Management of intersections with multi-modal high-resolution data [☆],[☆☆]

Ajith Muralidharan [1], Samuel Coogan [2], Christopher Flores, Pravin Varaiya [*]

Sensys Networks, Inc, Berkeley, CA 94710, United States

ARTICLE INFO

ABSTRACT

A high-resolution (HR) data system for an intersection collects the location (lane), speed, and turn movement of every vehicle as it enters an intersection, together with the signal phase. Some systems also provide video monitoring; others measure pedestrian and bicycle movements; and some have vehicle to infrastructure (V2I) communication capability. The data are available in real time and archived. Real time data are used to implement signal control. Archived data are used to evaluate intersection, corridor, and network performance. The system operates 24×7. Uses of a HR data system for assessing intersection performance and improving mobility and safety are discussed. Mobility applications include evaluation of intersection performance, and the design of better signal control. Safety applications include estimates of dilemma zones, red-light violations, and pedestrian–vehicle conflicts.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Poor management of intersections causes excessive delay and more frequent crashes. Conditions are worse in cities experiencing rapid growth in automobile ownership, but mature cities also face challenges as road capacity is taken away from vehicles to accommodate increasing bicycle and pedestrian traffic. Management today is handicapped by insufficient data. It can be more effective if it is based on high-resolution (HR) real-time data about the movement of vehicles, bicycles, and pedestrians. "If you don't know what's happening on your roads, don't expect to manage them well" is a truism.

A basic HR system should collect the location (lane) and speed of every vehicle as it approaches and enters the intersection, together with the signal phase. Some intersections may have video cameras for purposes of surveillance and enforcement, and these cameras can also be used to calibrate algorithms that process the raw HR data to estimate the traffic state. The system may have sensing capability beyond vehicle detection: there may be pedestrian occupancy sensors at crosswalks; sensors that count the passage of bicycles in designated lanes; sensors that classify vehicles (passenger cars vs trucks or buses); sensors that capture Bluetooth or WiFi addresses from radios in devices within the vehicles; and sensors that detect if a parking space is occupied.

Urban traffic varies over time and space. Efficient management requires control of the infrastructure that quickly predicts and reacts to this variable traffic. So the HR data system must operate 24 × 7. The data must be available in real time and archived for performance assessment and planning control strategies. This paper describes uses of a HR system for intersection management.

1.1. Urban traffic control in the U.S.

Urban traffic in the U.S. is regulated by 300,000 signalized intersections, whose performance is determined by their signal control algorithms. The performance is poor: the 2012 National Transportation Operations Coalition (NTOC) assessment of traffic signal control gives an overall grade of D+ (National Traffic Operations Coalition, 2012). Ninety percent of the signalized intersections use fixed time of day (TOD) plans, which are re-timed once in five years, based on three days of manual data collection. These traffic snapshots and the timing plans based on them completely miss the variability in the traffic to which the plans should adapt. Moreover, the infrequent snapshots do not permit the operating agencies or the public to determine whether the road network is performing well or poorly.

The 2007 NTOC report (National Traffic Operations Coalition, 2007) gives U.S. cities a grade of F for traffic monitoring and data collection. The report estimates the investment needed to maintain grade A at $3 per household, resulting in savings of $100 per year. These are major benefits, and known signal control algorithms can achieve these benefits. But current traffic detector deployments are too limited in spatial coverage and temporal granularity to provide the information these algorithms require.

The NCHRP report (Dowling and Ashiabor, 2012) found 98 percent of agencies use qualitative rather than data-driven techniques to determine the need for plan changes in TOD systems. The report is devoted to traffic signal operations and timing plans in the face of variability in demand and saturation rates. The report describes methodologies for robust operations, but observes that this 'data-rich' approach can today be followed only by "three dozen or so agencies in the United States capable of routinely collecting and archiving second-by-second advanced detector volume and occupancy for extended periods of time (on the order of one full year)."

Fixed-time controllers provide "walk/don't walk" signals for pedestrians at marked crosswalks; actuated signals may also provide push-buttons. However, virtually no city collects data on pedestrian crosswalk use and the extent of conflict between pedestrians and certain vehicle turn movements (e.g. 'right turn on red'). The absence of data prevents management to focus on safety. This lack of attention is costly. According to NHTSA (2010) 40 percent of nearly 6 million crashes in 2008 occurred at intersections, of which 7421 were fatal crashes and 733,000 involved injuries. The detailed study by New York DOT (2010) estimates that traffic crashes cost the city $4.29 billion annually, and pedestrians accounted for 52% of traffic fatalities during 2005–2009. Red-light and speeding violations are safety hazards, even when they don't result in crashes. Such violations are almost never measured, although knowing the frequency of these violations may suggest modifications in the signal plan or enforcement initiatives to reduce the frequency. Some red-light violations may be due to drivers finding themselves in a dilemma zone, whose presence and extent can only be determined from HR data.

Provisions for pedestrian crossings began in the early 1900s. But demand for exclusive bicycle lanes (as well as for transit vehicles and taxis) only recently became politically effective in the U.S. Again there is a virtual absence of data on bicycle
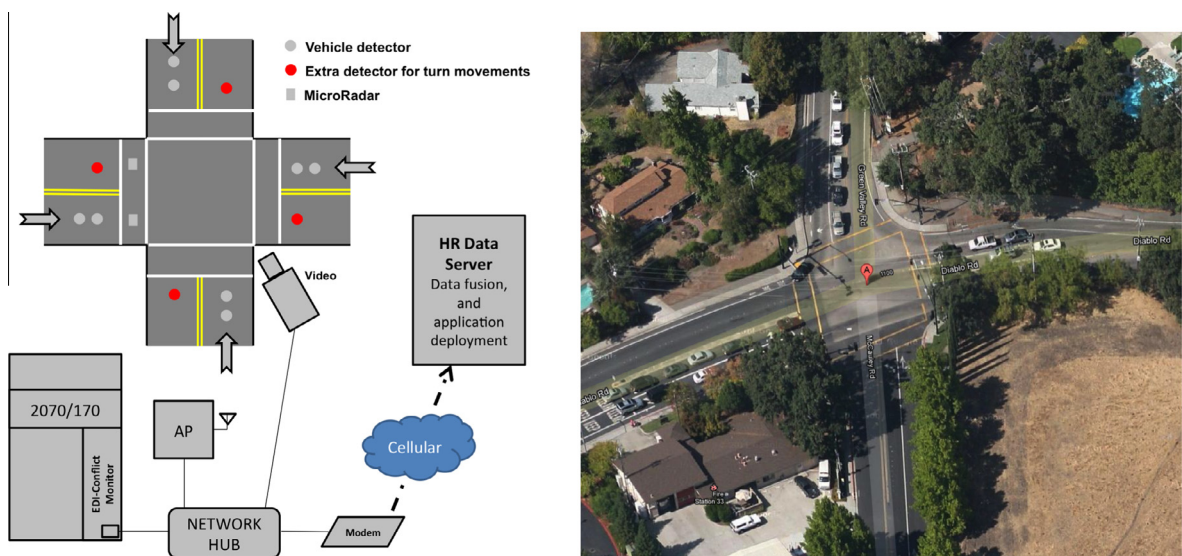

Fig. 1. HR system in Danville, CA.

counts or estimates of the utilization of transit-only lanes. A similar ignorance undermines understanding of the utilization and pricing policy of parking places. Not only are parking fees an important revenue source, but poor parking policies impose additional costs in the form of congestion caused by 'double-parking' and by drivers looking for empty spaces. Pierce and Shoup (2013) report on an experiment in San Francisco to price parking spaces in response to demand.

