated from further consideration in an Arizona study due to the proprietary data restrictions [3]. The report discusses experiences at three other agencies that had deployed SpeedInfo and at least at the time of the report, the extent of validation by these agencies was minimal. Needless to say, to the extent that studies of SpeedInfo's accuracy and performance exist, they appear to be small scale and largely unpublished.

Since ODOT adopted SpeedInfo other emerging technologies now offer a similar traffic data-stream without the need for the operating agency to maintain sensors. A prime example is INRIX, which, "aggregates traffic-related information from millions of GPS-enabled vehicles and mobile devices, traditional road sensors and hundreds of other sources," [4]. ODOT purchased a license to use INRIX traffic data, reportedly with the primary intent to monitor non-instrumented portions of their network. Although the INRIX process is proprietary, in the case of Ohio, the traffic data are believed to come almost exclusively from probe vehicles, e.g., automatic vehicle location (AVL) for fleet management and GPS equipped consumer electronic devices.

There has been an ongoing performance evaluation of the INRIX data in the I-95 Corridor Coalition, where INRIX has been an active partner in the consortium. Much of the evaluations consist of instantaneous comparisons offering only a "binary outcome," and the tests, "average errors over all time intervals, [so] only a few high-variance outliers are needed to invalidate the whole segment, even if the majority of the data points are within the acceptable range," [5]. Instead, [5] contemplated a time series evaluation. Although the authors do not mention it, Figure 1 in the paper shows a stable 5 mph bias between the INRIX measurements and the ground truth data. Several other papers evaluated INRIX data during strictly free flow periods, e.g., [6-7]. Here [6] examined a low volume, rural interstate and found the INRIX data had an accuracy of 80-90%, while [7] found INRIX exhibited a 6 mph bias relative to ground truth. Finally, [8] looked at INRIX performance under recurring congestion in an urban area and found INRIX exhibited both a bias and latency, though the authors did not quantify either. Table 1 in [5] presents 7.5 hrs of INRIX and ground truth data aggregated to 5 min and as it turns out, these data appear to be the same set used to reach the conclusions in [8]. Although the authors of [5] did not appreciate the detailed time series, we re-plotted these data and focused on the six major speed transitions evident therein, in this case the INRIX data exhibits a 10 min lag. After accounting for this lag, this particular INRIX dataset exhibited a bias of only 1.6 mph.

Our study sought to exploit a unique juncture as ODOT transitioned to the SpeedInfo data-stream. The Columbus Freeway Management System (CMFMS) had roughly 65 loop detector stations covering freeway segments that are also covered by the new SpeedInfo sensors. These loop detector stations are unique in their own right, rather than aggregating the data in the field, they send all of the vehicle actuations back to the traffic management center (TMC). Before aggregation at the TMC, the raw actuation data are archived [9]. With the actuation data, we can identify and correct for many errors at the loop detectors [10-12], yielding very accurate measurements (as have been verified by hundreds of GPS probe vehicle runs and manually generated ground truth from concurrent video for over 70,000 vehicle actuations).

The CMFMS loop detector stations were decommissioned approximately half a year after the concurrent SpeedInfo stations came on line. Using the archived data from this period, we were able to evaluate the SpeedInfo performance over a 14 mi corridor over several months while the two traffic data collection systems were operating. This period included potentially challenging conditions for the SpeedInfo radar, both recurrent (rush hour congestion and late night low flow), and non-recurrent (incidents and precipitation). Ultimately we studied 5 months of data at 14 of the SpeedInfo deployments. This long scope of time ensured that we would have a better chance of observing intermittent features that could easily go unobserved in a typical validation study against manual ground truth that lasts only a few hours. In our analysis we undertook a broad overview of the data, comparing speed measurements from SpeedInfo against the concurrent aggregated loop detector data and then we explicitly looked for any anomalous patterns such as latency and system outages.

Recognizing the potential to leverage this effort, the scope of this study was expanded to include the concurrent INRIX data from 2011 in the Columbus metropolitan area as well. According to ODOT, INRIX claims that the quality of their data in this region has improved significantly since the CMFMS was decommissioned in 2011. So as a result, we also consider a second batch of INRIX data from the same corridor, collected during the spring of 2013.

1.1 Research Objectives

This research project sought to use the archived coverage overlap to evaluate the SpeedInfo performance throughout the extended time period when the two systems are operating. The specific objectives as proposed were:

1) Collect the archived SpeedInfo data from ODOT as well as any pertinent configuration or deployment information such as the exact location of each sensor.
2) Correlate the location of the SpeedInfo sensors and detection zones with the loop detector stations.
3) Develop a method to easily time synchronize the two data sources.
4) Aggregate the loop detector data at the same rate as the SpeedInfo data, while also developing reasonable uncertainty bounds to account for clock drift and spatial offsets.
5) Compare all of the overlapping data from the two sensor systems, look for any biases or apparent discrepancies between the two datasets. Special care will be taken to separately examine periods that may be challenging to the SpeedInfo sensors, e.g., congestion, precipitation, and early morning.
6) When discrepancies are found, use the individual vehicle actuation data from the loop detectors to diagnose the problems.
7) Investigate and document potential biases inherent in the radar system to help ODOT work more effectively with the new data source.
8) Quarterly reports and final report.

After the study was underway, at ODOT's request we expanded the scope to include the INRIX traffic reports.

1.2 Overview

The remainder of this report is as follows, in Section 2 we present the analysis area and sensors used in this study. Section 3 presents the performance of the SpeedInfo sensors in three subsections: Section 3.1 presents the overall performance of the SpeedInfo sensors versus the loop detector system and discusses the results. In Section 3.2 we strictly focus on the SpeedInfo sensor to investigate outages observed in the data and the possible interaction with weather conditions. Section 3.3 discusses how the SpeedInfo performance varied across the study locations, depending on where a given sensor was mounted- either in the median or on one shoulder. Section 4 presents the performance of the INRIX data in three subsections: Section 4.1 presents the overall performance of the INRIX data versus the loop detector system and discusses the results. Section 4.2 examines repeated measurements reported by INRIX- although INRIX did not exhibit weather related performance drops, the system typically would report the same measurement for several successive sample periods. Section 4.3 examines the confidence score and confidence value reported by INRIX. Finally, Section 5 presents the conclusions and summarizes the results of this study.

2. Analysis Area

Figure 1a shows the freeway corridors (dashed lines) in the greater Columbus, OH area along with the SpeedInfo sensors (solid gray circles) and loop detector stations (solid black squares). When the SpeedInfo system was deployed, no effort was made to coordinate sensor placement with the existing loop detectors. So to minimize any offset between the two traffic monitoring systems, the I-71 corridor was chosen for this study because it has the highest density of loop detector stations: typically 3 stations per mile. Figure 1b shows the corridor, with the southern end at the I-70/I-71/SR-315 interchange in the central business district (CBD) and the northern end outside the I-270 beltway in the northern suburbs, covering roughly 13 mi, in the shape of a “J”. The loop detector station numbers are shown adjacent to the given detector station.

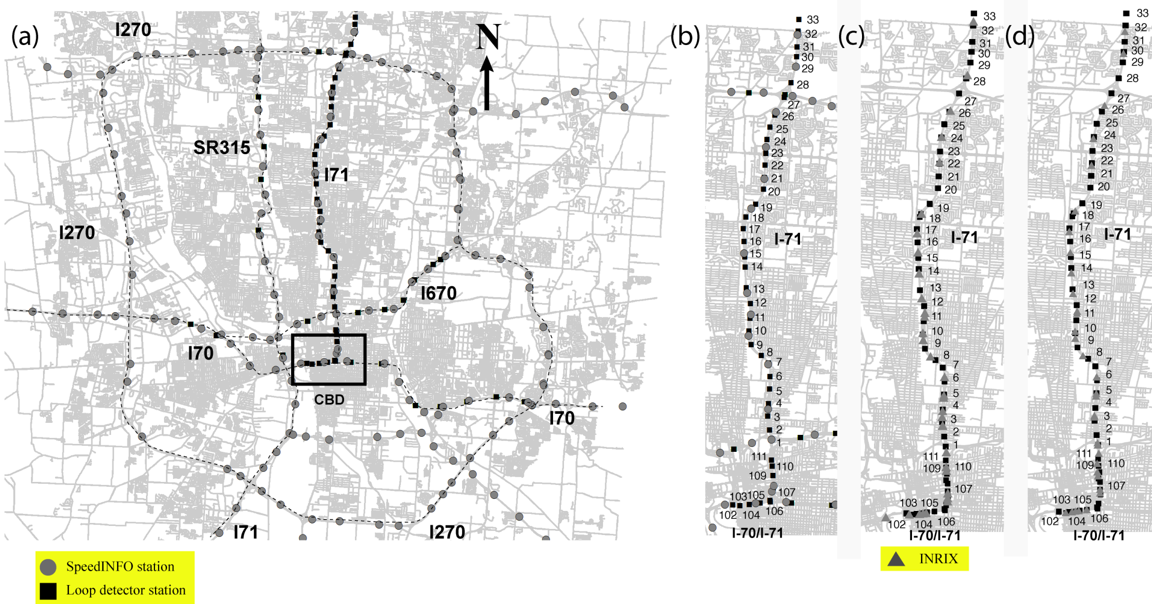

Figure 1: (a) Map of the greater Columbus, OH area showing the freeways and detector locations, (b) the I-71 corridor and SpeedInfo sensors used in this study (note, for clarity only the northbound loop detector stations are shown, while this study used both southbound and northbound stations in the evaluation), (c) the I-71 corridor and the ends of the northbound INRIX links used in this study, (d) the I-71 corridor and the ends of the southbound INRIX links used in this study.

The archived datasets investigated in this study come from Jan 2011 to May 2011 when both SpeedInfo and the loop detector system were operational. The native SpeedInfo dataset averages speed over 1 min intervals. On the other hand, speed from the loop detector system has a much finer resolution because it is based on individual vehicle actuations. As shown in Figure 1b, there are 42 northbound loop detector stations and 14 bidirectional SpeedInfo sensors along the study corridor (there are another 40 southbound loop detector stations not shown in this figure). From the map in Figure 1b, we selected the loop detector station closest to a given SpeedInfo sensor, as shown in Table 1, and we use this pairing throughout the rest of this study. The pair of SpeedInfo stations in each row comes from a single sensor[1]

[1] Excluding the two unidirectional SpeedInfo stations: 106 and 107, which only have one SpeedInfo station per row.

and the directional loop detector stations used in this evaluation were similarly co-located. So at these bidirectional stations any positional bias seen in the link would be balanced: for one direction whichever sensor (loop or SpeedInfo) was further downstream in the link would typically see a growing queue first, but this upstream/downstream relationship is flipped when observing the opposite direction. Table 1 also indicates how a given SpeedInfo sensor is mounted (M: median installation, NBS: northbound shoulder, SBS: southbound shoulder) and the heading of the sensor (SW: facing southbound or westbound, NE: facing northbound or eastbound). For brevity, we will refer to a given pairing in this table by the loop detector station number.

Table 1: Loop detector station associated with SpeedInfo

Northbound		Southbound		Installation	SpeedInfo Location
Loop	SpeedInfo	Loop	SpeedInfo		
26	12747	26	12748	M	363 ft north of loop, facing SW
24	12756	24	12755	M	1,097 ft south of loop, facing NE
21	12771	21	12770	M	203 ft south of loop, facing NE
19	12742	19	12741	M	881 ft south of loop, facing NE
16	12740	16	12739	M	1,668 ft south of loop, facing NE
13	12752	13	12751	M	740 ft south of loop, facing NE
11	12768	11	12769	NBS	450 ft north of loop, facing SW
10	12767	10	12766	SBS	716 ft south of loop, facing NE
7	12773	7	12772	SBS	194 ft north of loop, facing NE
3	12764	3	12765	NBS	941 ft north of loop, facing SW
1	12775	1	12774	NBS	0 ft from loop, facing NE
107	12759	-	-	NBS	288 ft north of loop, facing SW
-	-	106	13147	SBS	1,266 ft north of loop, facing SW
103	12661	103	12662	M	901 ft east of loop, facing SW
102	12665	102	12666	NBS	1,088 ft west of loop, facing SW

Where:

M: SpeedInfo sensor unit is located in the median

NBS: SpeedInfo sensor unit is located on the northbound shoulder

SBS: SpeedInfo sensor unit is located on the southbound shoulder

This process was repeated for the INRIX data, and Figures 1c-d show the ends of the directional INRIX links and the concurrent loop detector stations. Ultimately ODOT provided data from 20 northbound and 24 southbound INRIX links in this corridor. Unlike SpeedInfo, there are several INRIX links that span multiple loop detector stations. Most of these multiple station links are north of loop detector Station 14, and from the loop data, this portion of I-71 does not see significant recurring congestion. When an INRIX link spans multiple loop detector stations for this report we select the best of the available stations in the link, as follows:

1) Eliminate any loop detector stations that are known to have a chronic problem.
2) Give preference to dual loop detector stations over single loop detector stations.
3) Favor the station closest to the middle of the link.

However, since most of the links with multiple stations generally do not experience recurring congestion, criteria #2-#3 have little impact on the final results. The final mapping between loop detector stations and the INRIX links is shown in Tables 2 and 3, and we use this pairing throughout the rest of this study. The last column in Tables 2 and 3 presents the relative location of the loop detector station's location as a percentage of the INRIX link length (as measured from the upstream end), so, the closer [3] is to 100, the closer the loop detector station is to the downstream end of the corresponding INRIX link. The average of the ratio for southbound and northbound directions is 47 and 52, respectively. INRIX data consists of timestamp, speed, and two parameters for confidence of speed measurement, namely *confidence score* and *confidence value* (the details of these confidence statistics will be discussed in Section 4.3).

Table 2: Northbound loop detector station associated with INRIX

Loop Detector Station	Link ID	[1]: Link Distance (mi)	[2]: Loop detector location from upstream end (mi)	[3]: 100*[2]/[1]
33	122P04202	1.28	0.19	15
30	122+04202	1.39	0.60	43
27	122P04201	1.01	0.48	48
25	122+04201	0.71	0.32	45
23	122P04200	0.65	0.26	40
19	122+04200	1.47	0.32	22
17	122P04199	0.53	0.11	20
16	122+04199	0.48	0.25	51
15	122P04198	0.54	0.38	71
13	122+04198	0.65	0.20	32
12	122P04197	0.39	0.16	42
11	122P04196	0.43	0.29	68
10	122+04196	0.18	0.13	75
9	122P04195	0.48	0.24	51
7	122+04195	0.68	0.24	36
5	122P04194	0.48	0.02	4
4	122P04193	0.35	0.04	10
3	122P04192	0.21	0.07	35
1	122+04191	0.20	0.14	69
105	122+04632	0.33	0.21	63

Table 3: Southbound loop detector station associated with INRIX

Loop Detector Station	Link ID	[1]: Link Distance (mi)	[2]: Loop detector location from upstream end (mi)	[3]: 100*[2]/[3]
33	122N04201	1.44	0.45	31
30	122-04201	0.57	0.54	93
27	122N04201	1.44	1.08	75
25	122-04200	0.90	0.52	57
23	122N04200	0.69	0.33	47
19	122-04199	1.30	1.03	80
17	122N04199	0.57	0.49	86
16	122-04198	0.54	0.26	49
15	122N04198	0.52	0.13	24
13	122-04197	0.54	0.43	80
12	122N04197	0.47	0.32	67
11	122N04196	0.46	0.18	40
10	122-04195	0.18	0.06	33
9	122N04195	0.50	0.25	49
7	122-04194	0.70	0.43	61
5	122N04194	0.42	0.42	100
4	122N04193	0.35	0.32	91
3	122N04191	0.67	0.10	16
1	122N04190	0.22	0.01	6
112	122-04189	0.27	0.15	56
110	122N04188	0.35	0.08	23
106	122N04187	0.39	0.28	70
103	122-04631	0.20	0.18	90
104	122N04630	0.16	0.07	45

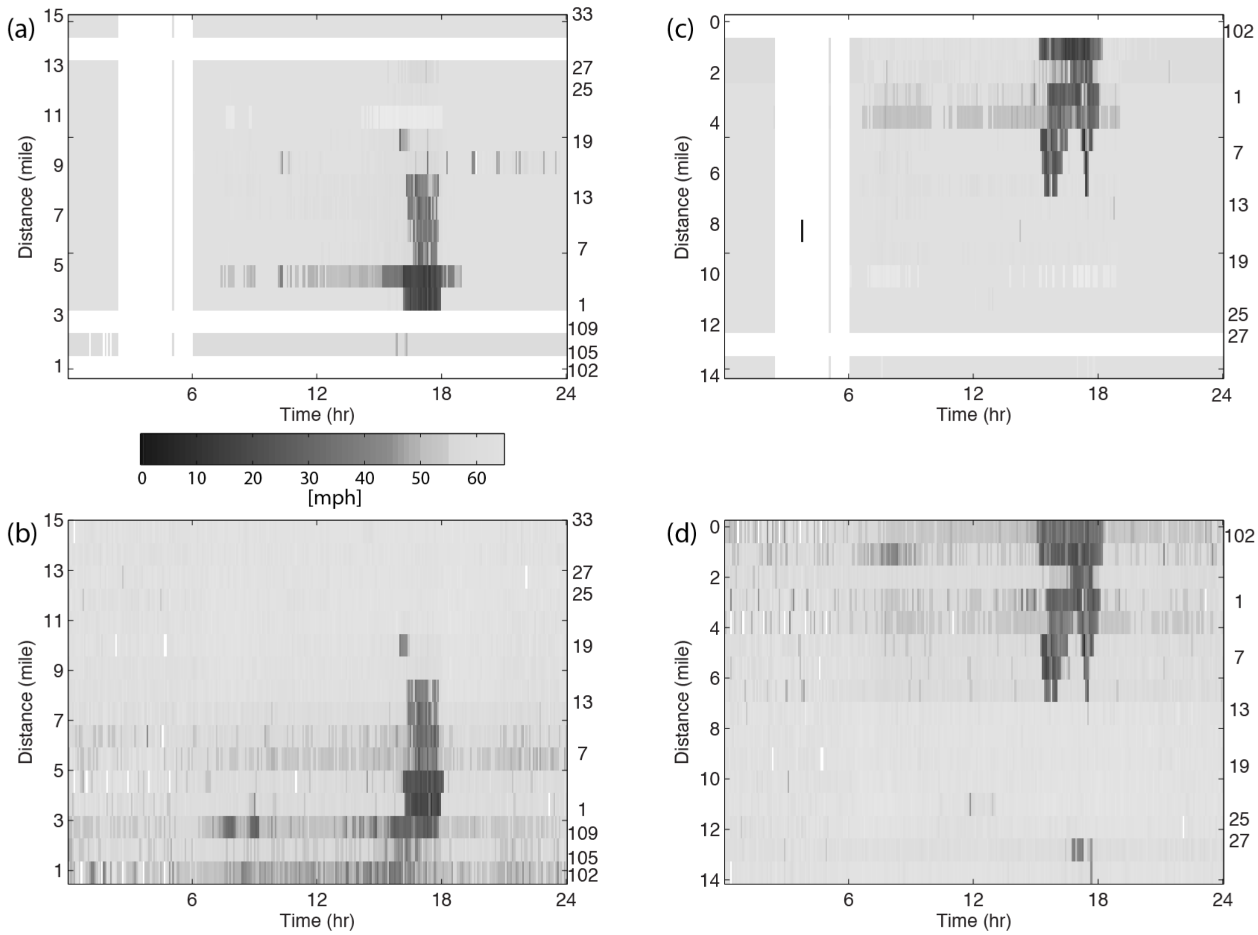

Figure 2: Matrix of 5 min median speeds along the I-71 corridor on a typical day (April 29, 2011), (a) from the northbound loop detectors (b) from the corresponding northbound SpeedInfo measurements, (c) from the southbound loop detectors, (d) from the corresponding southbound SpeedInfo measurements.

3. Evaluation of the SpeedInfo sensors

3.1 Evaluation of SpeedInfo system against loop detector system

Figure 2a shows the summary plot of the loop detector speed (5 min median speed across all lanes) over time and space from the northbound stations listed in Table 1 on a typical day. In this plot the lighter the color the higher the speed, with the exception that while denotes no data. Thus, Figure 2a shows that a heavy queue formed around mile 4.8 at approximately 16:00, grew upstream, and eventually dissipates around 18:00. Over this same period a lighter queue formed around mile 10.5, also grew upstream and overran the heavier queue around 17:00. Figure 2b shows the corresponding summary plot from the SpeedInfo sensor, using 5 min averages of SpeedInfo's native 1 min sampling period. Comparing Figure 2a with 2b, these two plots are consistent, showing queuing at roughly the same times and locations. At this resolution the overall performance from the SpeedInfo system in terms of detecting the start and end of congestion is comparable with the loop detector system. However, the measurements from the SpeedInfo system are typically slower speeds than the concurrent loop detector measurements, indicating a bias between the two systems that is discussed below. Although not evident at the current resolution of 5 min samples, we will soon see that at higher resolution the SpeedInfo sensor exhibits measureable latency. Figures 2c-d repeat the comparison for the southbound traffic on the same day, with

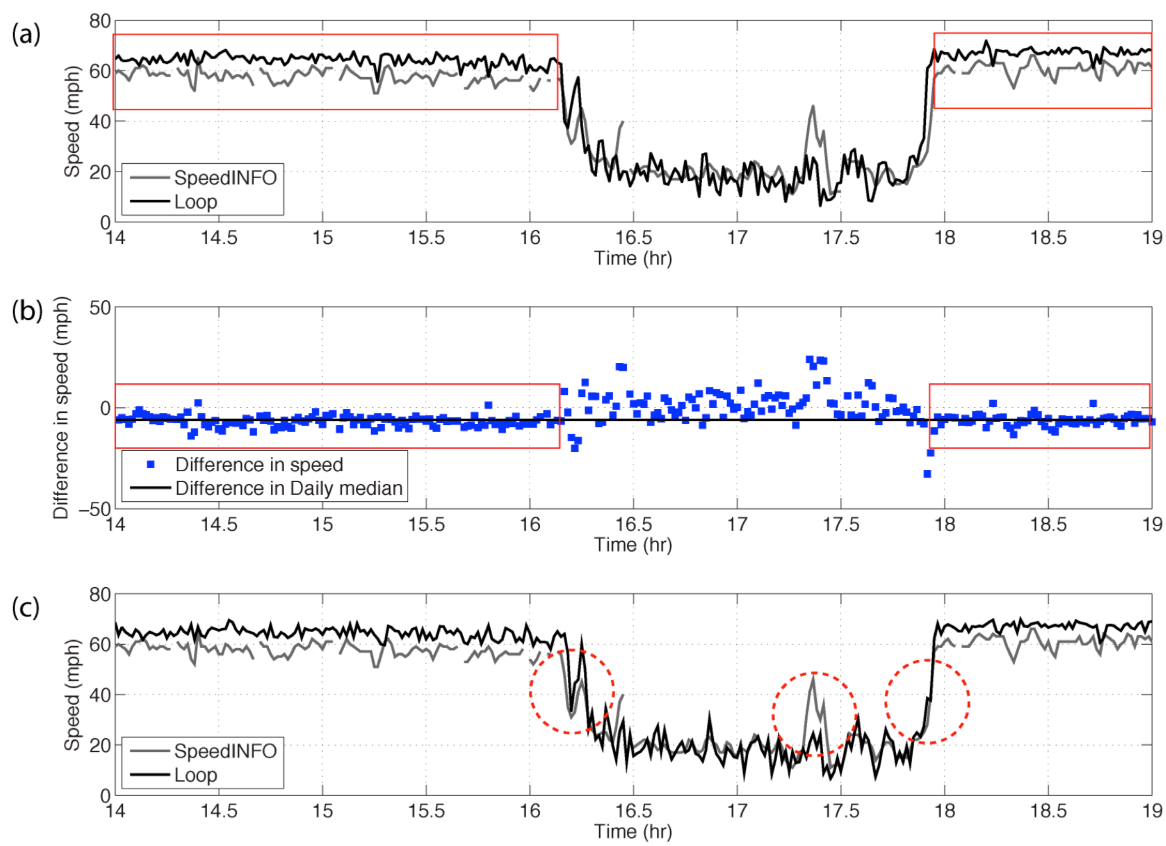

Figure 3: (a) 1 min aggregated speed as a function of time from the loop detector and SpeedInfo systems at Station 1 northbound, (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at -6 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 90 sec

similar performance. Although not shown here, we generated similar figures for days with incidents and found the performance in the presence of non-recurring congestion was similar to the recurring congestion shown in Figure 2.

To investigate any biases between the two systems, the time-series speed from a given loop detector station and corresponding SpeedInfo sensor are generated. Figure 3a shows the two time-series of speed from Station 1 northbound and the corresponding SpeedInfo sensor between 14:00 and 19:00 from the same day used in Figure 2. Now, however, the time-series speed is aggregated to 1 min, the native sampling rate of the SpeedInfo sensor. Figure 3b shows the difference between the two time-series speeds and the horizontal line at -6 mph shows the median of this difference. The square boxes highlight the free flow periods and show most individual differences are below zero, indicating a measurable bias during off-peak periods between the two systems. On the other hand, typical of the sites examined in this study, SpeedInfo tends to yield higher speeds than the concurrent loop detector during the congested periods, probably due to the fact that typical radar systems tend to favor faster moving targets (these systems typically ignore speeds below some threshold because they are too close to the background speed of zero mph) and the speed from the loop detectors is space mean speed while the speed from SpeedInfo is time mean speed (time mean speed is always equal to or greater than the space mean speed). From Figure 3b it is evident that the bias before and after the peak period is consistent with the daily median speed difference. Thus, the bias may reflect a poor calibration of one or the other sensor, and one should

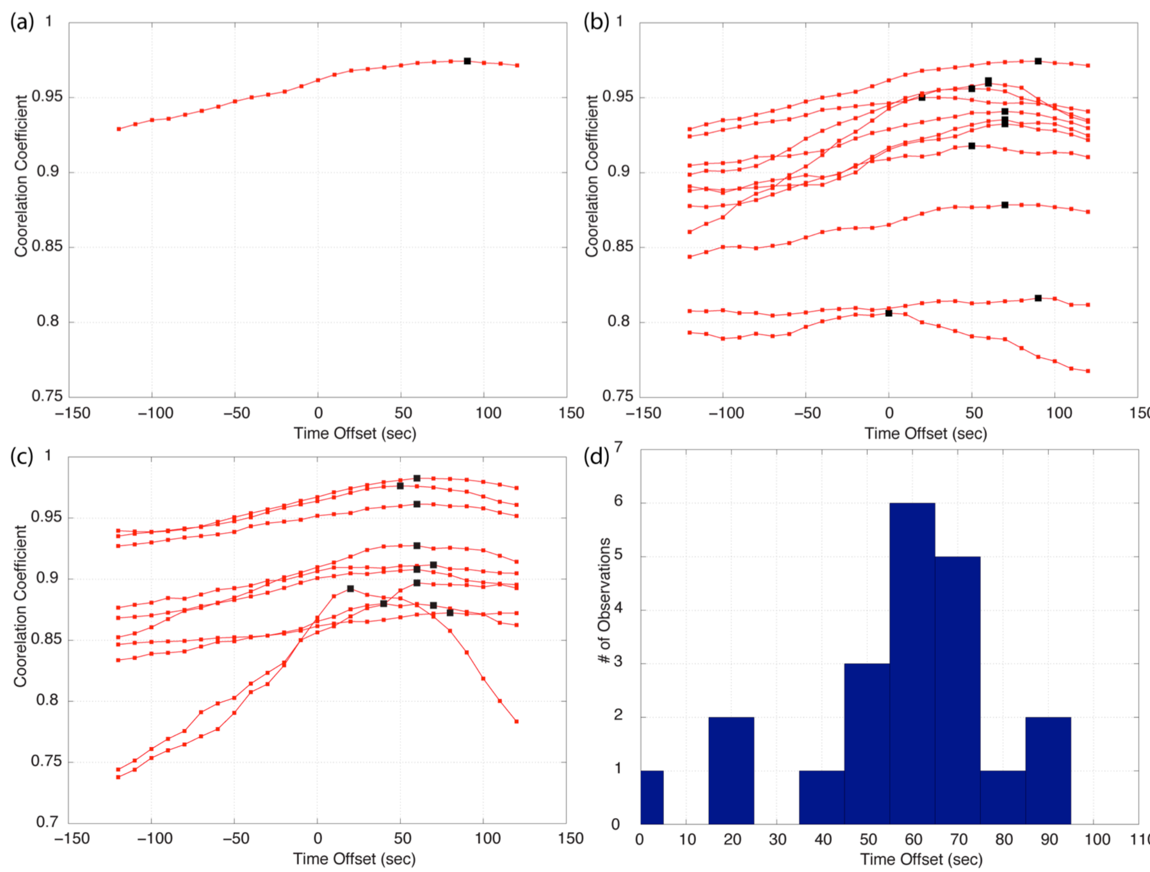

Figure 4: Correlation coefficient (CF) as a function of time offset (TO), Station 1 northbound, (a) April 29, 2011, (b) ten weekdays in April 2011, (c) ten weekdays in May 2011, (d) distribution of time offset from (b) and (c).

be able to easily correct the bias by a scale factor from either of those systems. This process is repeated in Appendix A for several stations in the study corridor with recurring congestion, with similar results.

Returning to Figure 3a, close inspection shows that the SpeedInfo measurements lag the loop detector speed by a few minutes; the clearest indication occurs when the speed drops around 16:12, and then the recovery back to free speed around 17:48. Since this queue grew upstream and receded downstream (as shown in Figure 2), whichever detector system is further downstream should respond to the onset of queuing first, and the recovery from queuing last. Given the fact that the SpeedInfo station associated with loop detector Station 1 northbound is facing northbound and the sensor is located at exactly the same location as Station 1 (see Figure 1b and Table 1), the SpeedInfo detection zone should be downstream of the loop detectors in this case, and thus, SpeedInfo should detect the growing queue around 16:12 before the loop detector. Figure 3a shows that the loop detector data responded to the onset of queuing before the SpeedInfo sensor, indicating a latency of a few minutes by the SpeedInfo sensor. Since the recovery is also seen first by the loop detectors, we can eliminate the possibility that the SpeedInfo detection zone was somehow upstream of the loop detector. This latency is within the specifications of the sensor deployment, but from an operational standpoint it is important to know about it for time sensitive applications, e.g., traffic responsive ramp metering. To measure the amount of latency, this study calculates a correlation coefficient (CF) via Equation 1, while shifting the time-series loop detector. In this study the amount of time shifted from zero is called the time offset (TO). If there is no latency, one would expect to see maximum CF when the TO is equal to zero. If there is measurable latency, the maximum CF is measured when the TO is greater than zero. The feasible range of TO used for SpeedInfo in this study is between -120 sec and 120 sec and the CF is calculated at 10 sec intervals within the range. Figure 4a shows CF as a function of TO for the case corresponding to Figure 3a. In this case the CF increases until the TO is equal to 90 sec and then decreases as the TO is increased further. In other words, when the time-series speed from the loop detector in Figure 3a is shifted by 90 sec to the right, the calculation of CF from two time-series yields its maximum value, corresponding to the best fit between the two time-series, as illustrated in Figure 3c. After shifting the loop detector data by TO both systems pick up the start and end of the congestion simultaneously, as highlighted with dashed circles in Figure 3c. To investigate the reproducibility of the latency from day-to-day, we repeat this comparison over an additional 20 weekdays exhibiting congestion from April 2011 (Figure 4b) and May 2011 (Figure 4c) using the same loop and SpeedInfo sensor pair. Figure 4d summarizes the results. Generally the SpeedInfo speed is reported later than the loop detectors, with an average latency of about 1 min (58 sec in April, 56 sec in May). As discussed earlier, since targets from the SpeedInfo station are generally downstream of the loop detector station, the true average latency should be a little longer than 1 min. This process is repeated in Appendix A for several stations in the study corridor with recurring congestion, with similar results.

$$CF(x,y) = \frac{Cov\ (x,y)}{\sqrt{Cov\ (x,x) \times Cov(y,y)}} \tag{1}$$

where,

x : Time-series speed from loop detector after shifting by TO

y : Corresponding time-series from SpeedInfo

cov (x,y) : Covariance of x and y

3.2 SpeedInfo outage issue

Because a prior study, [13], noted degraded SpeedInfo performance during the onset of rain (without independent speed measurements) and further that the sensor takes about two minutes before it begins filtering out the impacts of rain, we sought to explicitly investigate the impacts of precipitation over an extended period, i.e., the five study months. We found numerous outages lasting many minutes over the study period. We suspect that the outages are both intrinsic (e.g., power saving at night) and extrinsic (e.g., weather conditions). Strictly focusing on the SpeedInfo sensors, this section investigates when and how frequently the outages occur, and also investigates whether the outages are correlated with weather conditions.

First consider a single, sunny day: May 12, 2011, we calculate the total outage duration for every 10 min sample from each northbound SpeedInfo station in Table 1. For brevity, Figure 5a shows the median outage duration for a given 10 min sample across all of the northbound stations on this day, shown with one square for each 10 min sample period. The maximum duration of the median outages is 3 min, and this rate is steady from 22:00 to 5:00. ODOT reported that the solar powered SpeedInfo sensors are set to use a power saving mode at night and a lower reporting frequency, which is consistent with these findings. In contrast to the late night outages due to intrinsic settings, the rest of the day exhibits several median outages with 1 min duration, seemingly distributed randomly throughout the day.

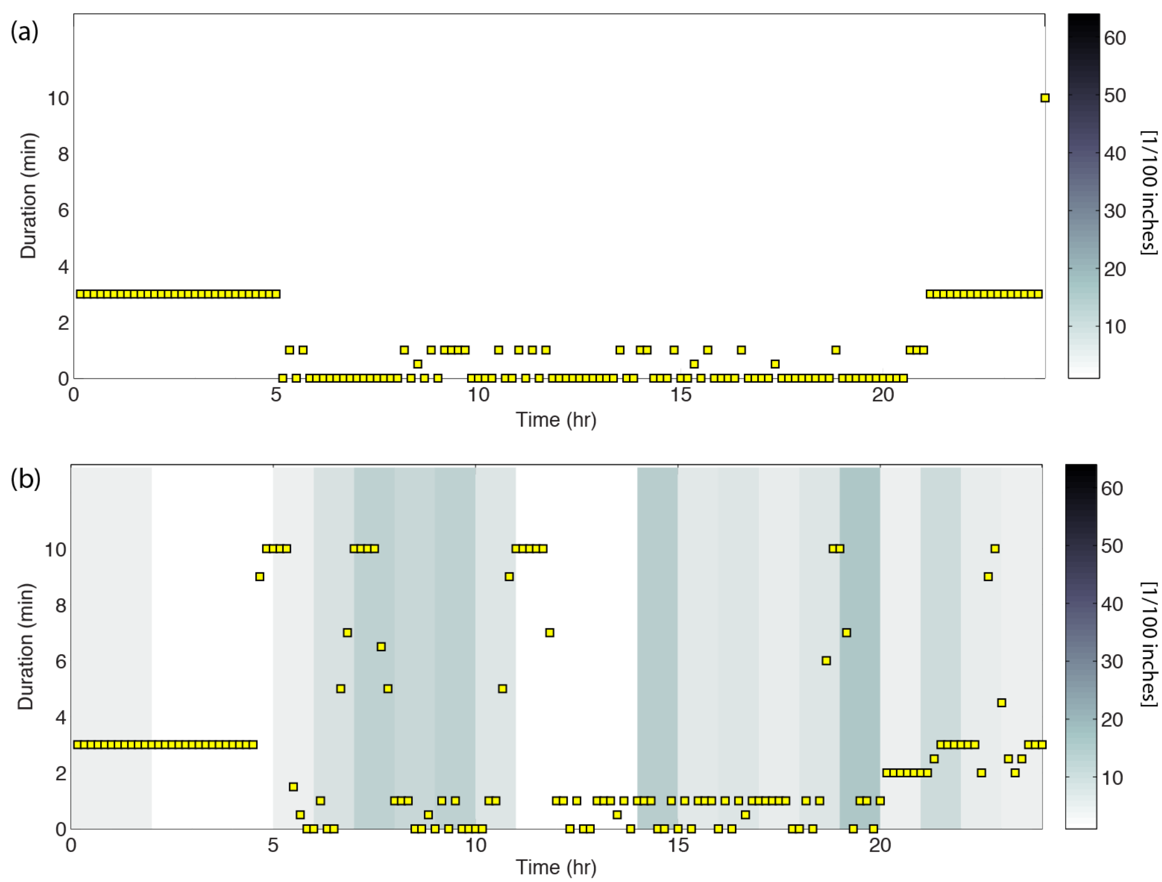

Figure 5: Median outage over a 10 min sample across all northbound SpeedInfo sensors along the I-71 corridor on (a) May 12 in 2011, a clear day, (b) May 17 in 2011, a rainy day.

Figure 5b repeats this analysis, showing the median outage duration every 10 min across the northbound SpeedInfo stations on I-71 for a rainy day. The amount of hourly precipitation is color-plotted based on the hourly summary from the Port Columbus International Airport, as reported by the National Oceanic and Atmospheric Administration (NOAA): The darker color, the heavier rain. Like the sunny day, one sees a nearly steady 3 min outage from 22:00 to 5:00. Unlike the sunny day, however, there are three clusters where the median maxes out 10 min[2] between 5:00 and 11:00 and another cluster at 19:00 (net outage lasting 20 to 50 min). On either side of these clusters are median outages of 5 min or more, capturing either the onset or recovery from the long duration outage. Comparing with Figure 5a, it appears that the 5 min or longer duration of outages is a distinctive feature that occurs when it is rainy, and that the long-term outages are correlated with weather conditions.

To check if the outages associated with precipitation are reproducible, this study examines an entire month of SpeedInfo data (May 2011). Unlike Figure 5, each station is kept separate, and for a given day we highlight every outage lasting 5 min or longer while again using color to denote the amount of precipitation. Figure 6a corresponds to Figure 5b; but, this time each individual station falls along a single horizontal line and we denote outages that last 5 min or longer with black points at the location of the given SpeedInfo station. The aforementioned four clusters of long-term outrages are evident across all of the stations (characterized by black points at multiple stations falling at the same time instant). Focusing on the start and end times of these long term outages across stations, we notice that the time when an outage starts and ends varies across stations on this day, which indicates that the outages are not simply due to communications or the central processor "going down." Figure 6b repeats the plot from Figure 6a for every day in May 2011, and presents the daily results in a calendar layout, with one plot per day. Reviewing Figure 6b, there are 5 additional rainy days where it appears that outages are associated with precipitation (i.e., May 2, 15, 16, 17, 19, 23). On the other hand, there are 6 non-rainy days (May 8, 9, 21, 24, 30, 31) where outages occurred across all stations that are not associated with any reported precipitation. In total we observed 22 long-term outages across all SpeedInfo sensors over 12 days in May 2011. Detailed investigations of each of the outages reveals that 15 outages correspond to weather conditions and show differing start and end times across different locations; 4 outages exhibit the exact same start and end times across all stations (suggesting some centralized cause); and the remaining 3 outages had the same start time but different end times. To verify a systematic correlation between outages and precipitation, one may need to study SpeedInfo stations distributed in two dimensions to follow storm fronts, rather than just the one dimensional north-south corridor used herein.

[2] Thus indicating that no data was received over the entire 10 min sample period.

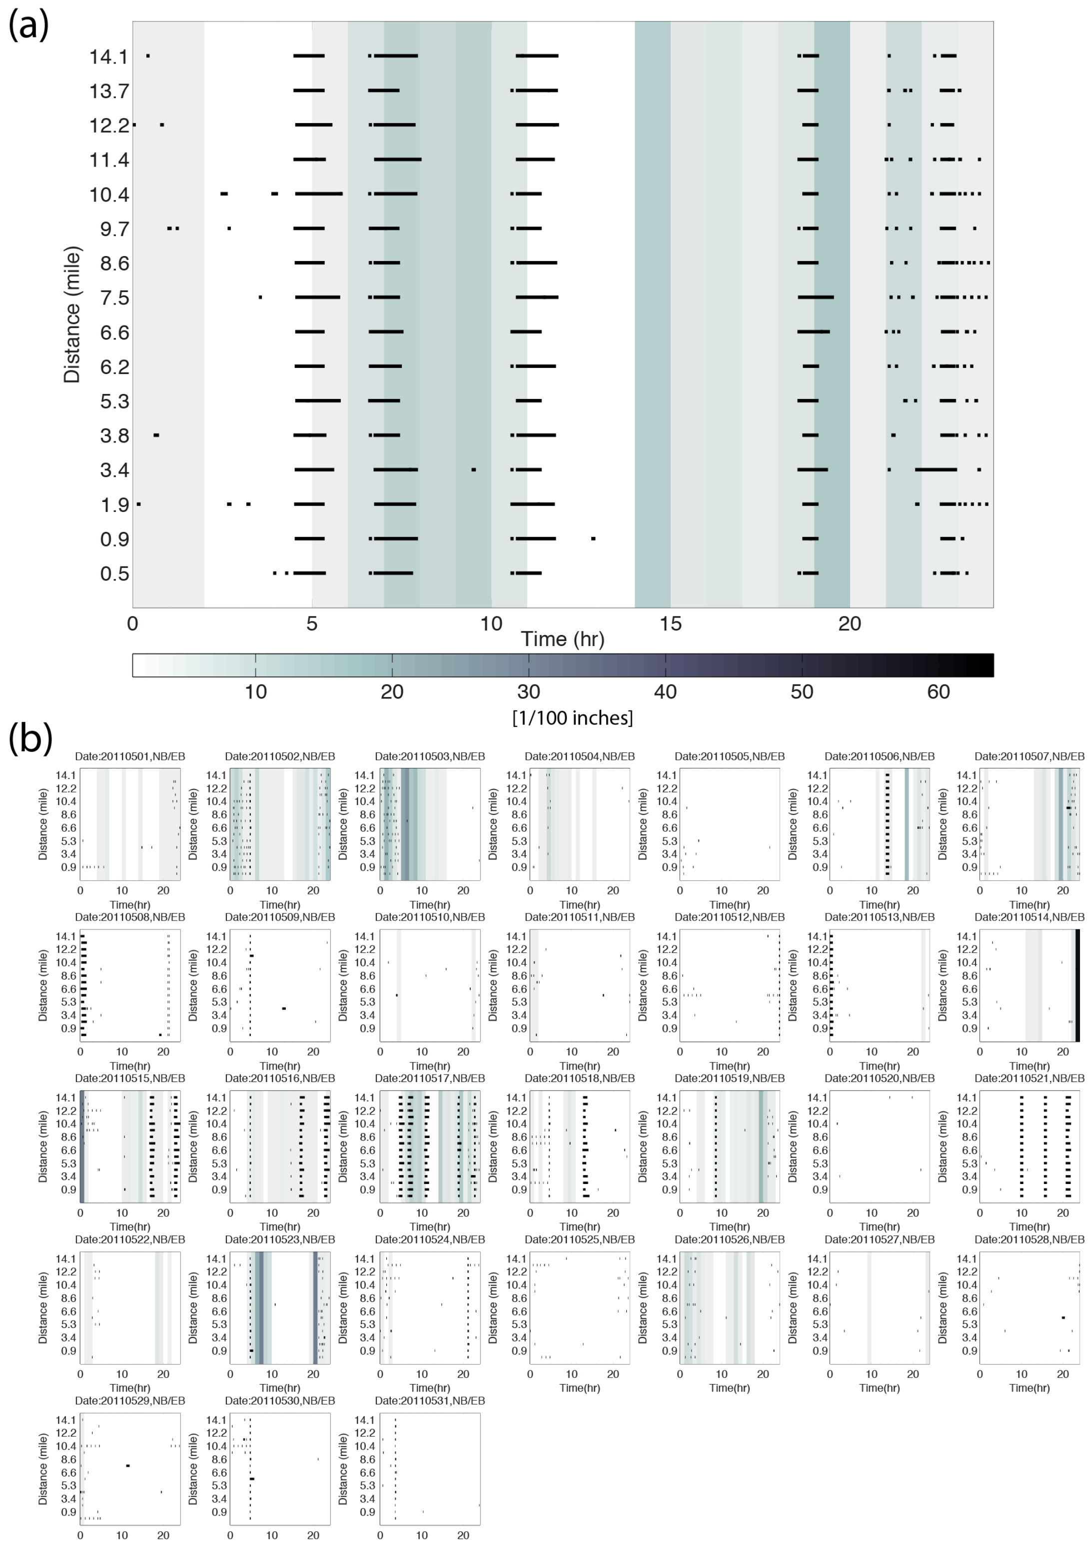

Figure 6: Outages lasting 5 min or longer shown with black points, across the individual northbound SpeedInfo stations in the I-71 corridor, (a) one day, May 17, (b) Monthly across all days in May 2011.