In summary, efficient management of urban roads is handicapped by the lack of high-resolution data. Current detection systems rarely provide data at the granularity of individual vehicles; there is virtually no data on pedestrians and bicycles; and data that relates traffic control policies to safety are almost non-existent. In the absence of an adequate data system, management will continue to be uninformed, and the public will be unable to call management to account.

A study of the uses of high-resolution data occupies most of the rest of the paper. Section 2 considers automated design of TOD plans. The point is to show how HR data clearly reveals variability of traffic. Quantifying the variability leads to timing plans that better match demand. HR data can be processed to provide a wide range of mobility measures summarized in Section 3. Some safety applications involving red light violations and pedestrian–vehicle conflicts are described in Section 4. Concluding remarks are collected in Section 5.

The applications discussed below use data from HR systems at two intersections, one in Danville, CA and another in Beaufort, SC. To appreciate how the data is collected briefly describe the hardware architecture of the HR system at Danville. The system components are from Sensys Networks, Inc.

1.2. HR system architecture

The Danville intersection is shown on the right in Fig. 1; on the left is a schematic of the system components. The four-way intersection contains three kinds of detectors: a pair of stop bar vehicle magnetic detectors spaced 6 feet apart permits speed measurement. An additional detector (red[3]) is placed in the departure lane to measure turn movements. One crosswalk is equipped with (square) MicroRadar pedestrian detectors. (The MicroRadar also detects bicycles and parked vehicles, but not at this intersection.) All detectors are wireless and communicate with the Access Point (AP) located near the 2070/170 controller. Signal phase data is obtained from the controller conflict monitoring card. There is in addition a PTZ camera. All measurements are time-stamped and synchronous to within 100 ms or 0.1 s. The data are sent to the HR data server via a cellular modem. The Beaufort intersection of Fig. 4 is a standard fully actuated four-way intersection with stop bar and advanced detectors; the Danville intersection has in addition the departure lane detectors, the MicroRadar pedestrian detectors, and the camera. An important requirement for a HR system is to synchronously collect all detection events, and to transfer the data to a server for processing.

2. Timing plans

The typical procedure for designing re-timing plans in the U.S. is based on two sets of data. The first set comprises background information such as intersection geometry, constraints on cycle length, min and max green times for certain movements (e.g. left turn), pedestrian clearance times, and current timing plans. The second set comprises manually collected turning counts for three days (typically weekday 7 AM–7 PM, and Saturday). Pressure hose vehicle counts for one continuous week may also be collected. Finally there may be up to five 'floating car' runs for AM, midday and PM peaks, before the re-timing effort. The data is input to a software package like Synchro to obtain the re-timing plan (splits, offsets, etc.). The next step is to implement the new plan. A post-implementation analysis based on another set of floating car runs is then carried out. The last step is to produce a performance report showing improvements in level of service (LOS) and coordination, and reduction in delay and emissions. Large benefits are invariably reported.

Lee County in Florida re-timed 50 intersections at a consultancy cost of $357,400 to achieve annual estimated benefits that include $15.3 M in time saved, $2 M in fuel cost reductions and 19% reduction in emissions valued at $.12 M (ITS International, 1/4/2013). Earlier examples of re-timing are recalled in Sunkari (2004), which estimates the cost of re-timing one intersection at $2500–$3500 (2004 prices) and a benefit-cost ratio of 40. The study cautions that the benefits of re-timing vanish as the traffic pattern changes. Ironically, the reported 40:1 benefit-cost ratios confirm that performance degrades greatly between infrequent re-timing plans. The degradation could be eliminated by plans that quickly react to changes in traffic patterns.

2.1. Traffic variability

We use data from two intersections in Danville, CA and Beaufort, SC, with HR systems that have been operational for more than two years.

2.1.1. Seasonal and daily variation

Fig. 2 shows for a particular movement the median 15-min counts (normalized to veh/hour) for Tuesday, Wednesday, Thursday for six months. There is a large increase in counts at 8 AM and 9 AM during February, March and September

---

[3] For interpretation of color in Figs. 1, 4, 10 and 11, the reader is referred to the web version of this article.

and a big reduction in July, August and December when a nearby school is not in session. Clearly, this intersection should have different timing plans for those three months. Traffic counts collected in a single week would not reveal this need.

Just as traffic varies from month to month, there is daily variation as well. Knowledge of this variation can be used to improve timing plans. From individual vehicle departure times during each phase we can measure the saturation flow rate for that phase. Hence we can accurately calculate the volume to capacity (VC) ratio for each phase: VC Ratio = Count/Green time × Saturation rate. Fig. 3 displays the 80th percentile VC ratios using data for 2014 for two (out of six) phases over 15-min intervals for Tuesday, Wednesday and Thursday, together with the level of service (LOS). The plot of Phase 4 shows the impact of the pickup times of the school near the intersection. The two plots immediately suggest "transferring" some green time from Phase 8 to Phase 4 around 5 PM to improve LOS of Phase 4 from C, while maintaining the LOS at Phase 4 of A.

2.2. Automatic selection of timing plans

The standard approach to designing TOD plans begins with measurement of average counts over three days, selects days of the week that will have different plans (e.g. weekday and weekend), and for each day, selects time intervals with different plans (e.g. AM and PM peak and off-peak periods). The selection of the days and time intervals is based on the traffic engineer's or consultant's intuitive understanding of the traffic patterns. If we have continuous measurements for one year (say), we could cluster the daily data to reveal the days of the week with significantly different traffic patterns, and then cluster the intra-day data to divide the day into periods with significantly different traffic. We illustrate the procedure using hourly data for an intersection in Beaufort, S.C. (see Fig. 4 (left)), for counts of the 12 movements (4 approaches, 3 movements per approach) for 164 days, Dec 2014 to May 2015. This represents each day's traffic by a 24×12 vector of hourly counts. We use a standard k-means clustering algorithm to partition these 164 vectors into k clusters, $k = 1, \ldots, 6$. For each k the algorithm partitions the vectors into k groups, $G_1, \ldots, G_k$ to minimize

$$\sum_{i=1}^{k} \sum_{v \in G_i} |v - \mu_i|^2.$$

Here v denotes a vector of hourly 12-movement counts, and μ_i is the average of the vectors in G_i.

The plot on the left in Fig. 5 shows the goodness-of-fit as the number of clusters is increased from 1 to 6. The fit steadily improves up to four clusters: cluster 1 corresponds to weekdays M–Th, clusters 2, 3, 4 correspond respectively to Fri, Sat, Sun. Cluster 6 splits up the M–Th cluster, but not in any revealing manner. So we conclude that there should be a separate TOD plan for M–Th, Fri, Sat, Sun, which is indeed what is currently done in this intersection. The other four plots show the 12 movement counts of the mean μ_i of the four best clusters.