3.3 Precision of SpeedInfo

From our past analysis of Doppler speed radar [14-15], performance seems to degrade on the far-side of the freeway, yielding systematically low speeds due to the cosine effect being accentuated by the target traversing a larger distance as it bisects the cone of view at the further range to the target. To see if these earlier findings were also evident in the SpeedInfo data, we compared the loop detector and SpeedInfo speed measurements by direction, and looked at how performance changed depending on whether the SpeedInfo sensor was mounted on the nearside shoulder, far-side shoulder, or in the median between the two directions. As illustrated in Table 1, a given SpeedInfo station is mounted on a pole either in the median (M) or on the roadside (NBS or SBS). Excluding the unidirectional Stations 106 and 107, there are 7 median installation SpeedInfo sensors and 6 roadside SpeedInfo sensors. Figure 7 compares the monthly median speed between SpeedInfo and the corresponding loop detector station, individually in each direction. For this comparison, we take the median speed measurement for both directions from the loop detector in the median lane for the M SpeedInfo sensors, and the shoulder lane for NBS or SBS SpeedInfo sensors mounted on the roadside. So the southbound traffic is on the nearside for the SBS deployments and far-side for the NBS deployments. The abscissa in Figure 7 denotes the loop detector station and the directional monthly median speed from a given loop detector is shown with a shaded circle. The corresponding directional monthly median speed from a SpeedInfo sensor is shown in the same column with a square if the sensor is installed in the median, or a triangle if the sensor is installed on the roadside. Regardless of the sensor, the darker markers correspond to the far-side direction (denoted "Direction B") at the stations where the SpeedInfo sensor is mounted on the roadside. Otherwise, in the case of an M SpeedInfo sensor the directions are assigned arbitrarily. Finally, to ensure that the results do not depend on lane utilization, we add a smaller circle to show the monthly median loop detector speed from the other lanes (e.g., lane 1 and 2 for a roadside installation and lane 2 and 3 for a median installation).

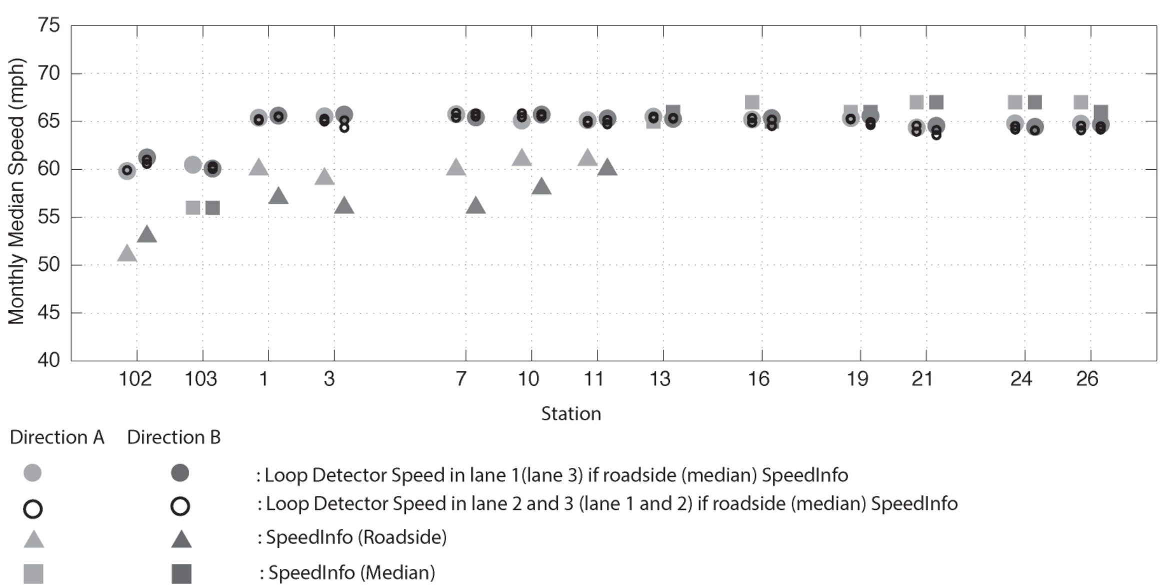

Figure 7: Directional monthly median speed comparison between a given SpeedInfo sensor and the corresponding loop detector station

Reviewing Figure 7, the directional monthly median speed from the loop detectors tends to be 5-10 mph faster than the corresponding measurement from the SpeedInfo sensors from Stations 102 to 11, consistent with the speed bias observed in Figures 2 and 3. While the two sensor systems yield similar monthly median speeds from Stations 13 to 26, potentially indicating that the bias may indeed be due to sensor calibration, and thus, potentially trivial to correct. Comparing the SpeedInfo monthly median speed from the two directions, for the roadside deployments the nearside lane measurements are faster than the far-side lane measurements five out of six times, and for these five cases the corresponding loop detector speeds do not show this nearside bias (the exception being Station 102, where both the loop detectors and SpeedInfo show the far side is faster). Like the loop detectors, the SpeedInfo sensors mounted in the median do not show any directional bias. Although we have not eliminated all possible confounding factors (e.g., sensor calibration), this comparison suggests that the SpeedInfo sensor performance depends in part on where the sensor is mounted relative to the roadway. The SpeedInfo sensors appear to offer more consistent speeds when installed in the median. For new deployments it may be worth mounting the sensor in the median whenever there is the option. However, the difference is small enough that for most applications it is not critical to move existing SpeedInfo sensors that are currently mounted on the roadside.

4. Evaluation of the INRIX data

4.1 Evaluation of INRIX data against loop detector system

Figure 8a shows the summary plot on a typical day for the loop detector speed (5 min median speed across all lanes) over time and space from the northbound stations listed in Table 2. As with Figure 2, in this plot the lighter the color the higher the speed, with the exception that while denotes no data. Thus, Figure 8a shows that a heavy queue formed around mile 4.8 at approximately 16:00, grew upstream, and eventually dissipates around 18:00. Over this same period a lighter queue formed around mile 8, also grew upstream and overran the heavier queue around 17:00. Figure 8b shows the corresponding summary plot from the INRIX data, using 5 min averages of the native 1 min reporting period used by INRIX. Comparing Figure 8a with 8b, these two plots are consistent, showing queuing at roughly the same times and locations. Although not evident at the current resolution of 5 min samples, we will soon see that at higher resolution the INRIX data exhibits measureable latency. Figures 8c-d repeat this comparison for the southbound traffic (via Table 3) on the same day, with similar performance.

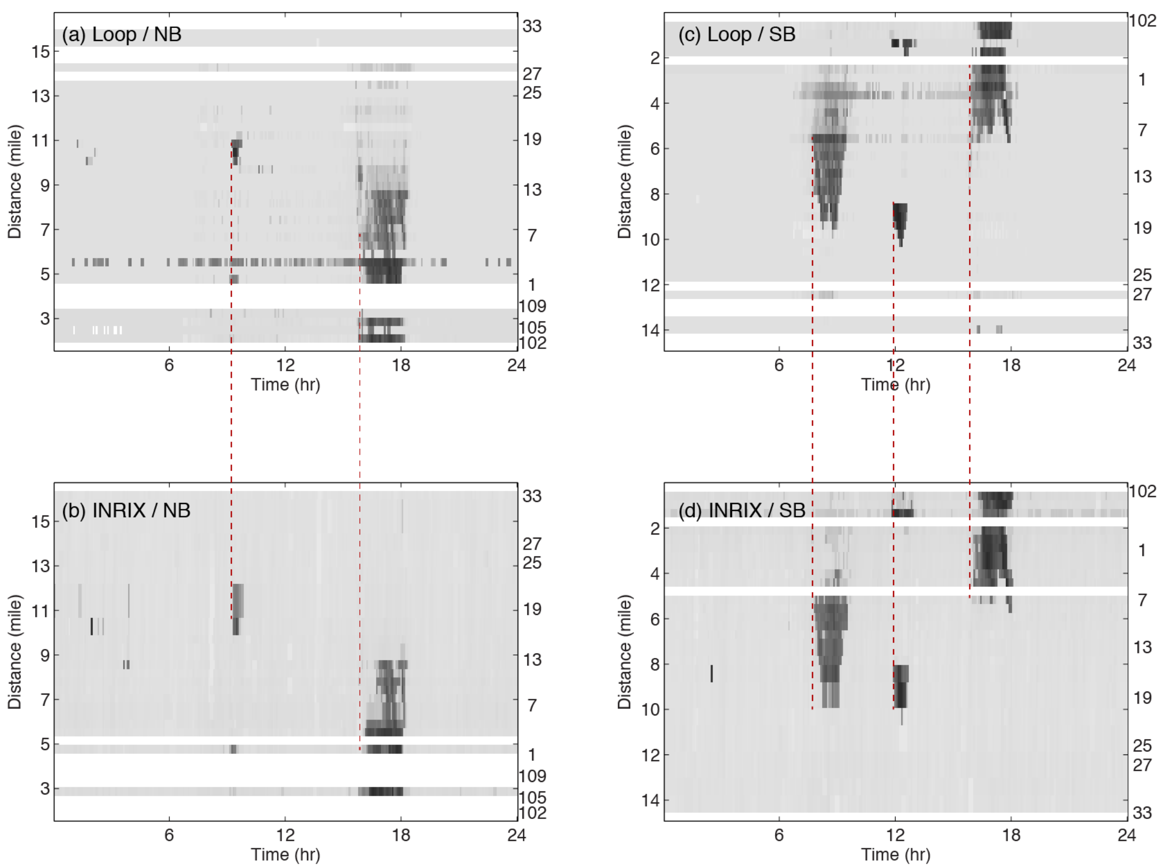

Figure 8: Matrix of 5 min median speeds along the I-71 corridor on a typical day (May 17, 2011), (a) from the northbound loop detectors (b) from the corresponding northbound INRIX measurements, (c) from the southbound loop detectors, (d) from the corresponding southbound INRIX measurements.

To investigate any biases between two systems, the time-series speed from a given loop detector station and corresponding INRIX data are generated. Figure 9a shows the two time-series speed from

Station 1 northbound and the corresponding INRIX data between 14:00 and 19:00 from May 26, 2011. Now, however, the time-series speed is aggregated to 1 min, the native reporting rate of the INRIX data. Figure 9b shows the difference between the two time-series speeds and the horizontal line at 5 mph indicates the median of this difference. The free flow periods show that most individual differences are above zero, indicating a measurable bias during off-peak periods between the two systems. From Figure 9b it is evident that the bias before and after the peak period is consistent with the daily median speed difference. Thus, the bias may reflect a poor calibration of one or the other system, and if so, it should be easy to correct the bias via a scale factor applied to either of those systems. This process is repeated in Appendix B for several stations in the study corridor with recurring congestion, with similar results.

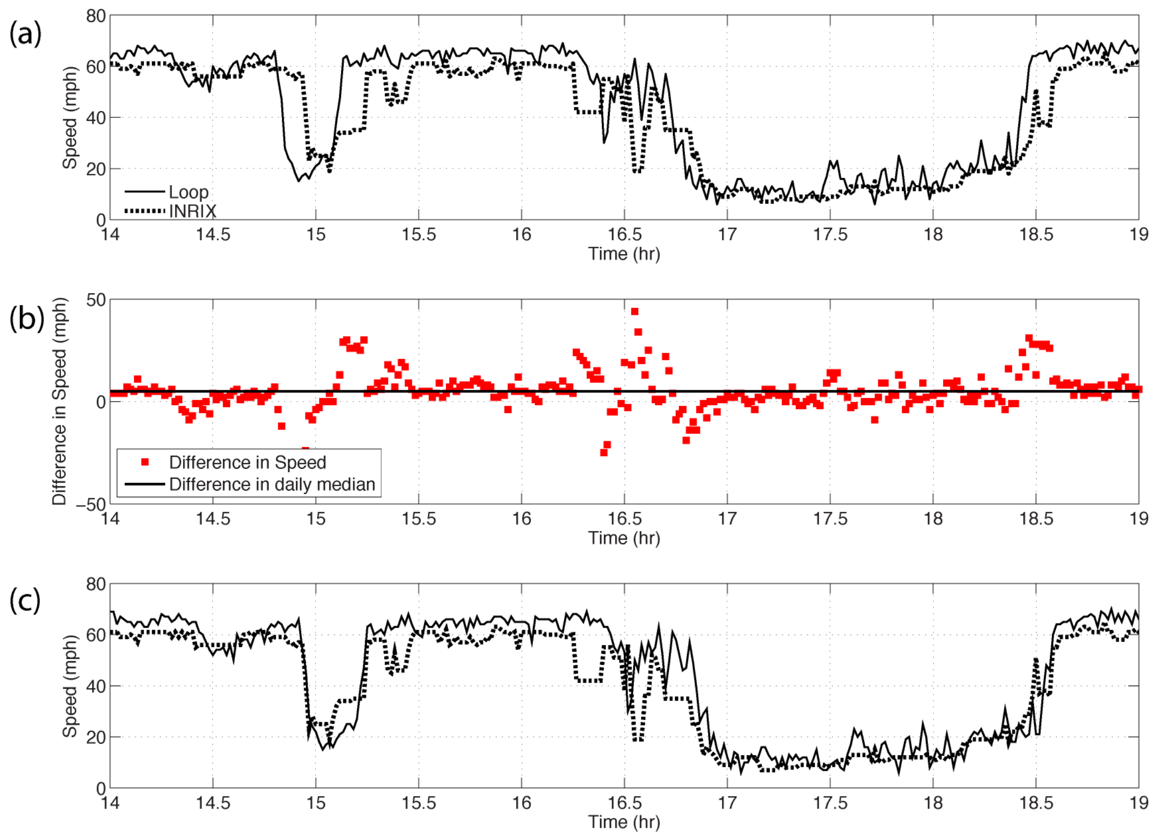

Figure 9: (a) 1 min aggregated speed as a function of time from the loop detector and INRIX systems at Station 1 northbound on a typical day (May 26, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at 5 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 420 sec.

Returning to Figure 9a, close inspection shows that the INRIX speed lags the loop detector speed by many minutes; the clearest indication occurs when the speed drops around 15:00, and then the recovery back to free speed around 18:30. As with the SpeedInfo comparison, if the plot reflects the impacts of a spatial offset, it should be unbalanced, with the downstream sensor responding to queuing first and recovering from it last. In this case the offset appears to be roughly balanced for both the onset of queuing and the subsequent dissipation. Figure 9a shows that the loop detector data responded to the onset of queuing before the INRIX data and similarly the recovery is also seen first by the loop detectors. Assuming the specifications from INRIX reported by the I-95 Corridor Coalition also apply to these data, namely that the lag time measure is less than or equal to eight minutes [16], then this latency is within the specifications of the sensor deployment, but just as with the SpeedInfo data, from an operational standpoint it is important to know about it for various time sensitive applications, e.g., traffic responsive ramp metering. Again, to measure the amount of latency, this study calculates a correlation coefficient (CF) via Equation 1, while shifting the time-series loop detector with 10 sec steps, though now we use a larger feasible range for TO. Figure 10a shows CF as a function of TO for the case corresponding to Figure 9a. In this case the CF increases until the TO is equal to 420 sec and then decreases as the TO is increased further. Figure 9c repeats the comparison from Figure 9a except that the loop detector data are now shifted to the right by 420 sec. After shifting the loop detector data by TO both systems pick up the start and end of the congestion simultaneously. To investigate the reproducibility of the latency from day-to-day, we repeat this comparison over 20 weekdays exhibiting congestion from April 2011 (Figure 10b) and May 2011 (Figure 10c) using the same loop and INRIX data pair. Figure 10d summarizes the results. Generally the INRIX speed is reported later than the loop detectors, with an average latency of about 340 (356 sec in April, 325 sec in May). The latency is longer than that observed in the SpeedInfo system in Figure 4. This process is repeated in Appendix B for several stations in the study corridor with recurring congestion.

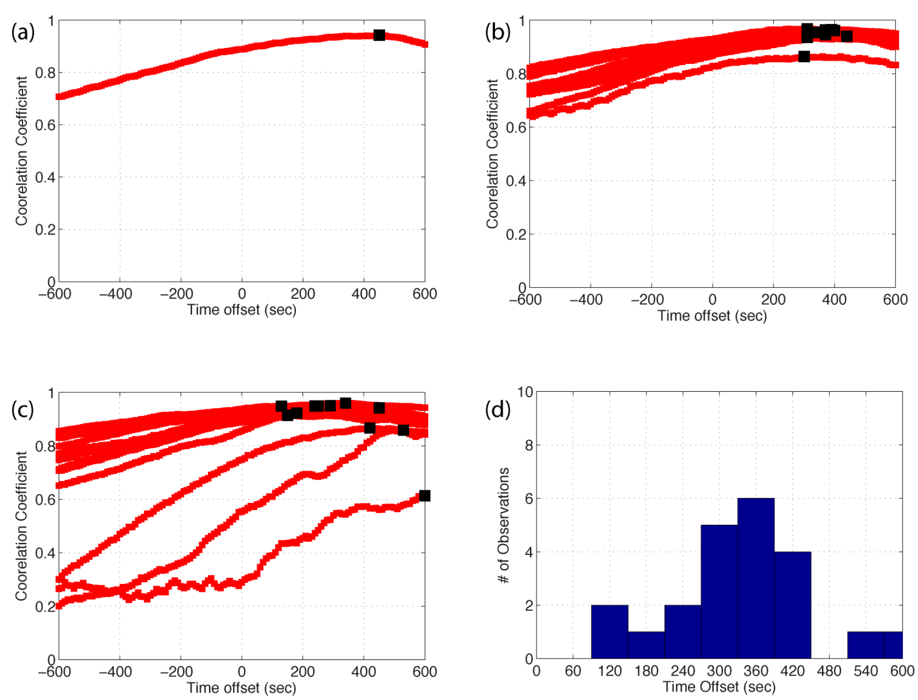

Figure 10: Correlation coefficient (CF) as a function of time offset (TO) for the INRIX data at Station 1 northbound, (a) May 26, 2011, (b) ten weekdays in April 2011, (c) ten weekdays in May 2011, (d) distribution of time offset from (b) and (c).

In contrast to Figure 9a, Figure 11 shows the two time series speeds on an atypical day. Notice that there are two speed drops evident in the INRIX data after the large recovery at 17:30 that are not present in the loop detector data. It is possible that these low speeds may come from slow moving vehicles traveling on an extended connector ramp that parallels I-71 northbound at this location; however, the further analysis in Appendix C does not support this suspicion.

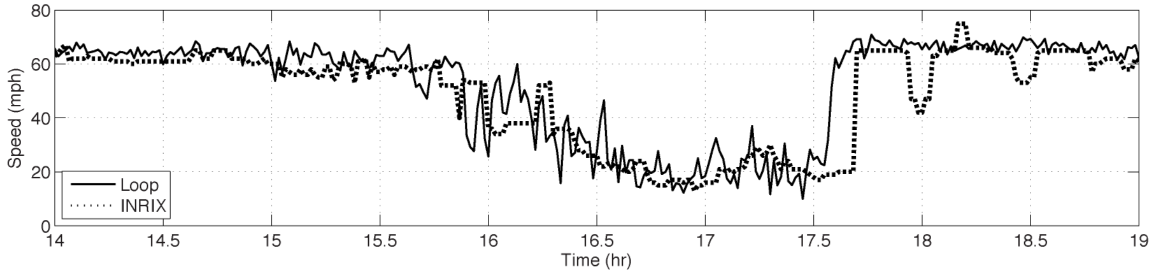

Figure 11: 1 min aggregated speed as a function of time from the loop detector and INRIX systems at Station 1 northbound on a non-typical day (April 15, 2011).

4.2 Repeated measurements reported by INRIX

Although INRIX did not exhibit weather related performance issues like SpeedInfo did, INRIX typically reports the same measurement for several successive reporting periods. Although the traffic pattern on April 15, 2011 was atypical, the data transmission was typical. Figure 12 shows that although INRIX reports measurements every minute, most of the time the reported speed is simply repeated from the previous sample (for clarity only repeated samples over 3 min are shown for the INRIX data, with the shorter durations suppressed in this figure). Figure 13 repeats this exercise on another day, except now the total number of repeated samples is shown for each instance whenever a measurement is repeated at least three times in a row.

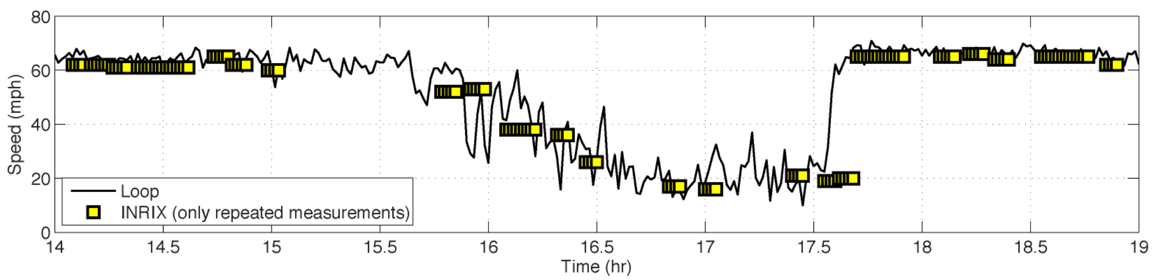

Figure 12: Replotting the data from Figure 11 this time only showing the INRIX data that repeat the value from the previous reporting period (for clarity only repeated samples over 3 min are shown for the INRIX data, with the shorter durations suppressed in this figure).