The next step is to take each of the four 'day' clusters $G_1, \ldots, G_4$ and divide each 24-h day into a number of disjoint time intervals, $T_1, \ldots, T_m$, with $\cup_i T_i = [0, 24]$. For each m we select the intervals to minimize the sum of squares,

$$\sum_{i=1}^{m} \sum_{t \in T_i} |\mu_i(t) - \bar{\mu}_i|^2.$$

Here $\mu_i(t)$ is the mean of a day cluster G_k for hour t and $\bar{\mu}_i$ is its average in the interval T_i. Obviously, this 'fit' improves with the number of intervals m, similarly to Fig. 5, and we select m appropriately. Fig. 6 shows the result of this procedure for the four clusters. A good design is to select 7 TOD (time of day) plans for M–Th (beginning at 0:00, 5:45, 7:10, 8:40, 14:40, 18:15 and 20:25), 8 plans for Fri, 4 plans for Sa, and 3 for Su. The figure shows the corresponding time intervals, for each of which we must design an individual timing plan (splits and cycle time).

We design this plan for interval T_i taking $\bar{\mu}_i$ to be the average volumes of the 12 movements, and we compare the result with the current splits. We calculate the 'optimum' splits and cycle time via a quadratic program that seeks to equalize the VC ratio for all 12 movements, constrained by a specified maximum VC ratio and min and max green. However, in the example we fix the cycle time to 130 s for comparison with the current practice, instead of using the optimum cycle time. Fig. 7 displays the two sets of splits arranged in a dual-ring diagram. The middle plot shows the measured VC ratios. The VC ratio for phase 8 is very low and the optimum split assigns a much smaller green time to this phase and a correspondingly larger time to phase 4. This makes intuitive sense since phase 4 has the largest flow. However, forming this intuition itself requires HR data. Instead of equalizing VC ratios we could have minimized delay expressed by the HCM formula.

2.2.1. Note on Synchro

Synchro is a software package used to design timing plans for a particular period. The input comprises average counts for all movements for that period, e.g. AM peak. Synchro accounts for variability in the arrivals simply by inflating the average rate λ vph to the 90th percentile, assuming the arrivals are Poisson. By making a Gaussian approximation to the Poisson distribution, the 90th percentile is $\lambda + 1.6\sqrt{\lambda}$. The right hand side of Fig. 4 shows histograms of inter-arrival times for 11 out of 12 movements (the missing movement from leg 3, lane 1 does not have an advanced detector and is not considered). Superimposed on the histogram is the exponential distribution in red. The numbers below the x-axis is the average inter-arrival time λ^{-1} in seconds. The histogram is close to the exponential for small values of λ, which is not surprising. However, for the right-turn movement from lane 4 of leg 2, which has the largest rate, the exponential is a poor fit. The exponential

Median turn movement count by month- EB, Left turn lane
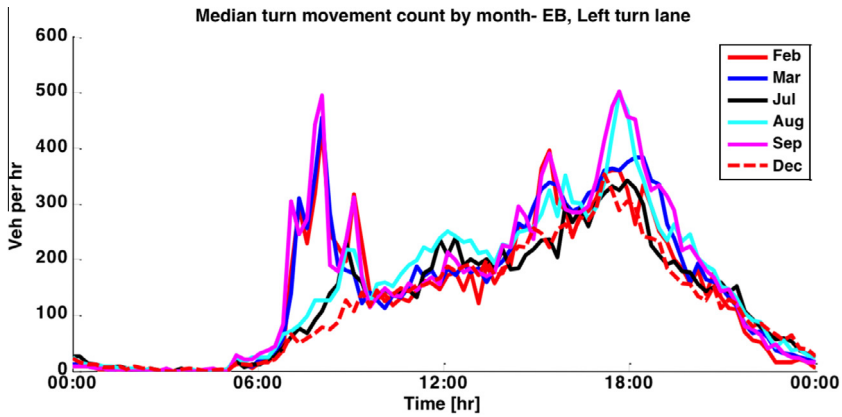

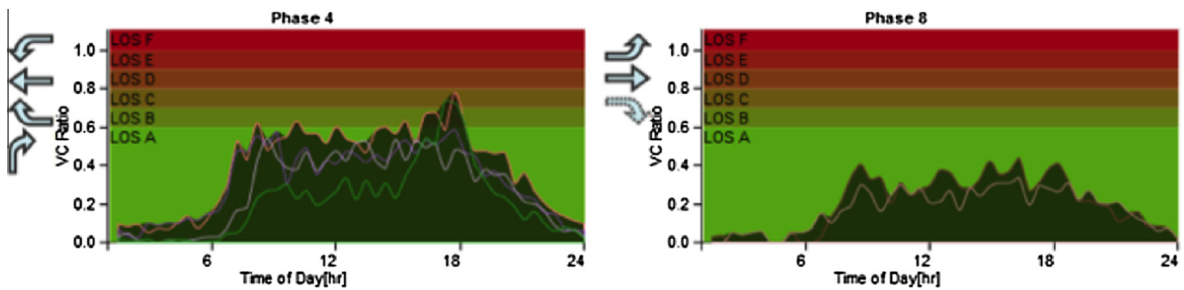Fig. 2. Seasonal trend in count in Danville, CA.

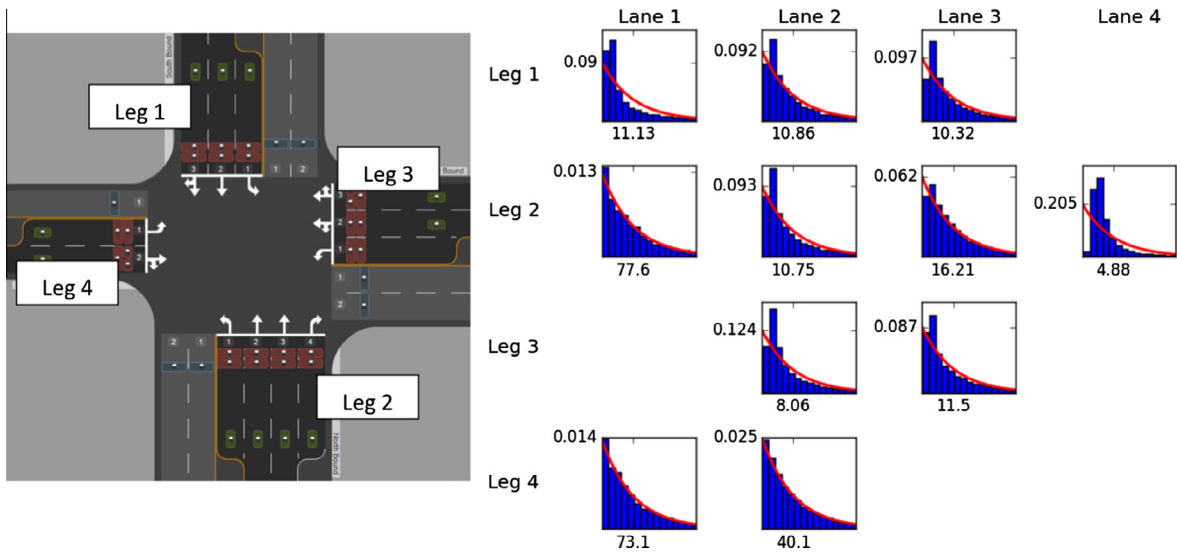Fig. 3. VC ratios and LOS for two phases in Danville, CA.