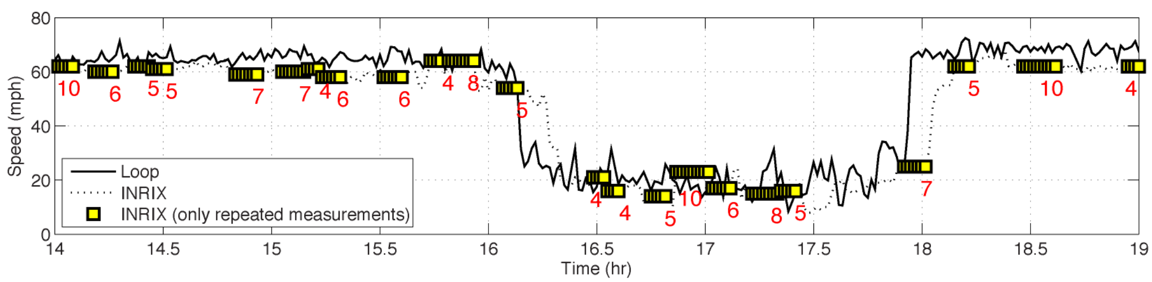

Figure 13: Time-series speed from loop detector and INRIX (April 29, 2011, Station 1 northbound) highlighting repeated speed measurements from INRIX (for clarity only repeated samples over 3 min are shown for the INRIX data, with the shorter durations suppressed in this figure).

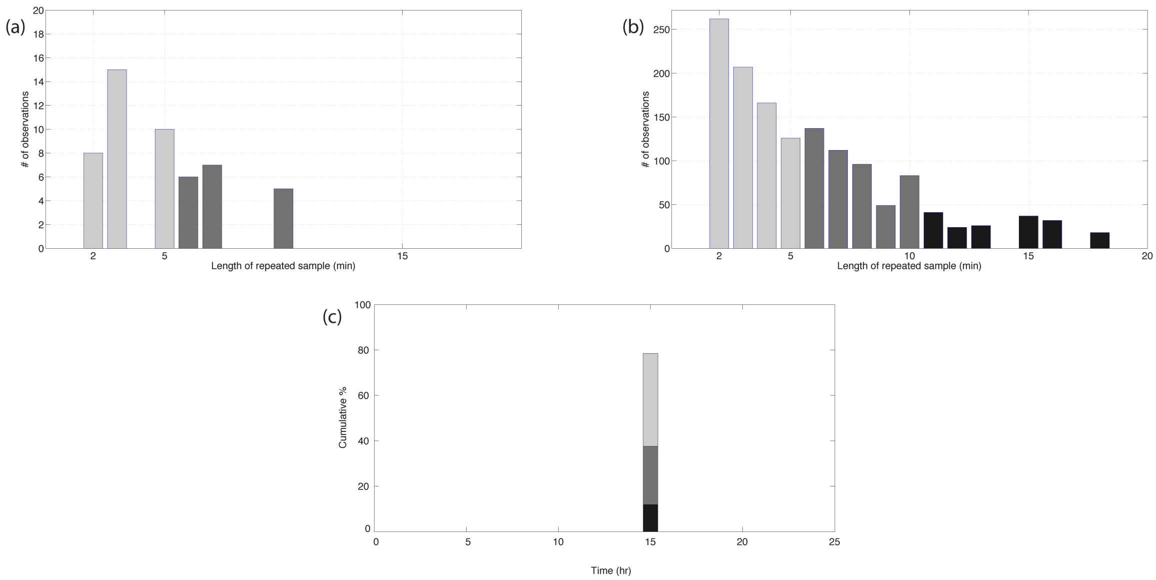

Figure 14: (a) distribution of length of time a report is repeated (by sample) between 14:00 and 15:00 on April 29, 2011 at Station 1 northbound, (b) bar chart showing the distribution between 14:00 and 15:00 from all days in April 2011, (c) binned cumulative distribution from the histogram in part (b).

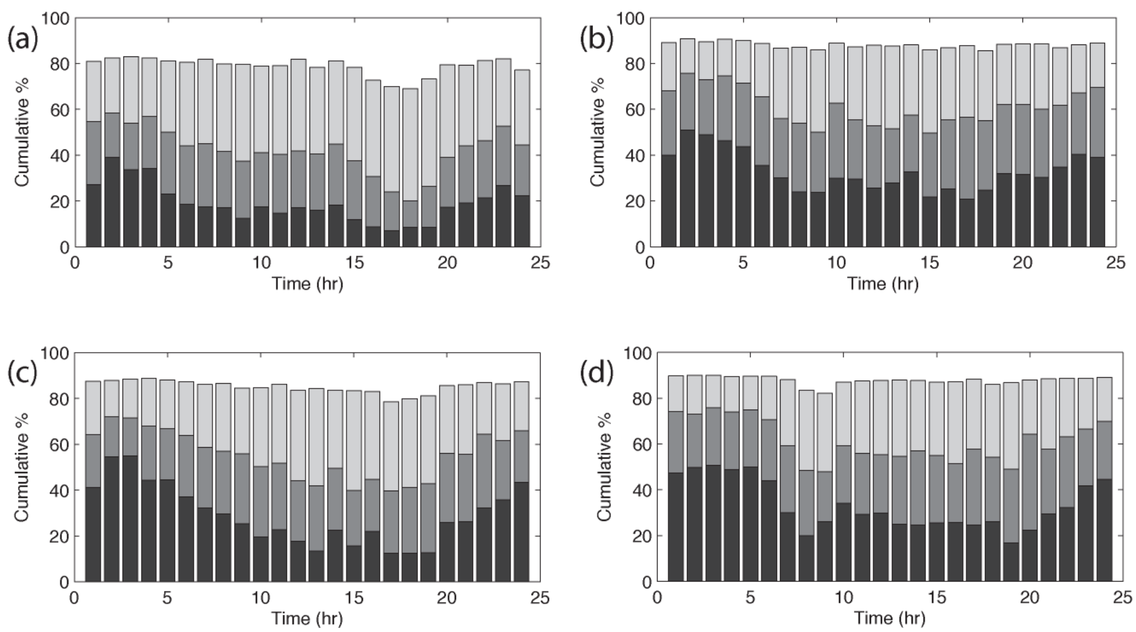

Figure 15: Binned cumulative percentage of monthly median durations with non-unique samples in April 2011 (a) Station 1 northbound, (b) Station 15 northbound, (c) Station 1 southbound, (d) Station 15 southbound.

For any directional station there should be 60 reported speeds per hour. For each sample we tally the duration of repeated samples, e.g., in Figure 13 the hour between 14:00 and 15:00 starts with 5 samples from a cluster of 10 repeated values, yielding 5 samples of 10 duration in this hour (the first 5 samples from the cluster fell in the previous hour) and repeating this tally for the entire hour results in the distribution in Figure 14a. Adding in the results from the same hour of the day for every day in the month at this station yields Figure 14b. Binning the results from Figure 14b in to: 2-5, 6-10, and 11+ and calculating the cumulative percentage yields Figure 14c. Then we repeat this process for every hour in the day at the station, yielding Figure 15a (Figure 14c corresponds to the 15th bar in Figure 15a). This performance varies by station, Figure 15b shows the corresponding results for Station 15 northbound, almost 5 miles further north. Between 8:00 and 20:00 the median repetition of a time report is in the 2-5 min bin at Station 1 and it falls in the 6-10 min bin at Station 15. Figures 15c-d shows the corresponding results from Stations 1 and 15 southbound. Figure 16 shows the results for all stations in May 2011. As noted above, the INRIX data quality has reportedly improved since 2011, so Figure 17 repeats this analysis for May 2013. Comparing a given station between 2011 and 2013 in all cases the cumulative distribution has shifted towards the shorter duration bins in 2013. Appendix D repeats this comparison for the southbound stations, with similar results.

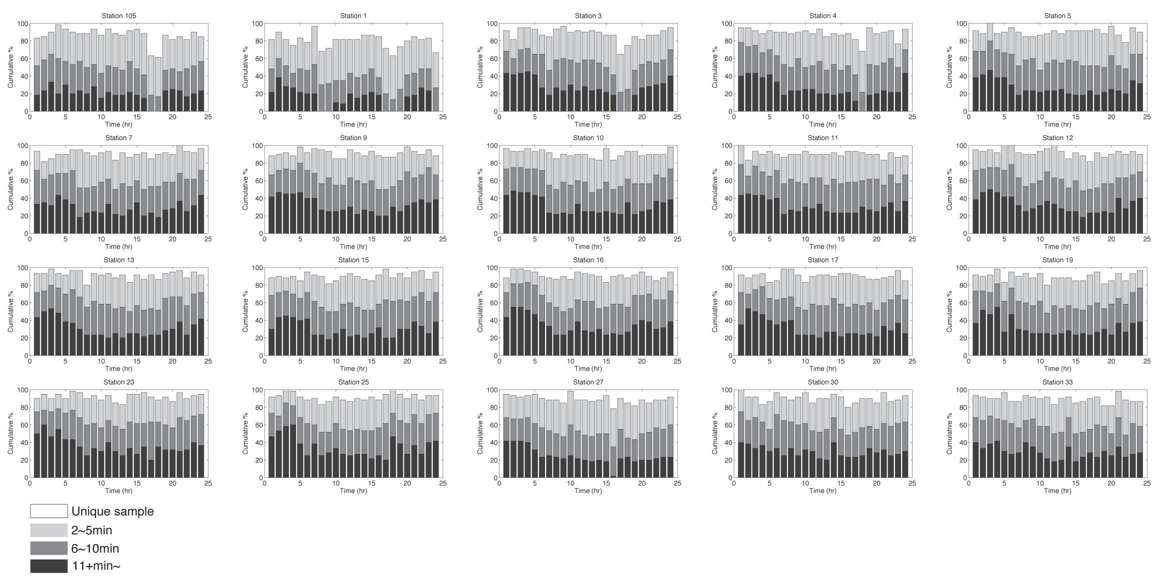

Figure 16: Binned cumulative percentage of monthly median durations with non-unique samples for each northbound station in May 2011.

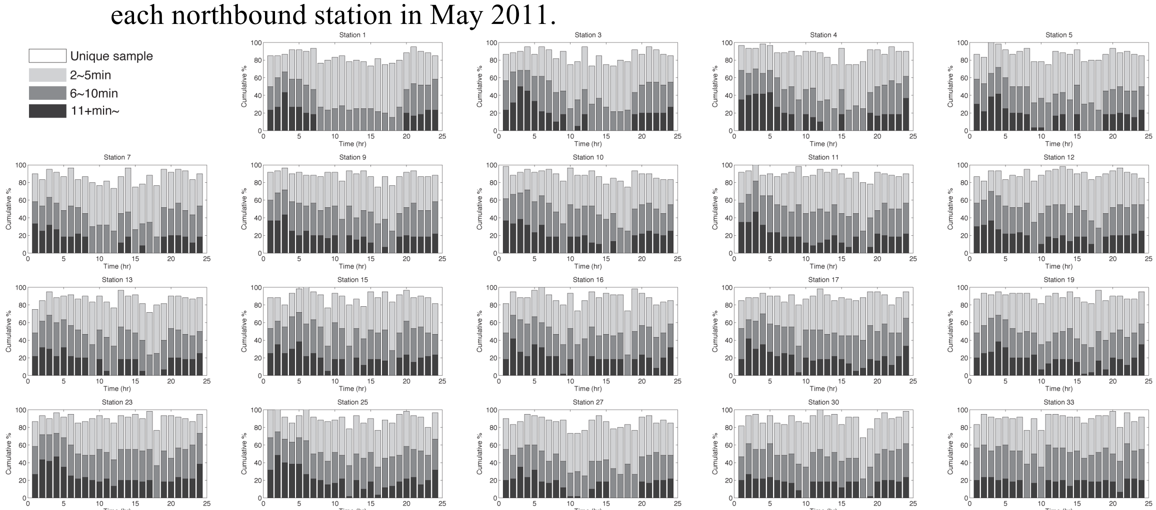

Figure 17: Binned cumulative percentage of monthly median durations with non-unique samples for each northbound station in May 2013.

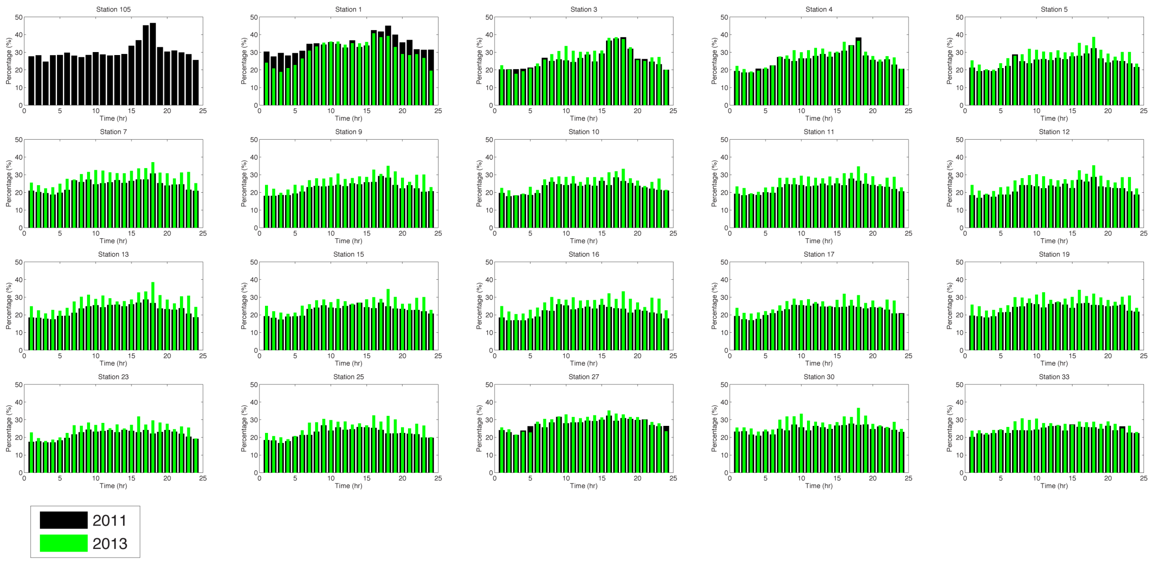

Figure 18: Percentage of samples with a change in speed by hour for all northbound stations in May 2011 and May 2013.

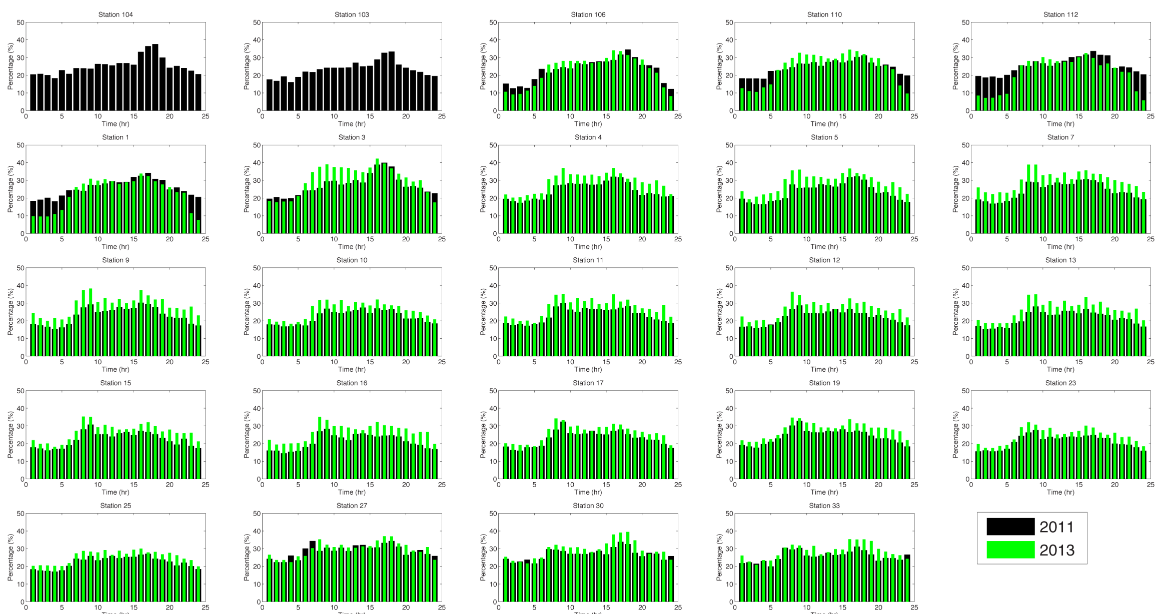

Figure 19: Percentage of samples with a change in speed by hour for all southbound stations in May 2011 and May 2013.

Another way of looking at this effective reporting period is to tally the total number of changes in speed per link, which is equal to the total number of "unique" speed measurements minus one, and calculate percentage of the samples that change every hour. So if 20% of the hourly samples are changed (12 unique samples out of 60 reported samples), it indicates that *on average* we see unique speed measurement every 5 min (equivalent to *an average* effective sampling period of 5 min during that hour), while 33% unique samples corresponds to *an average* of one unique speed measurement every 3 min. Figure 18 shows a bar graph of the hourly statistics in May 2011 contrasted against May 2013 by northbound station, and Figure 19 repeats this exercise for the southbound stations. Not surprisingly, the frequency of speed change tends to be at its lowest during the early morning and late evening, while the highest frequencies are typically seen during the morning or evening peak period. Between 2011 and 2013, except for Station 1, most stations showed roughly a 5% increase of the frequency of speed changes, suggesting that much of the improvement seen between Figure 16 in 2011 and Figure 17 in 2013 came from reducing the number of extremely long periods without a new measurement, thereby reducing the variance.

4.3 The confidence score and confidence values reported by INRIX.

INRIX reports two measures of confidence in a given speed measurement. The first is the *confidence score*, which can take one of three possible values:

30: high confidence, based on real-time data for that specific segment
20: medium confidence, based on real-time data across multiple segments and/or based on a combination of expected and real-time data
10: low confidence, based primarily on historic data

Reviewing the INRIX data from May 2011 and May 2013, we count the number of speed records associated with each of the confidence scores at each of the INRIX links in our study, as summarized in Table 4[3]. Since INRIX reports speed measurements and confidence score every minute, one would expect to find 44,640 confidence scores over a 31-day month. However, in May 2013 INRIX occasionally did not report a confidence score (or the archive extraction failed to retrieve it) for a given minute. We found no missing records in May 2011 and 151 missing records in May 2013, as enumerated by day in Table 5. Note that all stations exhibited the same pattern, if a given record was missing from one station it was missing from all stations, so the results are shown by day rather than by station in this table. The number of missing records is very small and again, may be due to the archive extraction. More importantly, Table 4 shows that over 94% of the samples have the highest confidence score. Figures 20-21 show the cumulative distribution of confidence score by time of day (minute) across all 31 days in the month, at a given station. So each of the 1,440 columns in a given subplot shows the frequency of the confidence scores for a specific minute of the day, across the 31 days of the month. Even during the early morning hours every station typically has a score of 30 at least 80% of the time. Compared to 2011, the frequency of scores of 10 and 20 drop in 2013. Appendix E shows the corresponding figures for the southbound stations.

Figure 22a revisits the data from Figure 9, only now highlighting the repeated values reported by INRIX. Figure 22b shows the corresponding confidence score reported by INRIX for all of the samples in this period, again highlighting those records corresponding to repeated values from INRIX. This example is typical of other days and stations: the repeated speed values usually have the highest confidence score, 30. In fact over the entire period the confidence score never deviated from 30. As such, the confidence score does not reflect how long a given measurement has been reported.

---

[3] The confidence scores were extracted from the INRIX archives at different times, and as a result, we did not collect the exact same stations in the two years. For the sake of comparison, this section only evaluates the stations for which we had data in both 2011 and 2013.

Table 4: Total number of speed records reported for a given confidence score (CS) by northbound station.

Station	May 2011				May 2013			
	CS:10	CS:20	CS:30	Total	CS:10	CS:20	CS:30	Total
St.1	2,679 (6.0%)	954 (2.1%)	41,007 (91.9%)	44,640 (100%)	2,821 (6.3%)	150 (0.3%)	41,518 (93.0%)	44,489 (99.6%)
St.3	2,006 (4.5%)	746 (1.7%)	41,888 (93.8%)		1,673 (3.7%)	190 (0.4%)	42,626 (95.5%)	
St.4	2,204 (4.9%)	828 (1.9%)	41,608 (93.2%)		1,620 (3.6%)	225 (0.5%)	42,644 (95.5%)	
St.5	1,827 (4.1%)	700 (1.6%)	42,113 (94.3%)		1,481 (3.3%)	128 (0.3%)	42,880 (96.1%)	
St.7	1,686 (3.8%)	656 (1.5%)	42,298 (94.8%)		1,262 (2.8%)	92 (0.2%)	43,135 (96.6%)	
St.9	1,651 (3.7%)	705 (1.6%)	42,284 (94.7%)		1,089 (2.4%)	110 (0.2%)	43,290 (97.0%)	
St.10	1,730 (3.9%)	740 (1.7%)	42,170 (94.5%)		1,020 (2.3%)	142 (0.3%)	43,327 (97.1%)	
St.11	1,739 (3.9%)	766 (1.7%)	42,135 (94.4%)		1,138 (2.5%)	158 (0.4%)	43,193 (96.8%)	
St.12	1,712 (3.8%)	812 (1.8%)	42,116 (94.3%)		1,242 (2.8%)	164 (0.4%)	43,083 (96.5%)	

Table 5: Number of missing confidence score records in May 2013 by day, all stations exhibited the same time series of either reported or missing scores.

Date	Stations	# of missing records
5	Across all feasible stations for both year 2011 and 2013 (Stations in Table 4)	48
12		12
14		9
22		2
23		2
24		38
25		8
27		1
28		6
29		8
30		16
31		1
Total per station		151

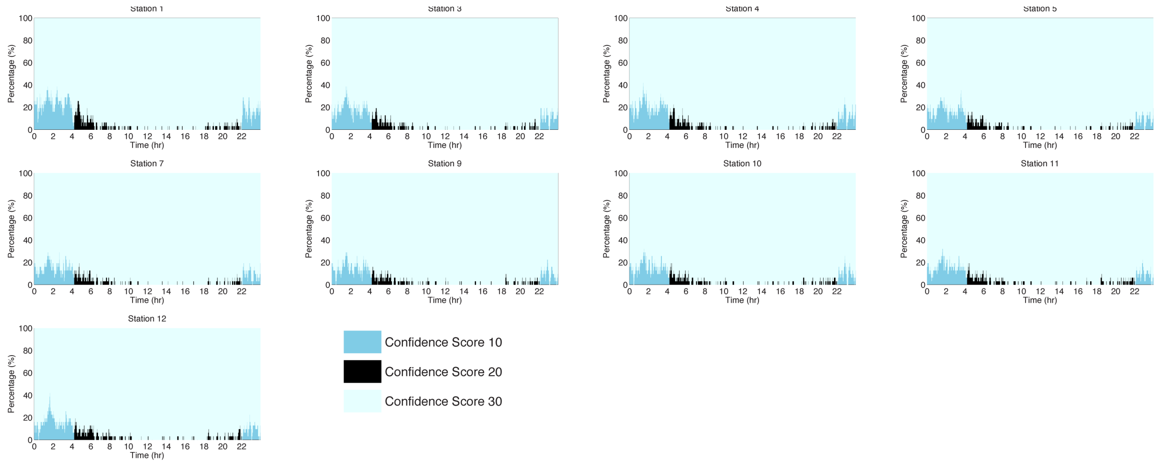

Figure 20: Percentage of samples with a given confidence score as a function of time over the entire month of May 2011 from the northbound stations.

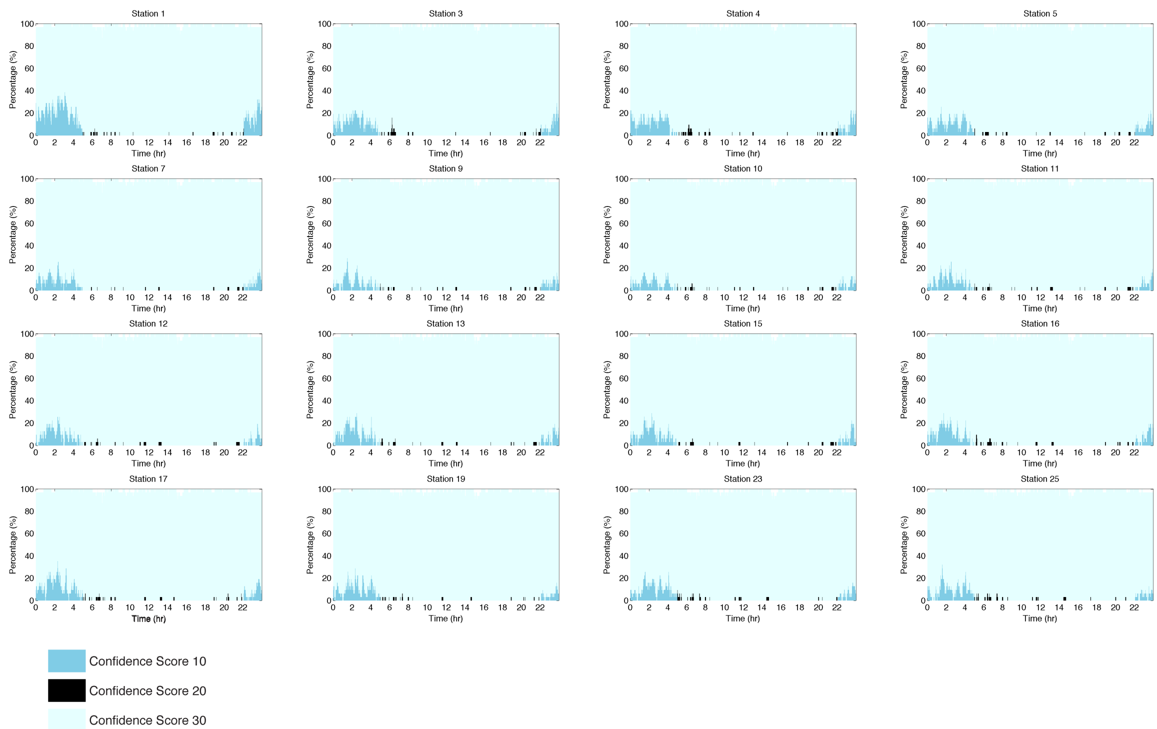

Figure 21: Percentage of samples with a given confidence score as a function of time over the entire month of May 2013 from the northbound stations.

Since 2011 INRIX added a second measure of confidence, the *confidence value.* Reportedly, the confidence value is based on a comparison against historical trends, the details of this evaluation are proprietary. The confidence value only applies when the confidence score is 30. Because the confidence value was introduced since 2011, we are limited to using the 2013 data, which do not have any corresponding loop detector measurements.

Figure 23a-b repeats the analysis from Figure 22 for May 2, 2013. Again Figure 23b shows that the repeated speed values have the highest confidence score, 30. In fact over the entire period the confidence score never deviated from 30. Figure 23c then repeats this analysis for the confidence value.

Similar to the confidence score, we find the repeated speed values usually have the highest confidence value, 100, though a few exceptions are evident, as shown in detail in Figure 24. Once more this example is typical of other days and stations. Here we see that when a repeated speed value starts out with a lower confidence value, the subsequent repetitions usually have progressively higher confidence values (though exceptions exist, e.g., the trend at 16.8 is decreasing). The exact nature of the confidence value measurement is proprietary, but like the confidence score, it appears that the confidence value does not reflect how long a given speed measurement has been reported.

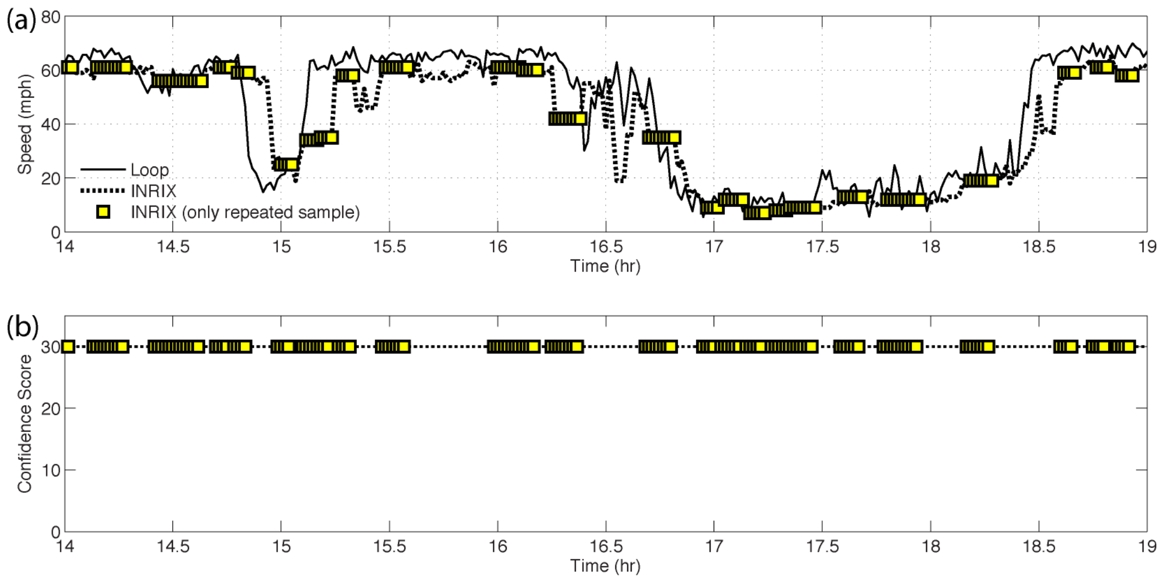

Figure 22: (a) The 1 min reported speed as a function of time from the loop detector and INRIX system at Station 1 northbound (May 26, 2011) highlighting the repeated speed measurements from INRIX (for clarity only repeated samples over 3 min are shown for the INRIX data, with the shorter durations suppressed in this figure), (b) the corresponding INRIX confidence score.

Investigating more groups of repeated samples with respect to confidence value, we found 8,943 groups with repeated speed samples over May 2013 at Station 1. The total number of speed measurements from these groups is 40,468, which is about 91% of the speed measurements reported for this station in the month (as per Table 4). Of the 8,943 groups, 7,099 had a constant confidence value where 83.45% are at 100. Only 1,844 consist of multiple confidence values. In 1,370 groups (74% of those with multiple confidence values) the confidence values increase monotonically. In 213 groups (11% of those with multiple confidence values) the confidence values decrease monotonically, while the remaining 15% of the groups show mixed trends.

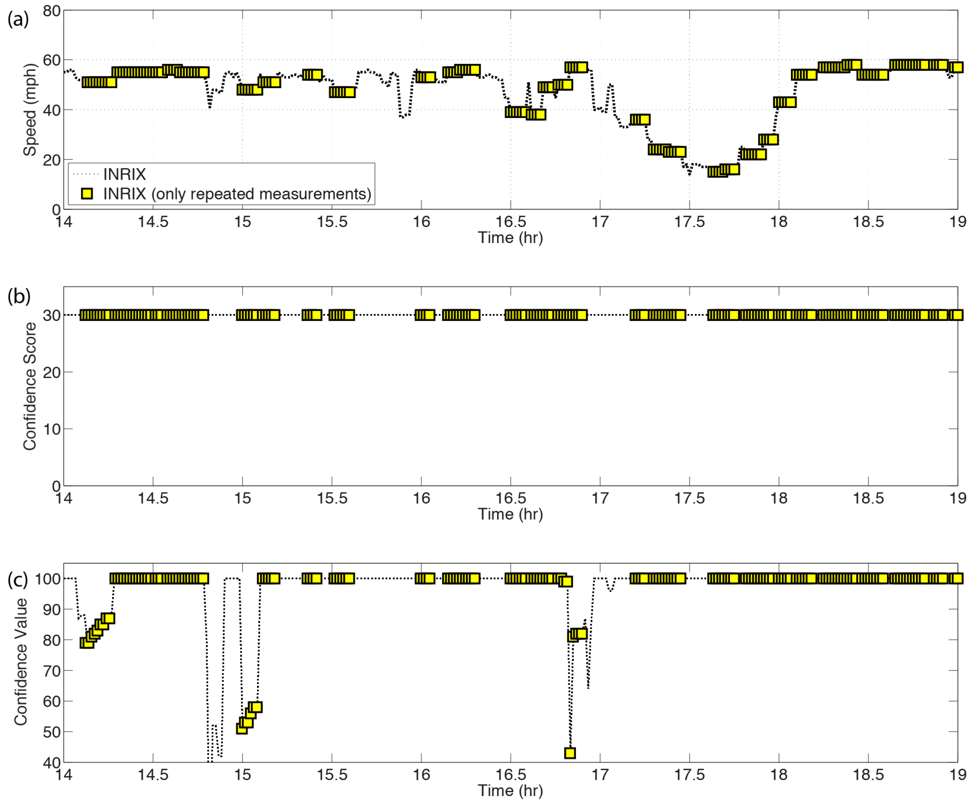

Figure 23: (a) The 1 min reported speed as a function of time from the INRIX system at Station 1 northbound (May 2, 2013) highlighting the repeated speed measurements from INRIX (for clarity only repeated samples over 3 min are shown for the INRIX data, with the shorter durations suppressed in this figure) (b) the corresponding INRIX confidence score, (c) the corresponding INRIX confidence value.

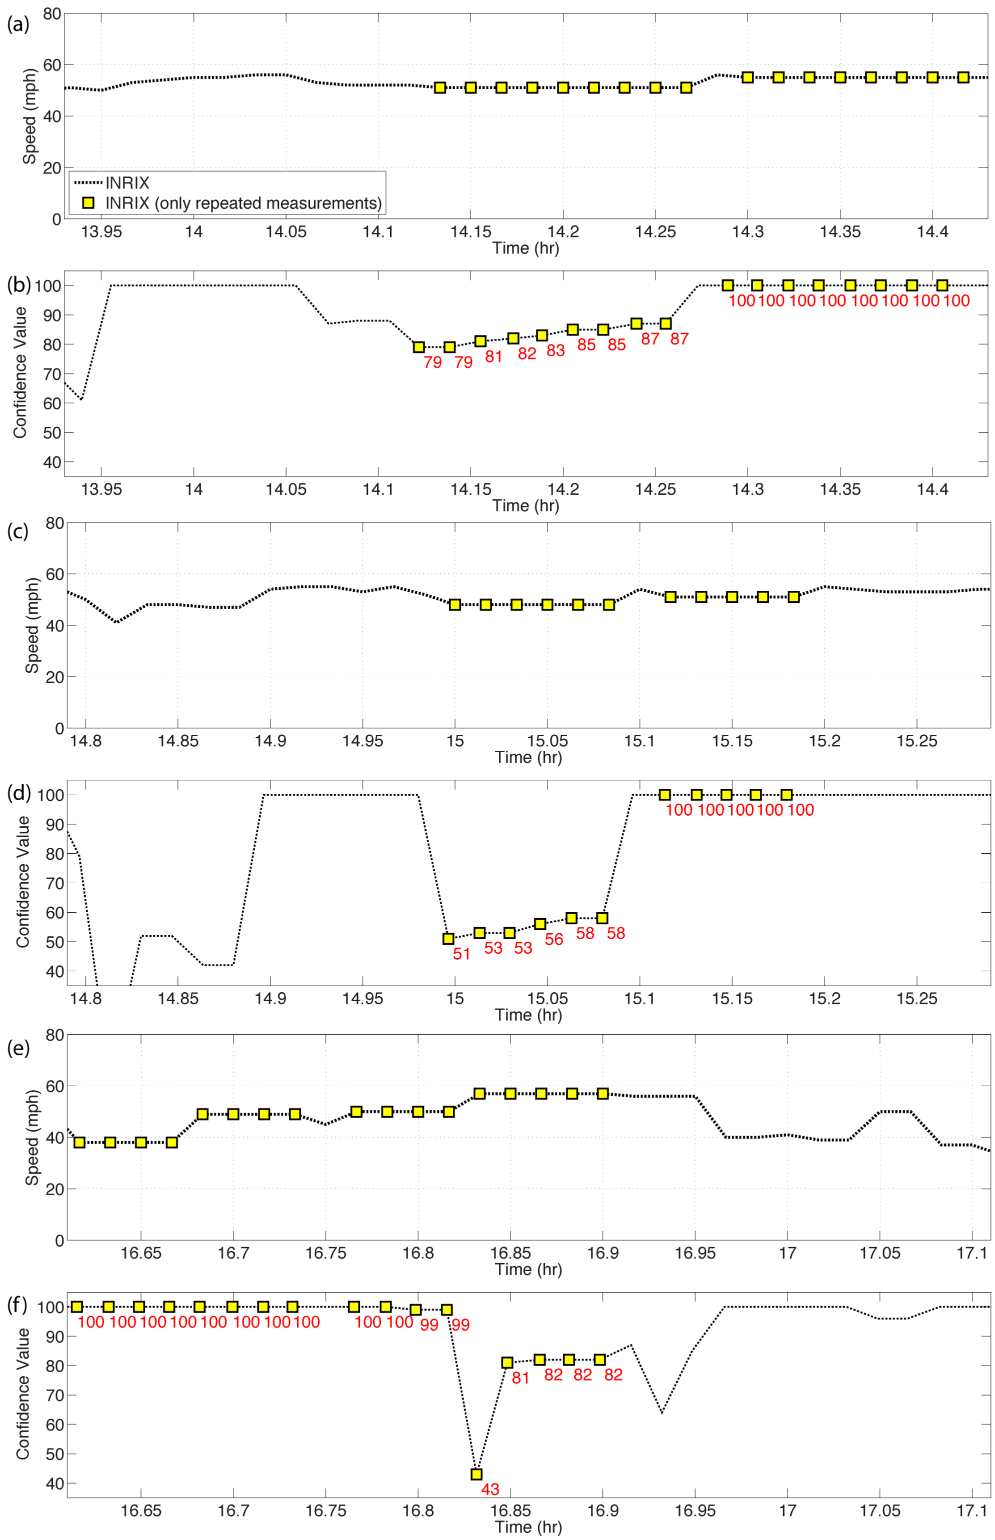

Figure 24: Details of from part (a) and (c) in Figure 23, (a-b) about 14:10, (c-d) about 15:00, (e-f) about 16:50.

5. Conclusions and Recommendations

Many real-time traffic-monitoring applications only require speed or travel time, e.g., traveler information or detecting the presence congestion. In recent years SpeedInfo has emerged with a different paradigm for traffic operations data collection: the company deploys its own Doppler radar sensors to measure aggregate speed and then sells the data. Unlike most prior traffic monitoring technologies, SpeedInfo only measures speed. From an operations standpoint the SpeedInfo model can be compelling compared to traditional traffic detection maintained by the operating agency and ODOT has deployed SpeedInfo statewide. Shortly after SpeedInfo arrived on the market, INRIX started collecting and selling real-time speed data collected from "a variety of sources," and believed to primarily come from probe vehicles in Ohio. ODOT has purchased the INRIX data to monitor otherwise non-instrumented highways in Ohio. From ODOT's standpoint, the two sources are functionally equivalent, representing third party, real-time data collected on ODOT facilities.

To date there has not been a comprehensive evaluation of the SpeedInfo sensor performance and only limited evaluations of the INRIX data, so this study set out to address those omissions and examined the performance of the emerging real-time technologies. For much of the analysis we used individual vehicle actuations from well tuned loop detectors as the baseline. Unlike prior studies, this work examined an extended period, with five months of contiguous data from concurrent detectors, allowing us to catch intermittent events as well as behavior under normal operation.

At the coarse level of 5 min aggregated speed data plotted in the time space plane, the SpeedInfo summary plots were comparable to the corresponding loop detector plots, both showing similar patterns of congestion, queue growth, and so forth. Upon examining the time series speed data aggregated at the SpeedInfo's native 1 min aggregation, two issues became apparent. First, the SpeedInfo sensors tended to measure speeds slower than the loop detectors, by about 5 mph during free flow periods. Although we have not eliminated all possible confounding factors, we found some evidence suggesting this bias might be correctable through further calibration. Second, the SpeedInfo measurements tended to lag the loop detector measurements by approximately 1 min. This latency is within the specifications of the sensor deployment, but from an operational standpoint it is important to know about it for time sensitive applications, e.g., traffic responsive ramp metering. Reviewing the periods when the SpeedInfo detectors were operational in May 2011 we observed 22 long-term outages (all in excess of 10 min, a few lasting several hours) across all SpeedInfo sensors over 12 days. Upon investigation, 68% of these outages were correlated with precipitation, though not all of the precipitation events triggered outages in the SpeedInfo data. Finally, we examined the performance differences between mounting the SpeedInfo sensors in the median, and on the shoulder. Indeed, the SpeedInfo sensor performance depends in part on where the sensor is mounted relative to the roadway. The SpeedInfo sensors offer more consistent speeds when installed in the median. Whenever there is the option, for new deployments it may be worth mounting the sensor in the median. However, the difference is small enough that for most applications it is probably not critical to move existing SpeedInfo sensors from the roadside to the median.

Following a similar process for INRIX, again at the coarse level of 5 min aggregated speed data plotted in the time space plane, the INRIX summary plots were comparable to the corresponding loop detector plots, both showing similar patterns of congestion, queue growth, and so forth. Upon examining the time series speed data aggregated at the INRIX's native 1 min reporting period, three issues became apparent. First, the INRIX speeds tend to lag the loop detector measurements by almost 6 min. Again, this latency appears to be within the specifications of the sensor deployment, but from an operational standpoint it is important to know about it for time sensitive applications. This latency is larger than observed in the SpeedInfo data; however, the bias relative to the loop detector speeds in the actual measurements appears to be smaller in INRIX. Second, although INRIX reports speed every minute, most of the time the reported speed is identical to the previous sample. After tracking the changes in reported speed, we observed an effective average sampling period of 3-5 min, with many periods of repeated measurements lasting in excess of 10 min. Third, INRIX reports two measures of confidence and we investigated whether these confidence measures reflected the repeated INRIX speed measurements or not. Most of the repeated measurements had no impact on the confidence measures, and the confidence measures do not provide any insight into how many times a reported value has been repeated.

6. Recommendations for Implementation of Research Findings

This research examined the performance of SpeedInfo and INRIX traffic speed data. Both sources generally performed within specifications. Both systems exhibited small biases that could likely be fixed with further fine tuning of the settings (e.g., a simple scale factor), though even without this fine tuning the values were sufficiently accurate for many applications. The biases tended to be larger for SpeedInfo, but still sufficiently accurate for congestion detection and traveler information applications. Other applications, such as traffic responsive ramp metering, might need further tuning of the sensitivity. Both sensors also exhibited a reporting lag, on the order of 60 sec for SpeedInfo and 360 seconds for INRIX. These lags are comparable to conventional 60 sec and 5 min aggregation periods used in loop detectors, but again, some applications might be sensitive to these lags, e.g., traffic responsive ramp metering.

Each system also exhibited unique behavior not found in the other system. SpeedInfo is sensitive to precipitation, it went offline for an extended period for tens of minutes on several occasions in the study month, concurrent with the onset of precipitation. When mounted on the shoulder, SpeedInfo also showed slightly degraded performance on the far side traffic. The impact does not seem to be large enough to merit redeploying the existing sensors. However, for future deployments it may be better to deploy them in the median when it is an option.

INRIX reported speeds every minute, but after excluding repeated values, the actual effective reporting period was more like 3-5 min, with occasional periods of repeated measurements lasting in excess of 10 min (note that this time is already reflected in the lag mentioned above). INRIX also provides two measures of confidence. These confidence measures do not appear to reflect the fact that a given measurement has been repeated.