Fig. 4. Detector layout for intersection in Beaufort, SC (left): each small white dot is vehicle detector. Histograms of inter-arrival times for different movements (right).

distribution has the largest values of inter-arrival times near 0, as well as a large tail, both of which are not physically meaningful; the histogram shows no inter-arrival times that are very short or very long. Hence for this case the empirical 90th percentile is much smaller than Synchro's estimate of $\lambda + 1.6\sqrt{\lambda}$. Synchro would assign a larger split to this movement (as well as a larger cycle length) than what HR data indicate is required.
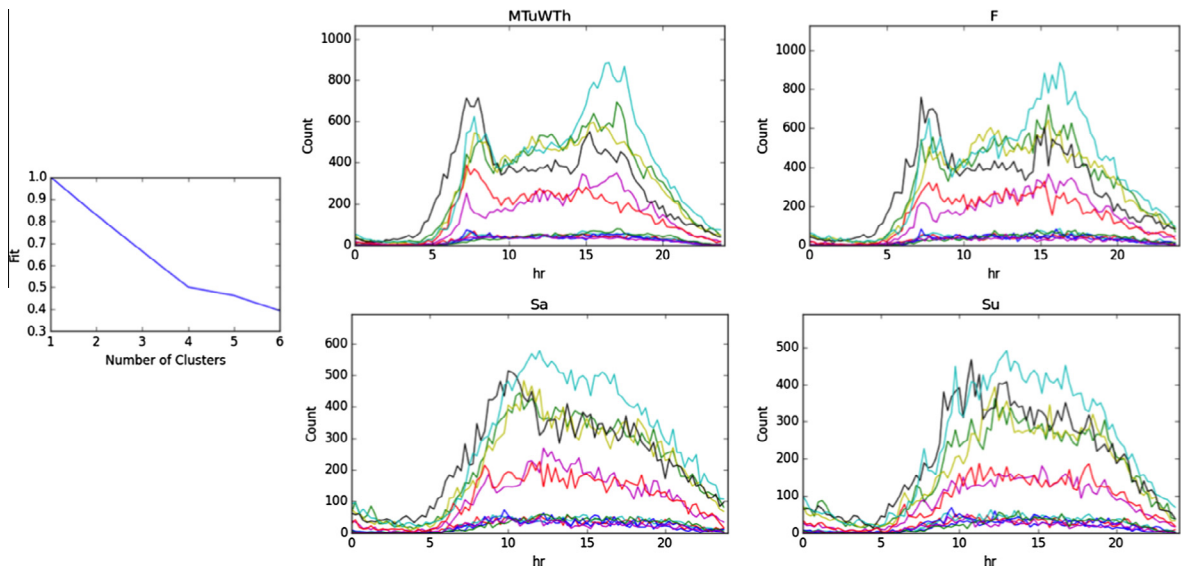
Fig. 5. Clustering of daily data for Dec 2014 to May 2015 in an intersection in Beaufort, SC.
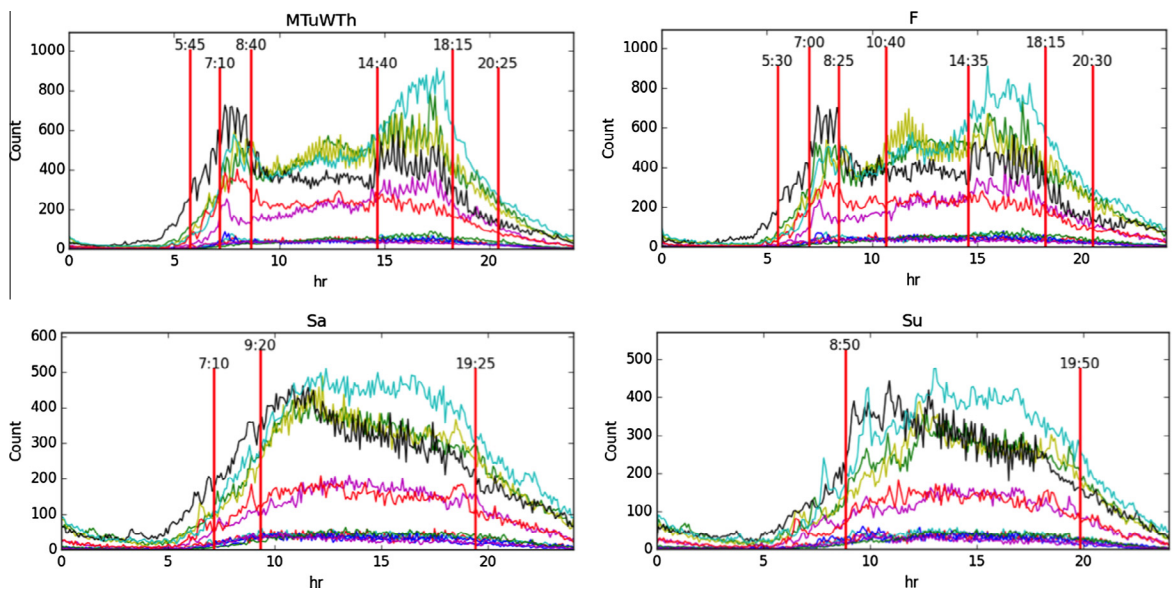

Fig. 6. Clustering of hourly data for Dec 2014–May 2015 in an intersection in Beaufort, SC.

The procedure described above to design the splits for time interval T_i uses only the average counts $\bar{\mu}_i$ in that interval. But the HR system provides *real time* counts, which one could use to *predict* future volume. These predictions could be used to adapt the splits to take the predicted traffic into account. We use a principal component analysis (PCA) for prediction. To illustrate the idea consider the North–South through movement (see Fig. 4), averaged over 15 min, for an entire day for 164 weekdays during Dec 2014–May 2015. Each day d gives a $96 = 4 \times 24$ dimensional vector, $v(d)$, from which we subtract the average $\bar{v} = \frac{1}{164} \sum_d v(d)$. We perform a PCA of the vectors $x(d) = v(d) - \bar{v}$. Fig. 8 displays the average, $\bar{v}$, and the top four principal components, as well as the PCA weights of each component, which rapidly decrease beyond four. Denote these components by x^1, x^2, x^3, x^4. Then the count vector $v(d)$ for any day d is represented as

$$v(d) = \bar{v} + w_1(d)x^1 + w_2(d)x^2 + w_3(d)x^3 + w_4(d)x^4 + \epsilon(d), \tag{1}$$