ODOT should continue monitoring the data quality from these two systems. Many tests could be done in house, e.g., on a corridor with recurring congestion,

1) at a location near the downstream end of a recurring queue set up an alarm to trigger off real-time congestion found in either INRIX or SpeedInfo.
2) have an operator watch a CCTV view a few miles upstream of the alarm location to make sure it is presently uncongested.
3) watching the location from #2, note the time that congestion sets in
4) watching the real time data from SpeedInfo or INRIX, note the time that the data series responds to the congestion
5) now swap the two locations, wait for free flow conditions to return upstream to trigger an automated alarm, and then manually monitor the thee systems at the downstream location as the queue recedes downstream.

In the case of vulnerability to precipitation, the operator could watch RWIS data or a CCTV camera as the weather starts to rain or snow. Then watch to see if the real-time SpeedInfo data is impacted, and if so, for how long.

The daily biases should be easier to detect, e.g., the median speed reported over an entire 24 hr period should be sufficiently close to the expected speed on that facility. Finally, in light of the findings of this study, ODOT should review the traffic monitoring specifications and make sure that they remain sufficient for all of the applications envisioned for the data.

Bibliography

[1] Buckeye traffic, http://www.buckeyetraffic.org/ (accessed September 12, 2013), and recently migrated to http://www.ohgo.com/ (accessed December 17, 2013)

[2] Hu, H., (2009) *Measuring the Effectiveness of Advanced Traveler Information Systems (ATIS)*, PhD Dissertation, North Carolina State University.

[3] Middleton, D., Longmire, R., Turner, S., (2007) *State of the Art Evaluation of Traffic Detection and Monitoring Systems, Volume I – Phases A & B: Design*, Final Report 627 (1), Arizona Department of Transportation

[4] INRIX traffic, http://www.inrix.com/ (accessed September 12, 2013)

[5] Aliari, Y., Haghani, A. (2012) Bluetooth Sensor Data and Ground Truth Testing of Reported Travel Times. *Transportation Research Record: Journal of the Transportation Research Board, Issue 2308*, pp 167-172

[6] Jia, C., Li, Q., Oppong, S., Ni, D., Collura, J., Shuldiner, P. (2013) Evaluation of Alternative Technologies to Estimate Travel Time on Rural Interstates. *Proc. of the Transportation Research Board 92nd Annual Meeting*, Transportation Research Board, 16p.

[7] Lattimer, C., Glotzbach, G. (2012) Evaluation of Third Party Travel Time Data. *Proc. of the ITS America 22nd Annual Meeting & Exposition*, ITS America, 8p

[8] Kim, K., Chien, S., Spasovic, L. (2011) Evaluation of Technologies for Freeway Travel Time Estimation: Case Study of I-287 in New Jersey. *Proc. of the Transportation Research Board 90th Annual Meeting*, Transportation Research Board, 17p

[9] Coifman, B., (2006) *The Columbus Metropolitan Freeway Management System (CMFMS) Effectiveness Study: Part 2 - The After Study*, Ohio Department of Transportation.

[10] Lee, H., Coifman, B., (2012) An Algorithm to Identify Chronic Splashover Errors at Freeway Loop Detectors, *Transportation Research-Part C*, vol. 24, pp 141-156.

[11] Lee, H., Coifman, B., (2012) Quantifying Loop Detector Sensitivity and Correcting Detection Problems on Freeways, *Journal of Transportation Engineering, ASCE*, vol. 138, no. 8, pp 1051-1061.

[12] Lee, H., Coifman, B., (2011) Identifying and Correcting Pulse-Breakup Errors from Freeway Loop Detectors , *Transportation Research Record 2256*, pp 68-78.

[13] Neeley, M., Bailey, T., (2008) WSDOT Field Tests of Wireless and Microwave Vehicle Detection Systems, http://westernstatesforum.org/Documents/WSDOT_FINALPresentationfortheWesternStatesForum_withvideo.pdf

[14] Sawyer, S., *Production of a Radar Monitoring System: Delivering Accurate Traffic Speed Data Over a Communication Network*, undergraduae honors thesis, Ohio State University, March 8, 2006.

[15] Denning, N., *Implementation and Validation of a Cost-Effective Networked Radar-Based Highway Traffic Speed Measurement System*, undergraduae honors thesis, Ohio State University, May 30, 2007.

[16] Haghani, A., Hamedi, M., Sadabadi, K., (2009) *I-95 Corridor Coalition Vehicle Probe Project: Validation Of Inrix Data July-September 2008, Final Report*, I-95 Corridor Coalition

Appendix A- Time series evaluation of SpeedInfo system against loop detectors at stations with recurring congestion

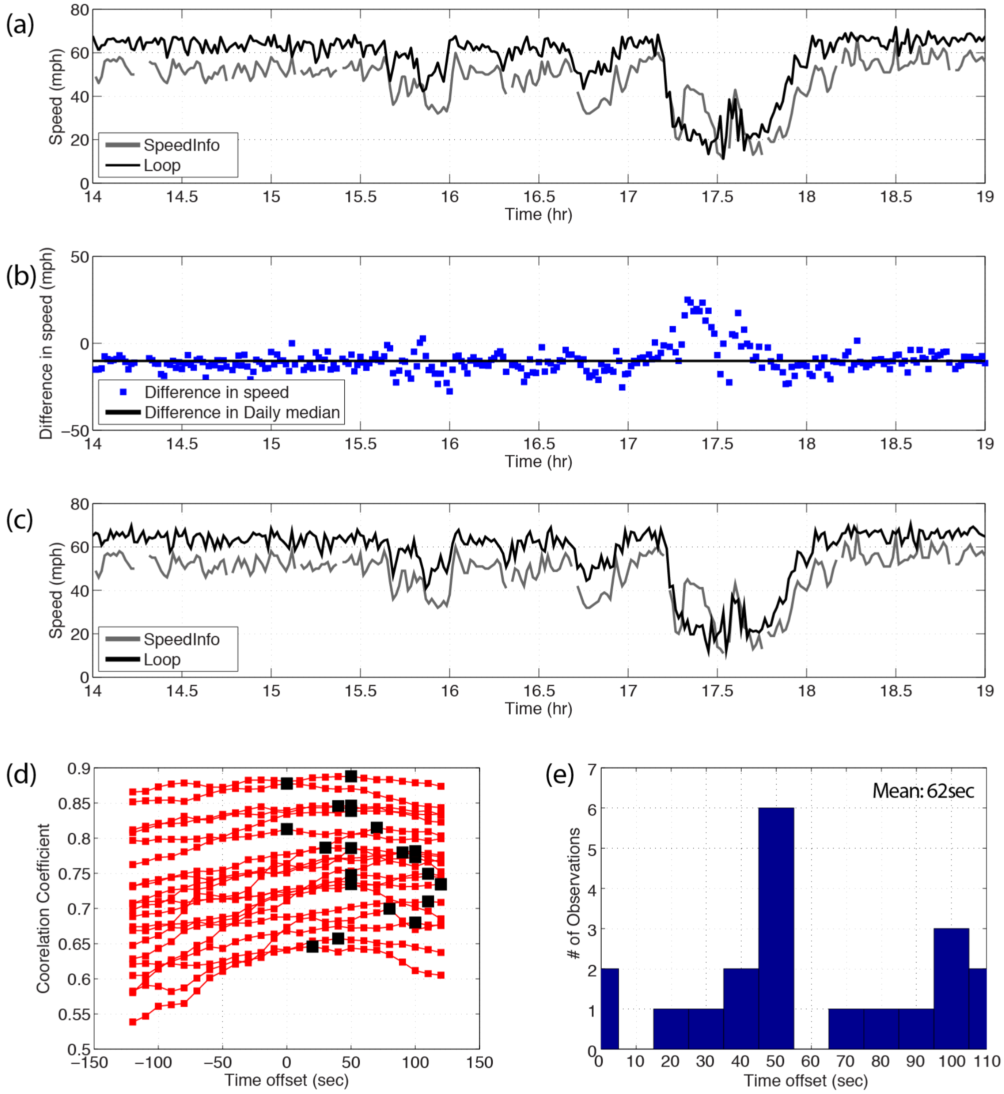

Figure A-1: (a) 1 min aggregated speed as a function of time from the loop detector and SpeedInfo systems at Station 1 southbound (April 6, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at -10 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 70 sec, (d) the correlation coefficient versus the time offset for 20 weekdays, (e) resulting distribution from (d), with mean of 62 sec.

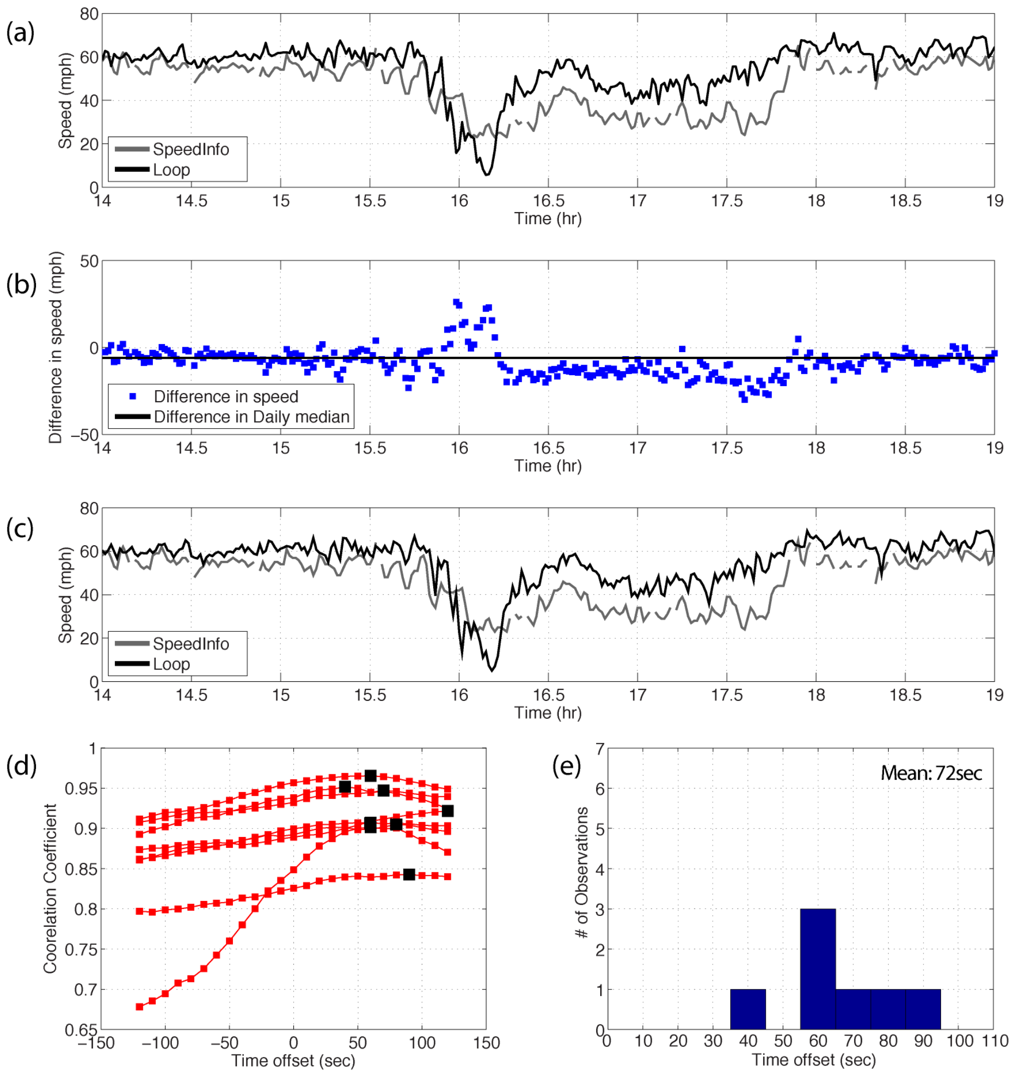

Figure A-2: (a) 1 min aggregated speed as a function of time from the loop detector and SpeedInfo systems at Station 103 northbound (April 5, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at -5 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 100 sec, (d) the correlation coefficient versus the time offset for 8 weekdays, (e) resulting distribution from (d), with mean of 72 sec.

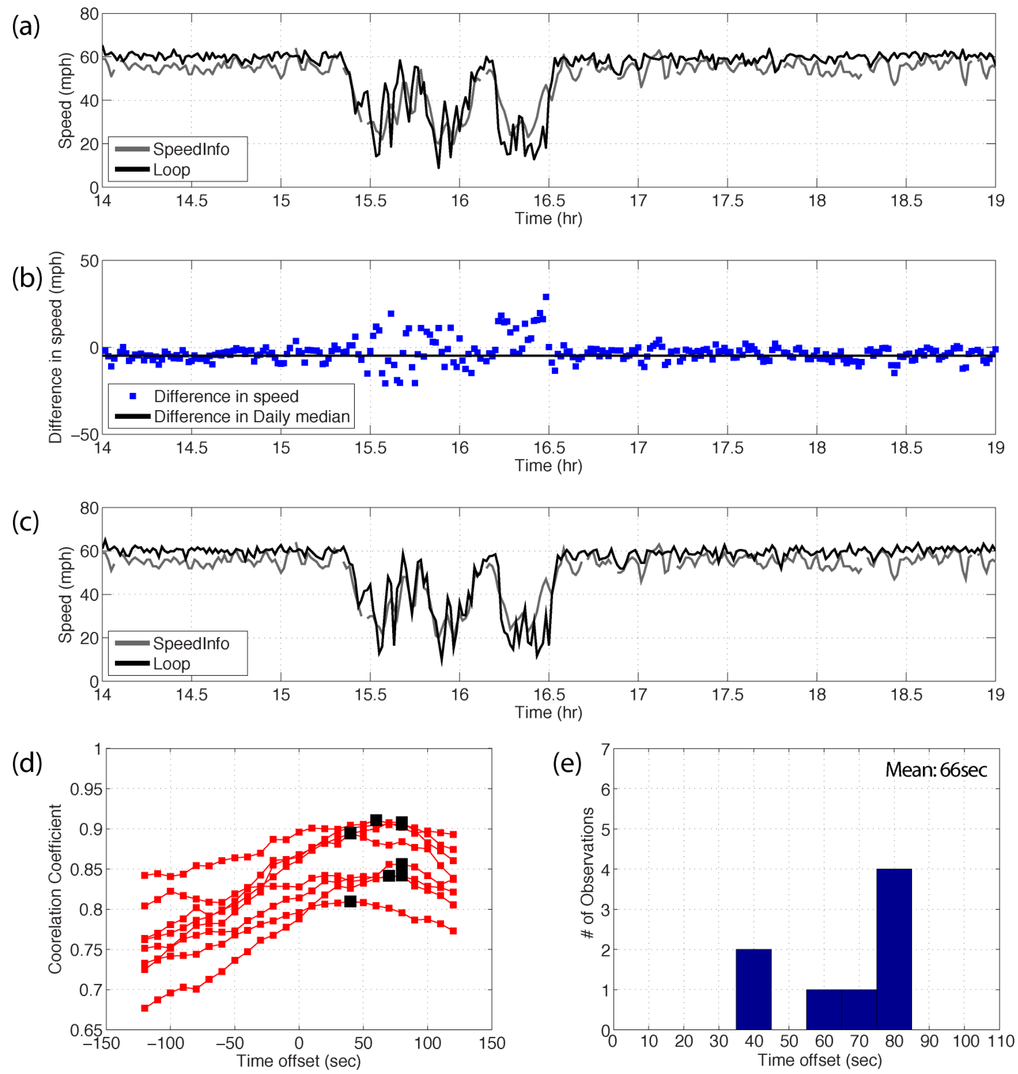

Figure A-3: (a) 1 min aggregated speed as a function of time from the loop detector and SpeedInfo systems at Station 103 southbound (April 6, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at -5 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 70 sec, (d) the correlation coefficient versus the time offset for 8 weekdays, (e) resulting distribution from (d), with mean of 66 sec.

Appendix B- Time series evaluation of INRIX system against loop detectors at stations with recurring congestion

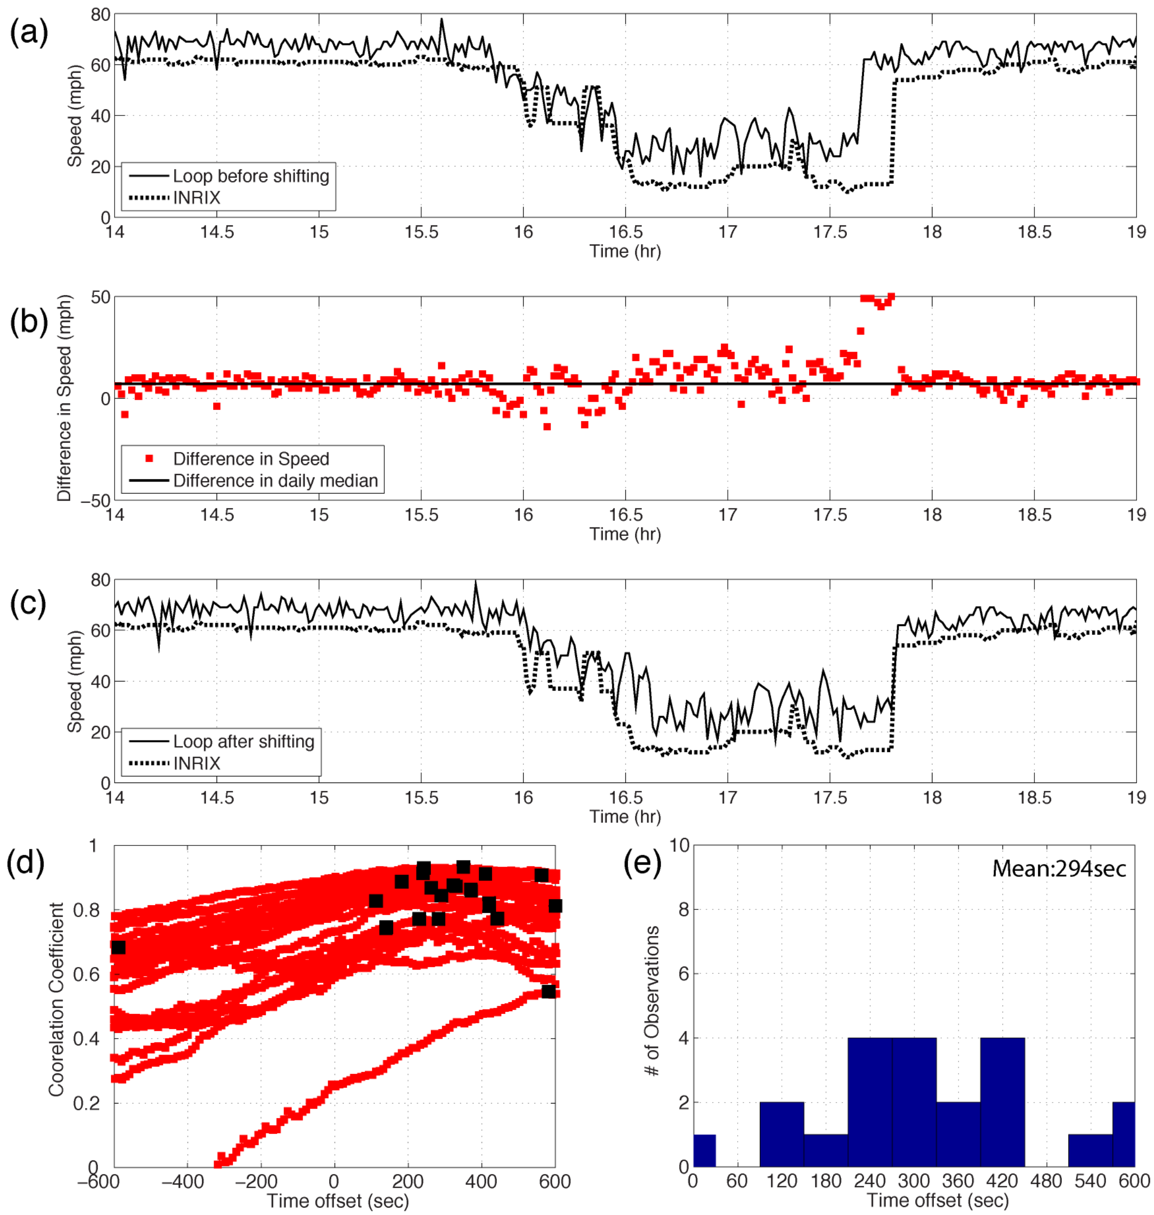

Figure B-1: (a) 1 min aggregated speed as a function of time from the loop detector and INRIX systems at Station 1 southbound (May 17, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at 7 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 600 sec, (d) the correlation coefficient versus the time offset for 20 weekdays, (e) resulting distribution from (d), with mean of 294 sec.