where the $w_i(d)$ are the coefficients of the orthogonal projection of $v(d) - \bar{v}$ on the 'basis' vectors x^1, x^2, x^3, x^4 and $\epsilon(d)$ is the residual.
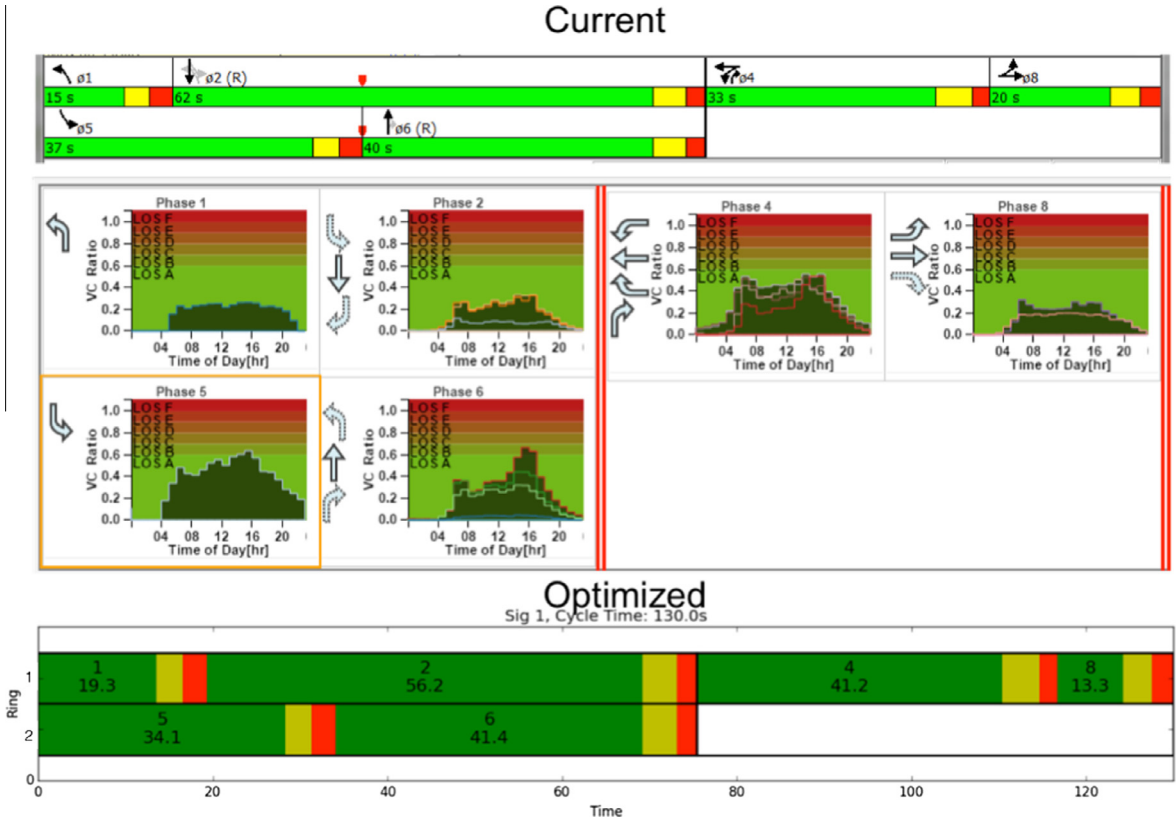
Fig. 7. Current splits and measured VC ratios (top, middle) and 'optimum' splits (bottom).

Principal components x^1, x^2, x^3 have an intuitive interpretation. x^1 is positive throughout and shows AM and PM peaks. So if $w_1(d) > 0(< 0)$, we expect an above (below) average AM and PM peak on day d. On the other hand, x^2 shows an AM and a PM peak and negative values for mid-day. So if $w_2(d) > 0(< 0)$, we expect below (above) average volume during mid-day. Similarly if $w_3 > 0(< 0)$, we expect below (above) average volume after 6 PM, and if $w_4 > 0(< 0)$ we expect a sharp increase (decrease) around 8 PM. One simple way to use these predictions is to change the start times of the fixed time of day plans. For example, if on a Monday we predict an above normal AM peak, we could start that plan earlier than the nominal time of 7:10 AM suggested in Fig. 6.

Of course, we don't know the weights, $w_i(d)$, in (1) at the beginning of the day. One idea we are investigating is this. Suppose at time t on day d, we have measured $v(d)_t = (v(d)(\tau), 0 \leqslant \tau \leqslant t)$. We project $v(d)_t$ on the (truncated) principal components $x_t^1, \ldots, x_t^4$:

$$v(d)_t = \bar{v}_t + w_1(d)_t x_t^1 + w_2(d)_t x_t^2 + w_3(d)_t x_t^3 + w_4(d)_t x_t^4 + \epsilon(d)_t, \qquad (2)$$

and we use the coefficients $w_i(d)_t$ in (1) to predict $v(d)(\tau)$, for $\tau > t$. Clearly $w_i(d)_t \to w(d)$ as $t \to 24$. However, we need to conduct an empirical study to investigate how good the prediction is.

3. Other mobility performance measures

The Highway Capacity Manual specifies three intersection mobility performance measures: vehicle delay, number of stops, and level of service. Estimating these measures requires per movement counts, saturation rates, proportion of vehicles arriving on green, and volume-capacity ratio, but cities do not measure these. HR data can be used to compute these parameters. Day et al. (2014) give the most comprehensive set of performance measures, including cycle time, queue service time, phase durations, and phase failures, focusing on local control that determines intersection capacity allocation and on system control concerned with progression quality. The study describes procedures to calculate each measure and gives examples to calculate these measures using high-resolution data. We illustrate two of these performance measures using data from the intersection in Danville, CA.
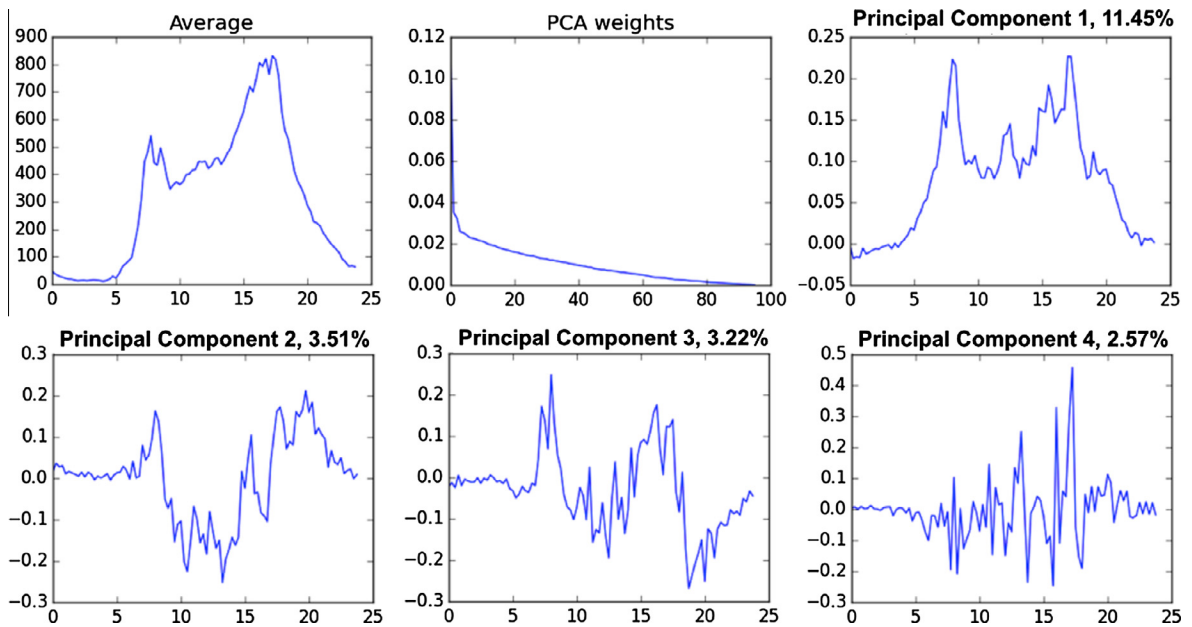
Fig. 8. Four PCA components of the North–South through movement and the average. The *x*-axis is hours.

3.1. Wasted green and max waiting time

Fig. 9 (left) is the histogram of wasted green in phase 4 within a cycle, defined as the time during which this phase is actuated and there is no vehicle present, but there is a vehicle waiting for the actuation of another phase. The histogram on the right is the maximum time a vehicle is waiting at the stop bar for the phase to be actuated. The histograms show that the median wasted green time per cycle is 6 s, and the median max wait time is 20 s; about 10% of cycles inflict a very large wait time, which may cause driver complaints.

3.2. Purdue coordination diagram (PCD)

The HR time series can be used to measure the quality of traffic progression at an intersection, and subsequently for signal offset re-timing calculations in order to coordinate signals at multiple traffic intersections to ensure smooth progression.

Fig. 10 is the Purdue Coordination Diagram (PCD) (Bullock et al., 2008). It is a visual representation of the quality of traffic progression through the Danville intersection. It displays the histogram and the cumulative distribution function (CDF) of vehicle arrivals at the intersection relative to the beginning of the green signal for a chosen time interval (in this case, the morning peak, 7 am–10 am). The plots immediately reveal the percent arrivals on green and vehicle waiting time – two important metrics characterizing traffic progression. The PCD shows that 60 percent of vehicles from the East and 40 percent from the West arrive on red and have to wait at the intersection.

4. Safety

We describe applications of HR data from the Danville intersection to assess moving violations, near crashes, and pedestrian–vehicle interaction.

4.1. Moving violations

Measurements of a vehicle's movement, its speed, and the signal phase, allows one to determine if its driver is making a moving violation. Three violations compromise safety: the vehicle is speeding, running a red light, or making a right turn on red (RTOR) without stopping.

Fig. 11 plots vehicles that leave the stop bar around the time the signal transitions through yellow before turning red. (The yellow signal in the Danville intersection lasts 3 s.) Vehicle speeds are plotted against the time of intersection entry, relative to the yellow signal start time. The speed measurements are quantized because of the detectors finite sampling frequency: the estimates are accurate to within ±3 mph for speeds under 20 mph but higher speeds have larger quantization errors. The red circles in Fig. 11 (top) capture five vehicles driven dangerously across the intersection at high speed within 3 s after the signal turns red—a rate of one violation per hour in the east direction. The high speed suggests that drivers find themselves in
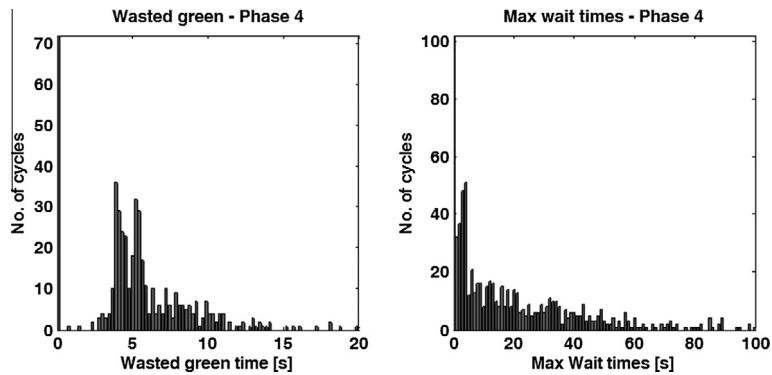
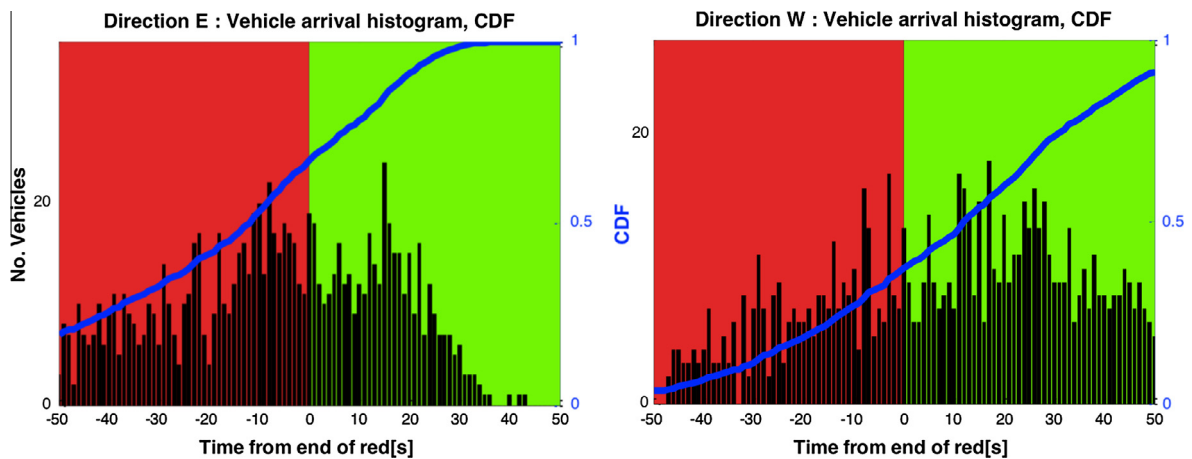**Fig. 9.** Histograms of wasted green and max waiting times.

Fig. 10. Histogram and cumulative distribution of vehicle arrivals.

a dilemma zone (Zhang et al., 2014). A right turn after stopping on red (RTOR) is allowed in the Danville intersection for periods when school is not in session. Fig. 11 (bottom) shows the 56 RTOR violations–several at high speed–amounting to a rate of 7 per hour.

Since we know the movement of all vehicles, one can determine whether these violations threaten another vehicle that legitimately enters the intersection from a conflicting direction at the same time. The Danville intersection has a camera that can be triggered by such simultaneous events to record a short video clip that includes the events. Fig. 12 shows three video frames from such a clip. The leftmost frame shows a vehicle entering the intersection during green, together with a red light violator who is quite far from the intersection. The HR data shows that the first vehicle entered at 5.8 s into its green and the violator entered at 5.9 s into its red phase, traveling at 30 mph. The high speed of the violator can also be inferred by comparing the vehicle locations in the left and middle frames. The first vehicle swerved to prevent a collision as seen in the rightmost frame. HR data can provide such automatic recording of 'near crashes' by appropriately defining trigger events. It may be possible to trigger a collision warning for a connected vehicle application.

Fig. 13 shows red light violations by time of day over one year. There is an extraordinarily high violation rate shortly after 3 PM, when the school lets out. Drivers could be encouraged to drive more carefully by having an officer stand at the corner for one or two weeks, accompanied by an advertising campaign. Subsequent data would reveal the effectiveness of such educational efforts.