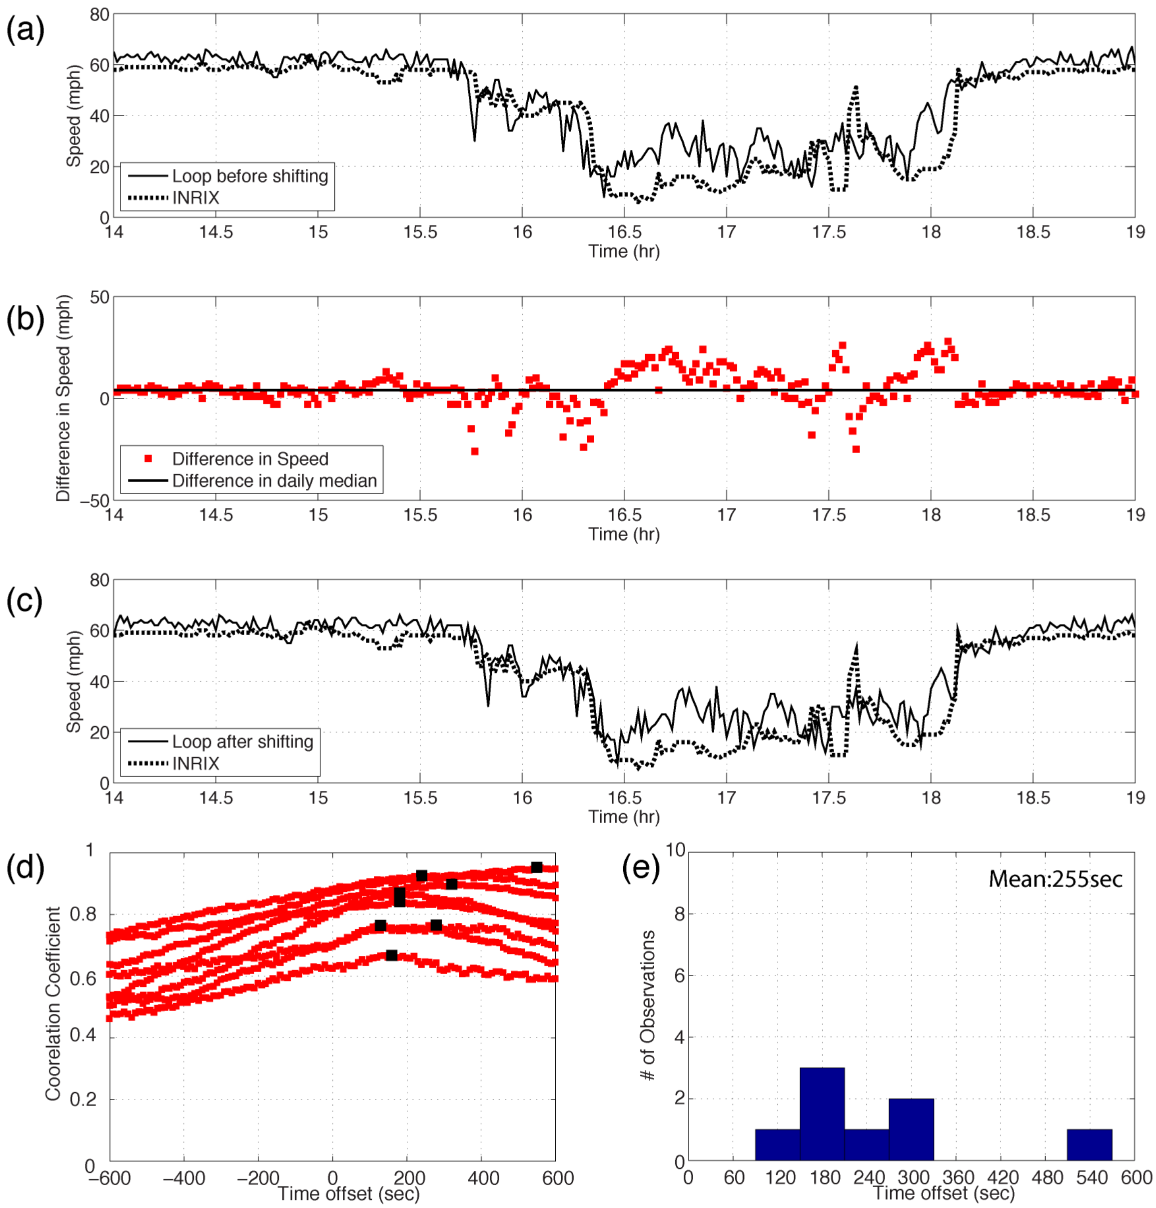

Figure B-2: (a) 1 min aggregated speed as a function of time from the loop detector and INRIX systems at Station 105 northbound (May 17, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at 4 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 240 sec, (d) the correlation coefficient versus the time offset for 8 weekdays, (e) resulting distribution from (d), with mean of 255 sec.

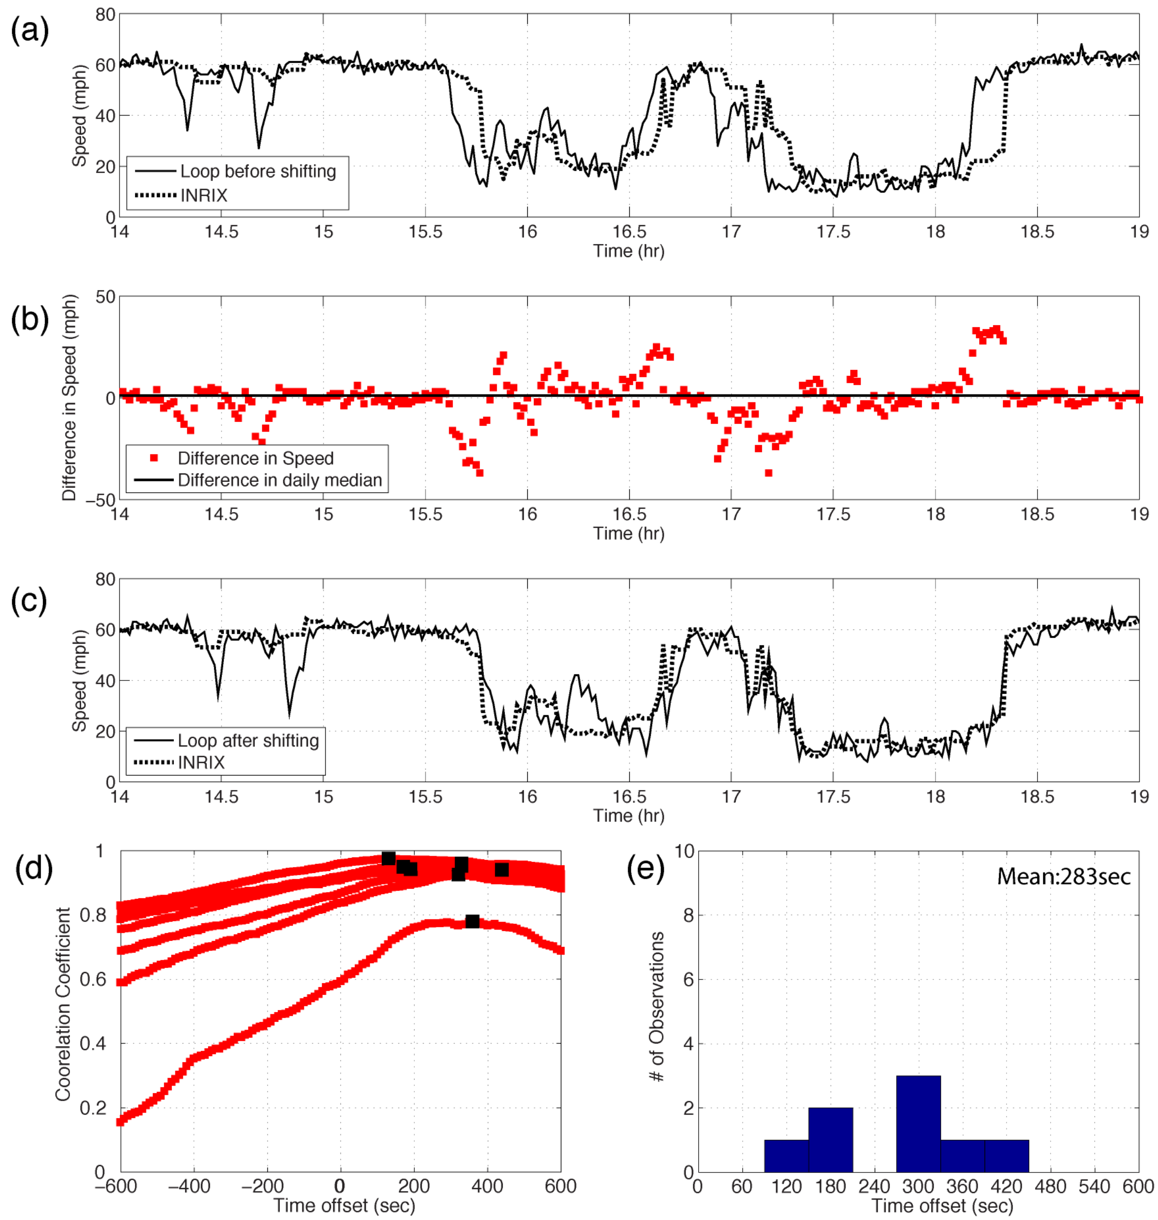

Figure B-3: (a) 1 min aggregated speed as a function of time from the loop detector and INRIX systems at Station 103 southbound (May 26, 2011), (b) difference between the two concurrent time-series speeds in (a) and showing day's median bias with a horizontal line at 1 mph, (c) time-series speed corresponding to (a) after shifting the loop detector time-series by 540 sec, (d) the correlation coefficient versus the time offset for 8 weekdays, (e) resulting distribution from (d), with mean of 283 sec.

Appendix C- Potential bias of INRIX measurement due to the parallel road

In Figure 11, we recognized abrupt speed drops at 18 hr and 18.5 hr from INRIX while corresponding loop detector reports generally free flowing speed. Figure C-1 shows the road network near Station 1, and INRIX links downstream and upstream of it. As highlighted with a bold line in Figure C-1b, there is a freeway connector parallel to I-71. Meanwhile there is another parallel freeway connector downstream of Station 1 as illustrated in Figure C-1c. In fact both of these connectors merge to I-71. To check whether INRIX may inadvertently measure speed from the parallel connector when its traffic condition differs from I-71, we investigate the speed from 5 neighboring INRIX links as well as speed from loop detectors shown in Figure C-1 and see if the speed drops in Figure 11 are associated with conditions on the parallel road.

Figure C-1 (a) Two INRIX links downstream of northbound Station 1, (b) the INRIX link corresponding to northbound Station 1, and (c) two INRIX link upstream of northbound Station 1.

From the five INRIX links shown in Figure C-1, we look at time-series speed and compare with corresponding loop detector speed (if applicable). As can be seen in Figure C-2, those speed drops observed in Figure C-2c are not shown in Figure C-2a while one or both of the drops seems evident in other links in Figure C-2b, C-2d, and C-2e. However, the two parallel connectors highlighted in Figure C-1b and C-1c are independent. Thus it is unlikely that the speed drops Figure C-2c are due to bias from local traffic conditions in the parallel connector. Meanwhile, we also compare loop detector speed from Station 1 per lane with the corresponding INRIX link speed, as shown in Figure C-3. None of the individual lanes in the loop detector data exhibit speed drops coincident with the two INRIX speed drops.

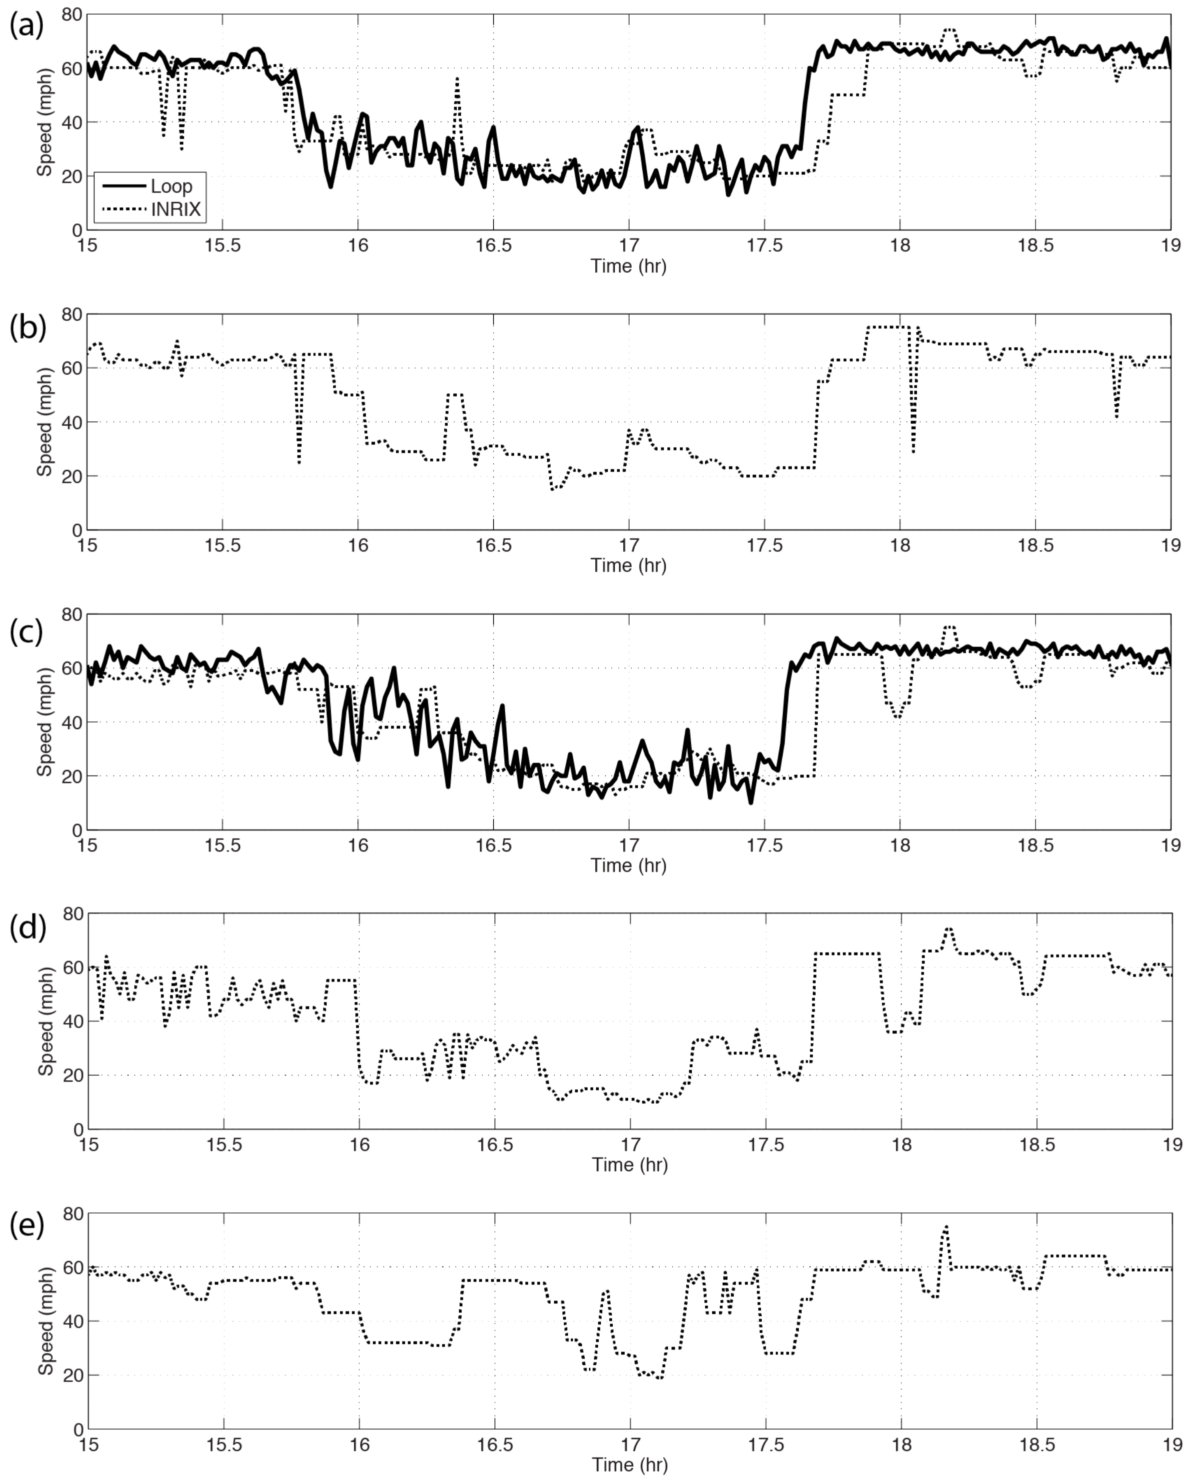

Figure C-2: Speed from loop detectors and corresponding INRIX links on April 15, 2011 northbound, (a) INRIX link 122P04192, (b) 122+04192, (c) 122+04191, (d) 122P04190, (e) 122+04190.

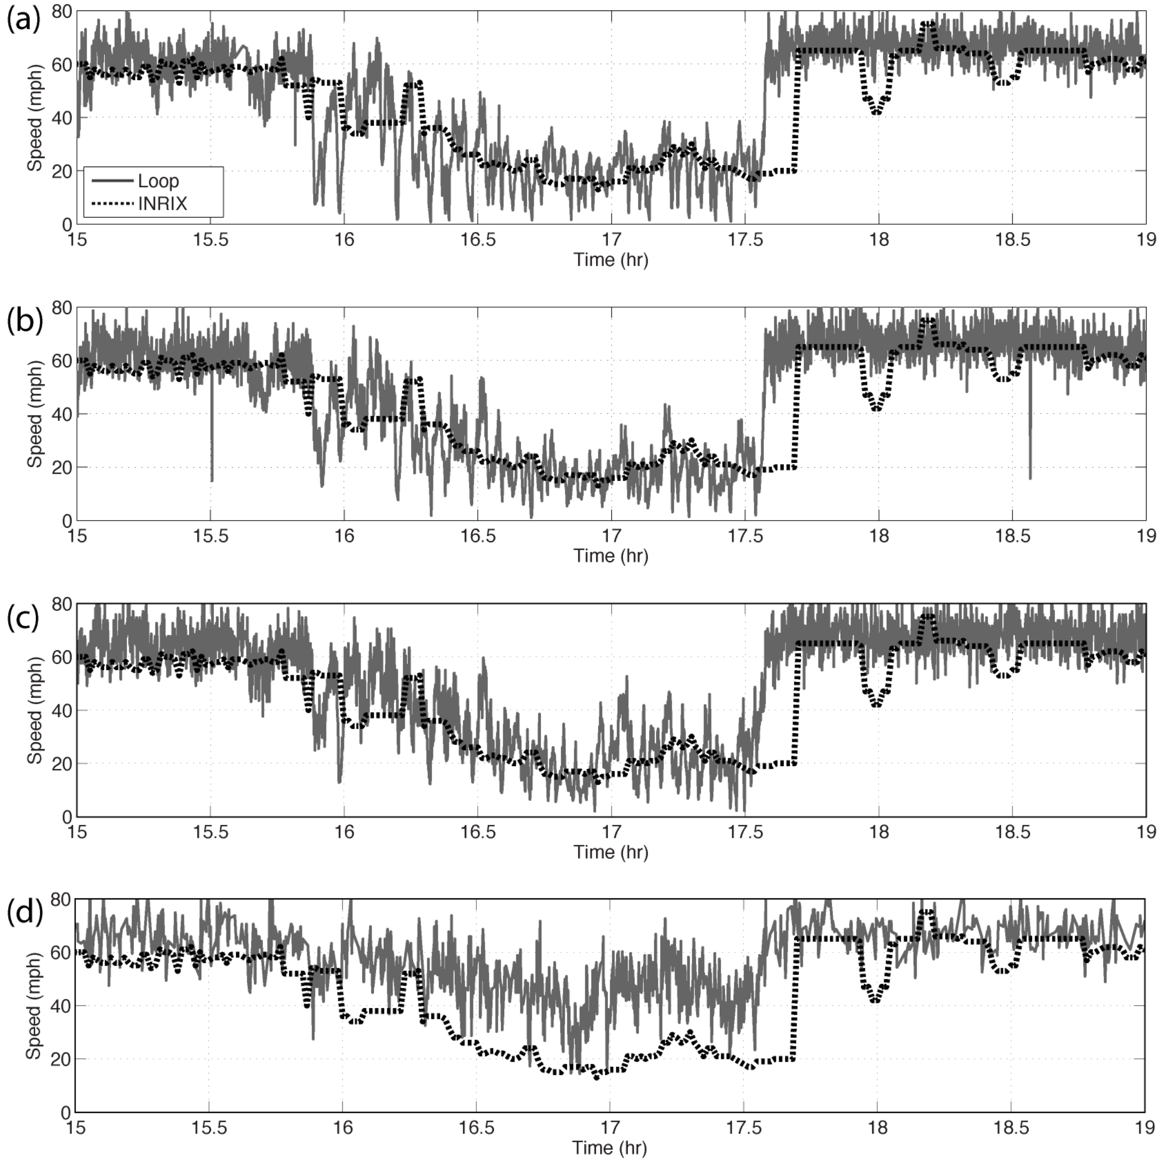

Figure C-3: Speed from INRIX (all lanes) and the corresponding loop detector lane by lane (a) lane 1, (b) lane 2, (c) lane 3, (d) lane 4, April 15, 2011, Station 1, northbound

If the speed drop in Figure C-2c is not associated with traffic conditions in the parallel connector, another possibility would be that the speed drop comes from a specific lane in I-71. As can be seen in Figure C-3, however, none of the lanes at Station 1 exhibits the speed drop corresponding INRIX link. While we could not find clear evidence of the speed drops based on this investigation, we found that it is a common phenomenon, occurring over several different days, e.g., see Figure C-4 and C-5.

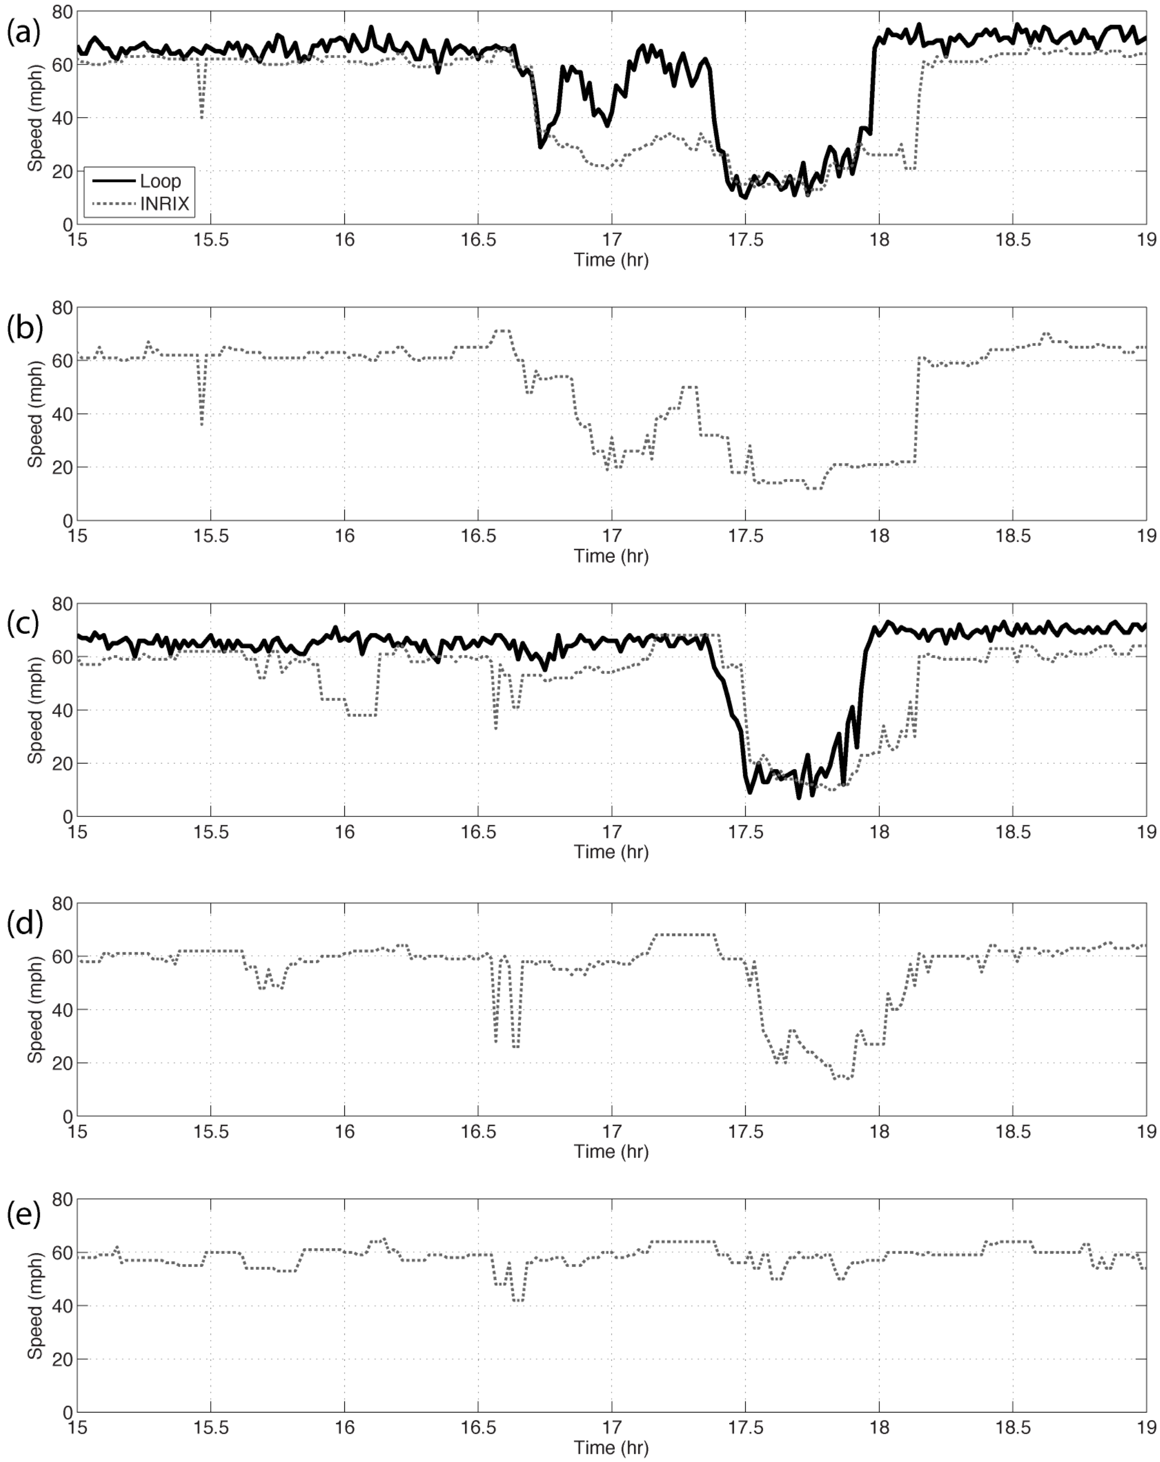

Figure C-4: Speed from loop detectors and corresponding INRIX links on May 3, 2011 northbound, (a) INRIX link 122P04192, (b) 122+04192, (c) 122+04191, (d) 122P04190, (e) 122+04190.

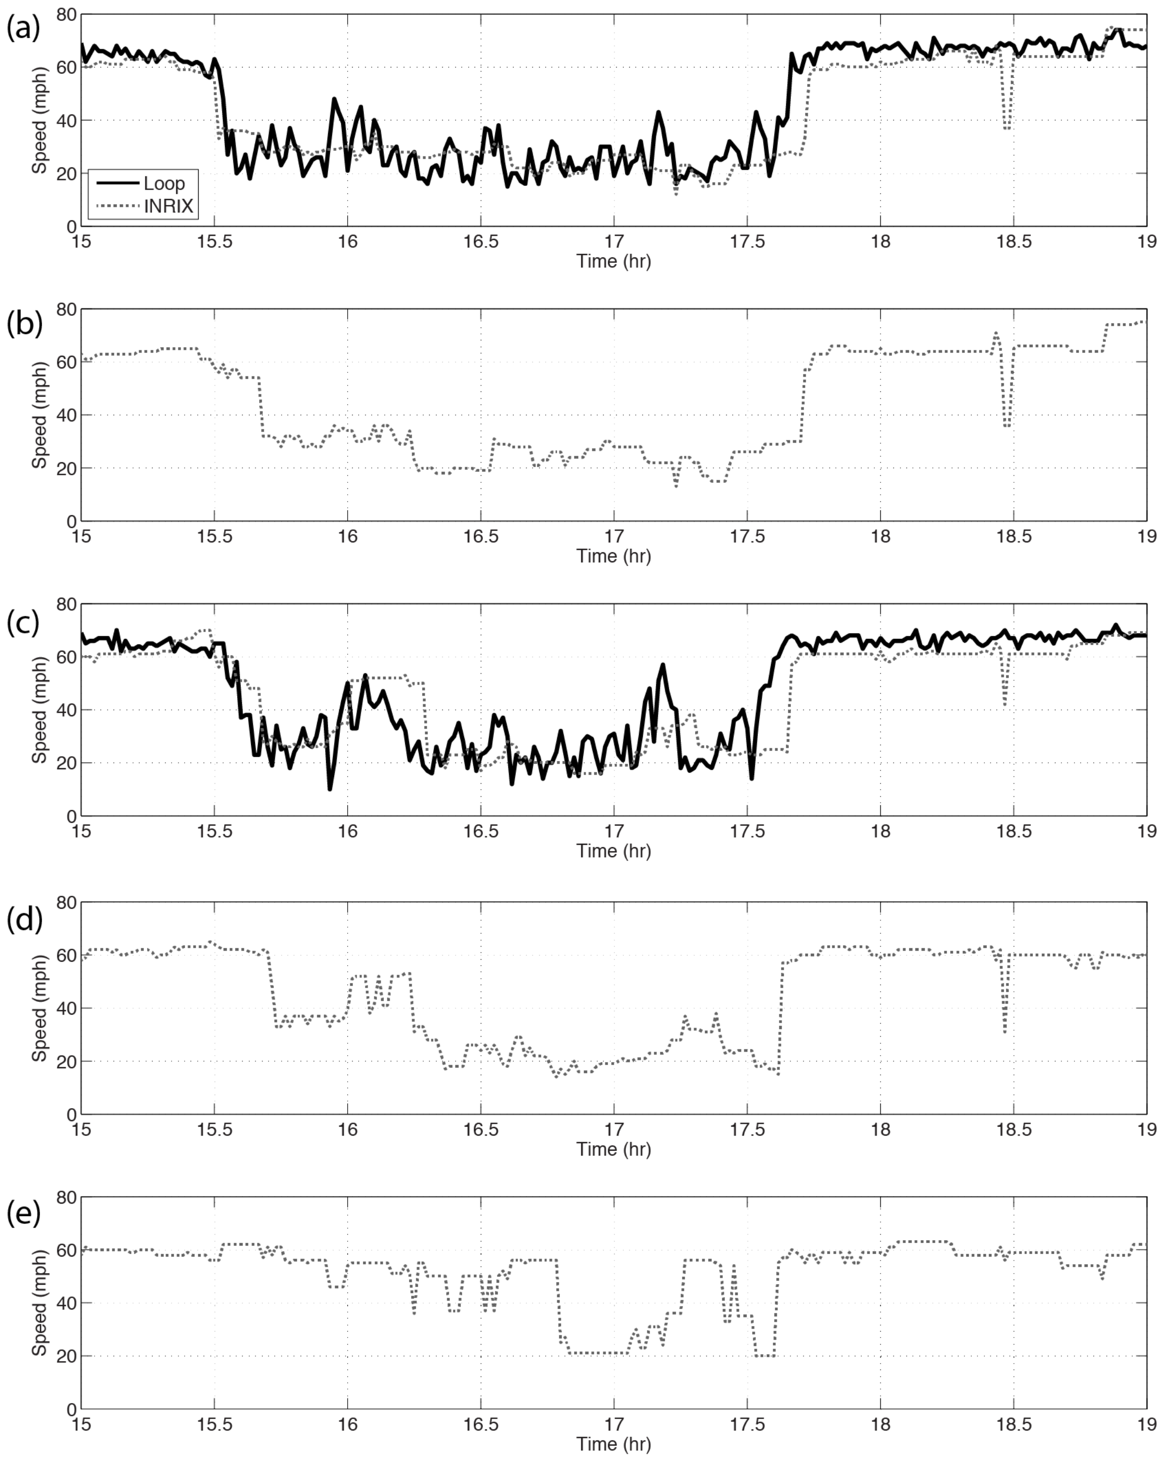

Figure C-5: Speed from loop detectors and corresponding INRIX links on May 27, 2011 northbound, (a) INRIX link 122P04192, (b) 122+04192, (c) 122+04191, (d) 122P04190, (e) 122+04190.

Appendix D- Repeated measurements reported by INRIX

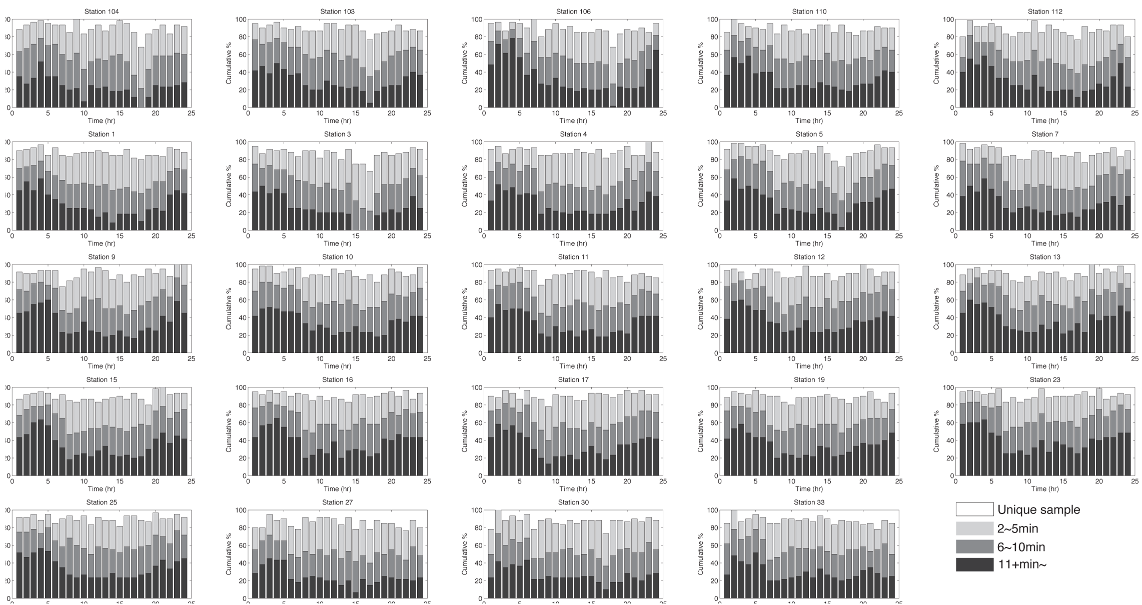

Figure D-1: Binned cumulative percentage of monthly median durations with non-unique samples for each southbound station in May 2011.

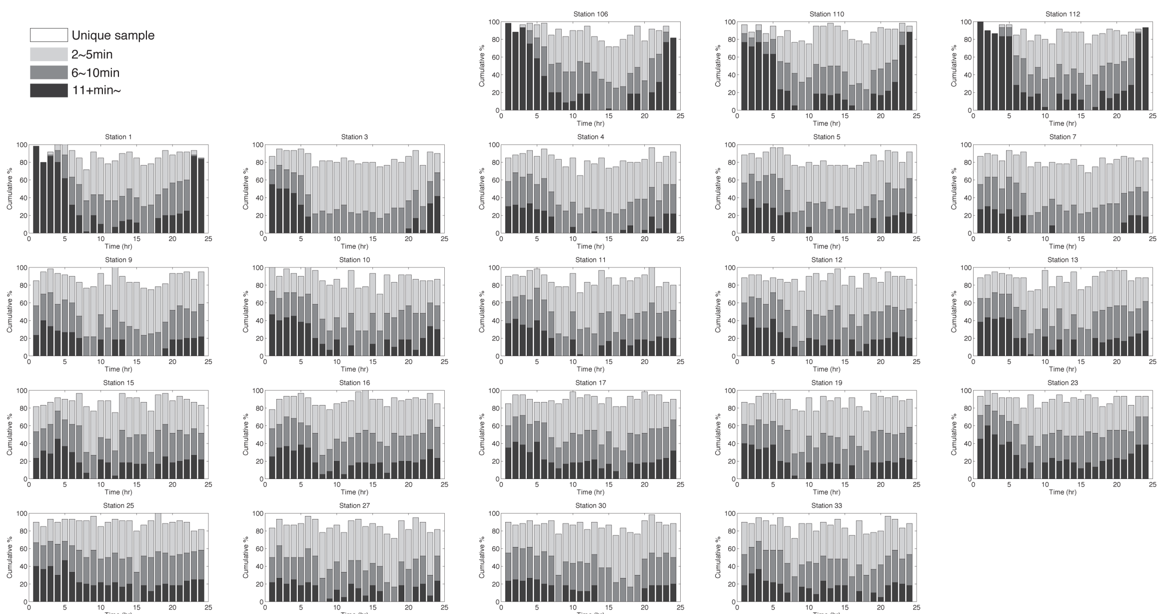

Figure D-2: Binned cumulative percentage of monthly median durations with non-unique samples for each southbound station in May 2013.

Appendix E- INRIX confidence score by southbound station in 2011 and 2013

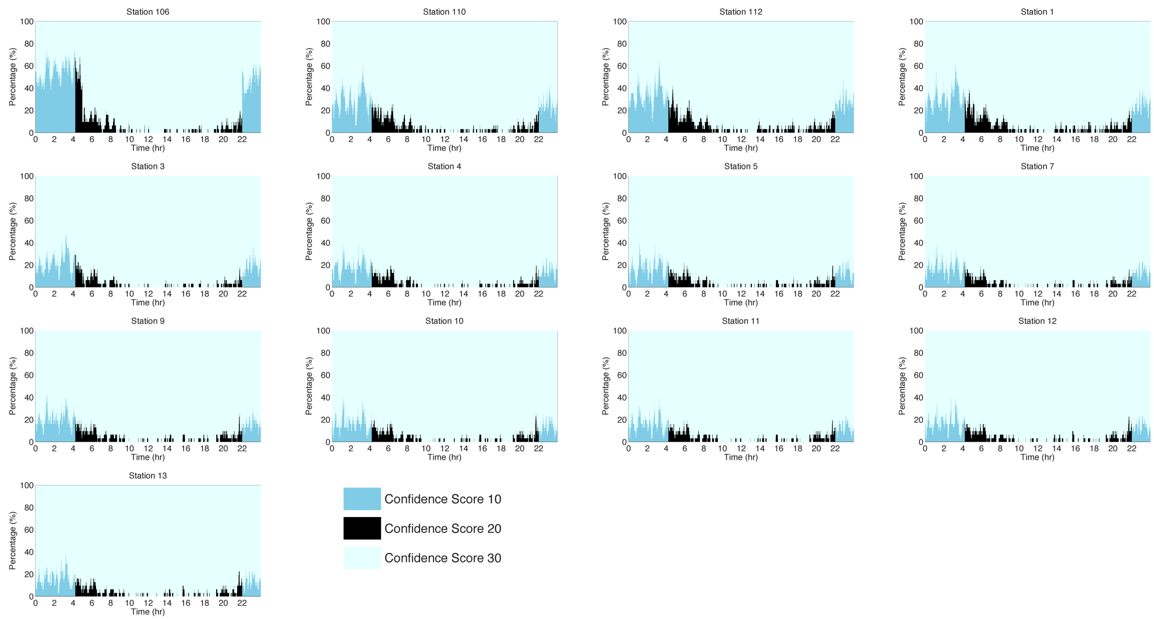

Figure E-1: Percentage of samples with a given confidence score as a function of time over the entire month of May 2011 from the southbound stations.

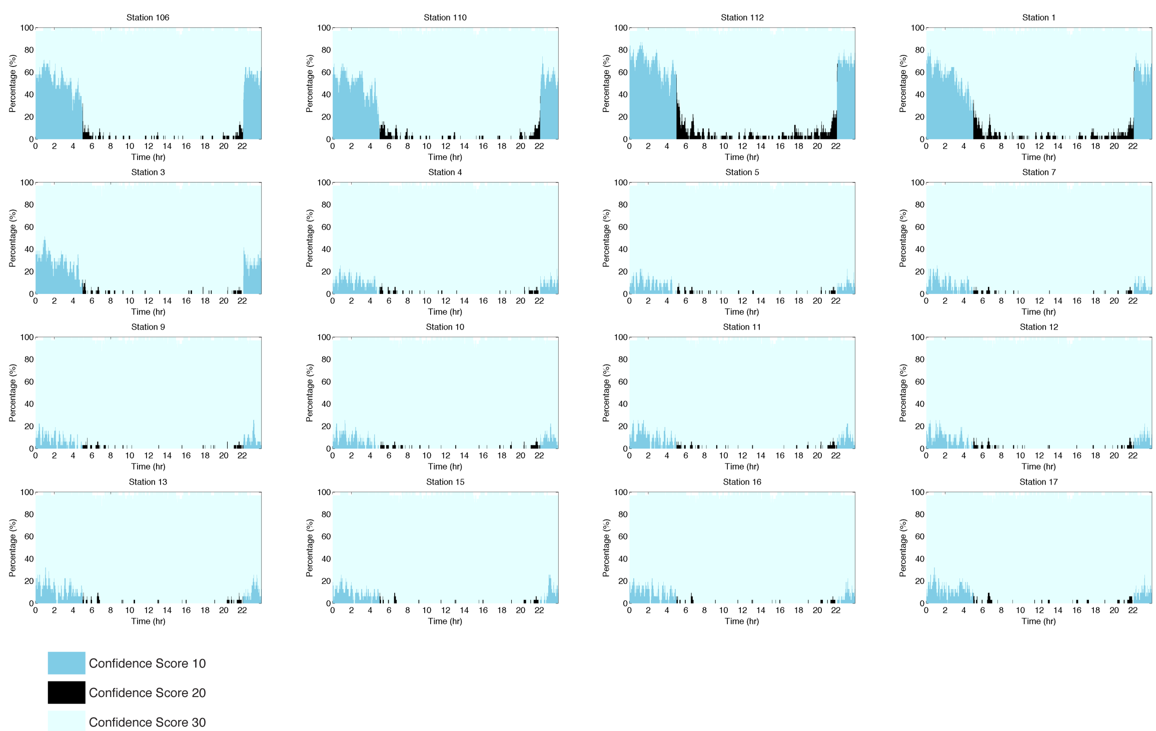

Figure E-2: Percentage of samples with a given confidence score as a function of time over the entire month of May 2013 from the southbound stations.

Table E-1: Total number of speed records reported for a given confidence score (CS) by southbound station.

Station	May 2011				May 2013			
	CS10	CS20	CS30	Total	CS10	CS20	CS30	Total
St. 106	6,356 (14.2%)	2,118 (4.7%)	36,166 (81.0%)	44,640 (100%)	7,161 (16.0%)	561 (1.3%)	36,308 (81.3%)	44,030 (98.63%)
St. 110	3,445 (7.7%)	1,808 (4.1%)	39,387 (88.2%)		6,808 (15.3%)	774 (1.7%)	36,448 (81.6%)	
St. 112	3,941 (8.8%)	2,223 (5.0%)	38,476 (86.2%)		9,031 (20.2%)	1926 (4.3%)	33,073 (74.1%)	
St. 1	3,646 (8.2%)	2,185 (4.9%)	38,809 (86.9%)		7,691 (17.2%)	1,337 (3.0%)	35,002 (78.4%)	
St. 3	2,653 (5.9%)	1,138 (2.5%)	40,849 (91.5%)		3,945 (8.8%)	309 (0.7%)	39,776 (89.1%)	
St. 4	2,176 (4.9%)	1,181 (2.6%)	41,283 (92.5%)		1,622 (3.6%)	271 (0.6%)	42,137 (94.4%)	
St. 5	2,053 (4.6%)	1,284 (2.9%)	41,303 (92.5%)		1,341 (3.0%)	253 (0.6%)	42,436 (95.1%)	
St. 7	2,038 (4.6%)	1,076 (2.4%)	41,526 (93.0%)		1,376 (3.1%)	216 (0.5%)	42,400 (95.0%)	43,994 (98.55%)
St. 9	2,452 (5.5%)	1,351 (3.0%)	40,837 (91.5%)		1,269 (2.8%)	242 (0.5%)	42,481 (95.2%)	
St. 10	2,027 (4.5%)	1,260 (2.8%)	41,353 (92.6%)		1,318 (3.0%)	298 (0.7%)	42,376 (94.9%)	
St. 11	1,956 (4.4%)	1,201 (2.7%)	41,483 (92.9%)		1,555 (3.5%)	279 (0.6%)	42,158 (94.4%)	
St. 12	2,011 (4.5%)	1,237 (2.8%)	41,392 (92.7%)		1,471 (3.3%)	296 (0.7%)	42,225 (94.6%)	
St. 13	1,978 (4.4)	1,207 (2.7%)	41,455 (92.9%)		1,821 (4.1%)	233 (0.5%)	41,938 (93.9%)	