4.2. Pedestrians

The Danville intersection has one crosswalk that is instrumented with several MicroRadar sensors that can detect pedestrians. A sensor emits a microwave pulse at a rate of 8 Hz and a range of six to eight feet. The strength of the reflected wave is proportional to the cross section of any object in the range. Fig. 14 shows the experimental site. MicroRadar sensors are located at the bars in the crosswalk and record when they detect an object (pedestrian) within range as indicated in the images. We summarize two of the findings of an experiment reported in detail in Muralidharan et al. (2014).
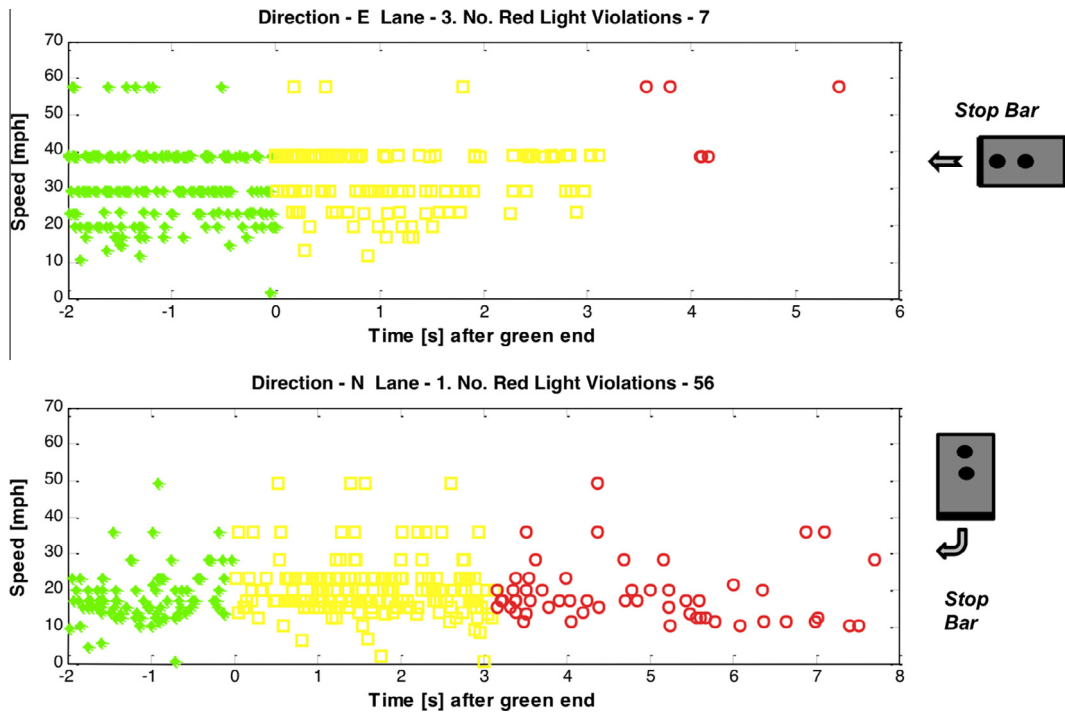
Fig. 11. Vehicle speeds during signal phase transitions over an 8-h period on November 18, 2011. Each marker corresponds to one individual vehicle.

red light violator

Fig. 12. A near crash between a red light violator and a conflicting entry. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
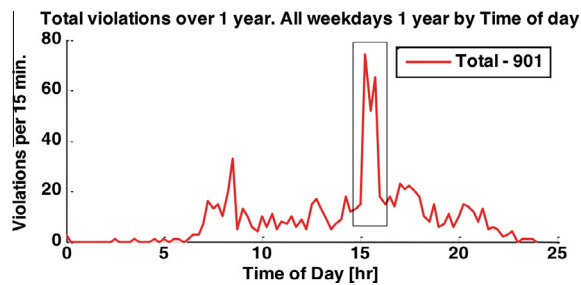

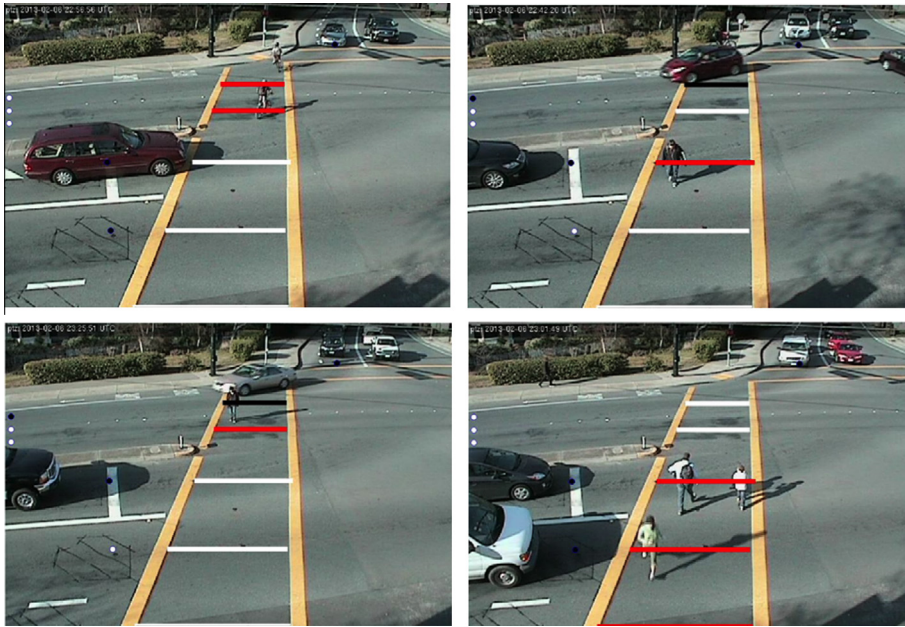**Fig. 13.** Trends in red-light violations.

Fig. 14. Five bars represent sections of the crosswalk covered by pairs of MicroRadars. White bars denote no events, Red bars denote pedestrian detection and Black bars denote vehicle detection events. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
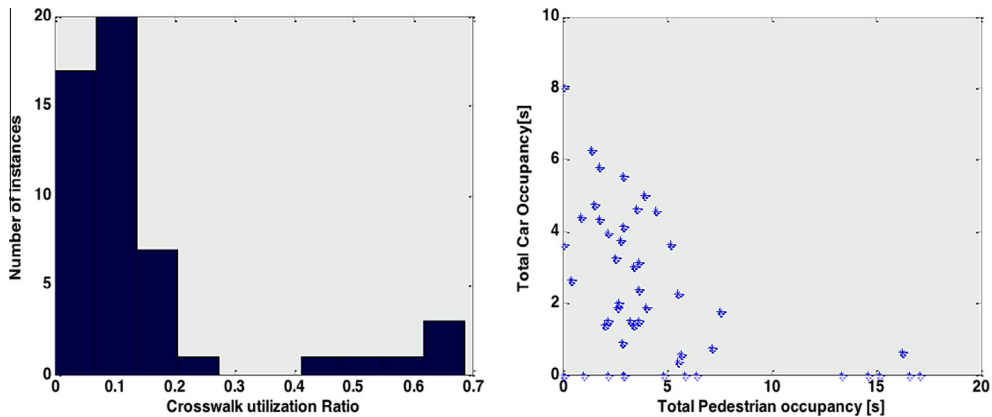

Fig. 15. Crosswalk utilization ratio across different pedestrian cycles (left). Scatter plot of total pedestrian and car occupancy times during pedestrian signal active times (right).

At this intersection the pedestrian walk signal is actuated during each cycle. Fig. 15 displays an analysis of the detector data during several walk cycles. The plot on the left is a histogram of the crosswalk utilization ratio, defined as

$$\text{Crosswalk utilization ratio} = \frac{\text{Total pedestrian occupancy during ped signal time}}{\text{Total ped signal time}}.$$

The numerator is the time during which any MicroRadar detected an object while the ped signal was active, and the denominator is the total time for which the ped signal was active. The histogram shows that in most ped actuation times no pedestrians used the crosswalk, while during a few cycles, the crosswalk was heavily used. This suggests that it may be worth considering replacing the fixed ped actuation by a push-button that actuates the signal on demand.

On the right of Fig. 15 is a scatter plot of the total pedestrian occupancy times and total car occupancy times over different pedestrian cycles observed during the experiment. Such a plot can be used to quantify pedestrian–vehicle conflicts in the crosswalk. Conflict-free situations correspond to points lying on the x-axis (when no cars are present) or the y-axis (when no pedestrians are present). Points away from the x-axis and y-axis represent increasing levels of conflicts. These plots

can also be generated for individual MicroRadars, providing further insight into possible conflict hotspots within the crosswalk.

5. Conclusions

Inefficient and unsafe operation of intersections is costly, increasing delay, emissions and injuries. In rapidly growing cities these costs will increase rapidly. They will also increase as cities try to provide multimodal transportation choices to all citizens, whether it's by walking, bicycling, transit, or driving. Unless well managed, the commitment of more road surface to pedestrians and cyclists will cause greater congestion and increase pedestrian–bicycle–vehicle conflicts and accidents. But careful management of urban roads requires the capacity to evaluate, predict and identify means to improve intersection mobility and safety. Development of such a management capacity requires a high-resolution (HR) data system that collects and archives real time measurements of every vehicle in the intersection. Without HR data, management will continue to operate in the dark.

The paper reports analyses of data from HR systems at two intersections. The analysis shows trends in seasonal and daily patterns of traffic that cannot be discerned from spot measurements made every few years. A wide range of performance measures can be obtained on a continuing basis. These measures include standard VC and LOS measures, as well as measures of waiting times, wasted green and progression quality. Safety-relevant measures include frequencies of red-light, speeding and right-turn-on-red violations. If pedestrian and bicycle detectors are available, conflicts with vehicles can also be estimated.

Changes in traffic patterns immediately suggest re-timing plans that adapt to those changes. The paper has focused on fixed time control because of its ubiquity, but HR data is even more valuable in the evaluation and implementation of actuated and adaptive control. The hardware needs for a HR data system are not significantly more than what is needed for a fully actuated intersection. The HR system described in the paper can be readily enhanced to take advantage of communications with connected vehicles via a DSRC modem. The modem can in turn be employed to send messages to vehicles to warn them of approaching vehicles or to communicate impending changes in the signal.

It may be worth noting that every application presented in the paper requires HR data. The applications summarized in Figs. 3, 8–15 require event-by-event data; those summarized in Figs. 1, 2, 4–7 require archived time series on the order of one year. Most commercial controllers do not make such HR data available. Secondly, the applications provide information that is actionable, suggesting re-timing plans, modification of signal control to reduce wasted green or improve coordination, and reduction of safety hazards from red-light violations. Lastly, the HR system described here is very practical: its cost is comparable to the sensor requirements of a fully actuated intersection; at the same time it automates the data analysis.

The analysis of the rich HR data presented here suggests other uses, of which we mention three. First, the data can be used to estimate queue lengths as done by Lee et al. (2015). Important signal control algorithms rely on good queue length estimates (Aboudolas et al., 2009; Varaiya, 2013). Second, Coogan (2015) significantly extends the PCA analysis both to predict traffic patterns and to use the prediction to proactively change timing plans. Third, in the context of an arterial corridor, the data can be used to optimize offsets (Coogan et al., 2015).

References

Aboudolas, K., Papageorgiou, M., Kosmatopoulos, E., 2009. Store-and-forward based methods for the signal control problem in large-scale congested urban road networks. Transp. Res. Part C 17, 163–174.

Bullock, D.M., Day, C.M., Haseman, R., Brennan Jr., T.M., Wasson, J.S., Sturdevant, J.R., 2008. Evaluation of arterial signal coordination: methodologies for visualizing high-resolution event data and measuring travel time. Transp. Res. Rec. 2192, 37–49.

Coogan, S., 2015. History's principal components repeat themselves: using data for traffic predictive control. In: Presented at IPAM Workshop on Traffic Control, UCLA, October 27, 2015.

Coogan, S., Gomes, G., Kim, E.S., Arcak, M., Varaiya, P., 2015. Offset optimization for a network of signalized intersections via semidefinite relaxation. In: IEEE 54th Conference on Decision and Control (CDC), pp. 2187–2192.

Day, C.M., Bullock, D.M., Li, H., Remias, S.M., Hainen, A.M., Stevens, A.L., Sturdevant, J.R., Brennan, T.M., 2014. Performance Measures for Traffic Signal Systems: A Outcome-Oriented Approach. Technical Report. Purdue University, Lafayette, IN. http://dx.doi.org/10.5703/1288284315333.

Dowling, R., Ashiabor, S., 2012. Traffic Signal Analysis with Varying Demands and Capacities. Draft Final Report Technical Report NCHRP 03-97. Transportation Research Board.

ITS International. Benefits of Florida's Traffic Signal Retiming, 1/4/2013.

Lee, S., Wong, S.C., Varaiya, P., 2015. Stage-based max pressure adaptive traffic control policy at isolated signalized junctions. In: 94th Annual Meeting of the Transportation Research Board, Washington, D.C., USA.

Muralidharan, A., Flores, C., Varaiya, P., 2014. Pedestrian detection and real-time performance measures at intersections. In: 93rd Annual Meeting of the Transportation Research Board, Washington, D.C., USA.

National Traffic Operations Coalition, 2007. National Traffic Report Card, 2007. <http://www.ite.org/REPORTCARD/> (accessed 14.08.08).

National Traffic Operations Coalition, 2012. National Traffic Report Card, 2012. <http://www.ite.org/reportcard/>.

New York DOT, 2010. The New York City Pedestrian Safety Study & Action Plan. <http://www.nyc.gov/html/dot/html/about/pedsafetyreport.shtml>.

NHTSA, 2010. Crash Factors in Intersection-related Crashes: An On-scene Perspective. Technical Report DOT HS 811 366. National Highway Traffic Safety Administration.

Pierce, G., Shoup, D., 2013. SFpark: pricing parking by demand. Access Mag. (Fall), 43

Sunkari, S., 2004. The benefits of retiming traffic signals. ITE J., 26–29

Varaiya, P., 2013. Max pressure control of a network of signalized intersections. Transp. Res. Part C 36, 177–195.

Zhang, Y., Fu, C., Hu, L., 2014. Yellow light dilemma zone researches: a review. J. Traffic Transp. Eng. (English Edition) 1 (5), 338–352.