Coping with Real-World Challenges in Real-Time Urban Traffic Control

Xiao-Feng Xie
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
xfxie@cs.cmu.edu

Stephen F. Smith
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
sfs@cs.cmu.edu

Gregory J. Barlow
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
gjb@cmu.edu

Ting-Wei Chen
College of Information
Liaoning University
Shenyang, Liaoning, China 110006
twchen@lnu.edu.cn

5995 words + 5 figures + 1 tables

November 9, 2013

ABSTRACT

In urban road networks, the use of real-time adaptive traffic signal control systems faces two typical challenges. First, various sources of uncertainty and disturbance can significantly degrade the accuracy of real-time flow predictions. Second, the optimization of vehicle flows must also give active attention to other transportation modes such as bus transit and pedestrian flows. In this paper, these challenges are investigated in the context of a recently implemented system called SURTRAC (Scalable URban TRAffic Control), which has now been running continuously in an actual urban environment for more than a year. SURTRAC takes a decentralized, schedule-driven approach to real-time traffic control and its design aims at urban (grid-like) networks with multiple, competing dominant flows that shift through the day. Motivated by observations of the system in operation, several strategies are proposed for strengthening the basic SURTRAC algorithm to better deal with real-world uncertainties and disruptive events, as well as multi-modal traffic demands. We evaluate the effectiveness of these strategies using both simulations and analysis of data collected from the pilot deployment.

INTRODUCTION

Efforts to apply adaptive traffic signal control systems in urban environments face two typical challenges. The first challenge, which must be met in all application contexts, stems from their dependence on real-time information. To be effective in practice, a traffic control method requires quite accurate knowledge of traffic flows (*1, 2*). However, given that prediction of local traffic flows must be accomplished with a limited number of sensors (*1, 3, 4*), various sources of uncertainty can degrade the accuracy of flow prediction. The quality of detectors can be highly dependent on proper installation, and for some sensing technologies, performance is also influenced by dynamic environmental factors, e.g., weather and traffic conditions (*5, 6*). Mis-counting or over-counting can be caused by arbitrary lane-changing behavior of human drivers. In urban settings, there are additional uncertainties and disruptions. Road closures might trigger significant changes in traffic flow patterns. Temporary lane blockages might be caused by turning trucks, stopping buses, or on-street parking in progress. Unsignalized intersections with side streets and alleys can contribute additional undetected flows. These nontrivial perturbations are often ignored in existing work.

A second challenge in applying real-time adaptive signal control systems, which is more specific to urban road networks, is that passenger vehicles, transit, and pedestrians must share the right of way, and traffic signal control must give active attention to and be compatible with other travel modes. Improving performance of transit operations has attracted growing interest (*7, 8, 9, 10, 11*). Likewise, pedestrian flow is being given increasing priority in many urban environments (e.g., (*12*)) and there is increasing emphasis on multi-modal urban traffic environments (*13, 14, 15*), as compared to the vehicle-centric focus in conventional designs.

In this paper, we consider these challenges in the context of the SURTRAC adaptive traffic signal control system (*16, 17, 18*), which was deployed on a nine-intersection network in the East Liberty region of Pittsburgh, PA in June 2012. SURTRAC (Scalable Urban Traffic Control) implements a totally decentralized approach to real-time adaptive signal control, adopting much the same conceptual framework as promoted by earlier online-planning approaches to adaptive signal control (*2, 19, 20*). SURTRAC is distinguished by its use of a novel formulation of the intersection scheduling problem (*16*), which enables intersections to compute near-optimal schedules over an extended planning horizon in real-time. Local schedulers at each intersection in a controlled road network operate asynchronously in rolling horizon fashion, communicating outflows to their neighbors to increase visibility of future incoming traffic and achieve coordinated behavior. When necessary, additional coordination mechanisms are applied to adjust local schedules to compensate for mis-coordination with its neighbors (*17*).

An evaluation of the East Liberty deployment of SURTRAC performed shortly after installation indicated that significant performance improvement was achieved relative to the coordinated-actuated signal plans that were previously in place (*18*). Nonetheless, limitations due to various uncertainties, disruptions and multi-modal demands have become apparent as we have observed its continued operation, and over the past year a number of extensions have been investigated and incorporated to achieve more effective and stable operations, based on analysis and/or simulations using field data. Before discussing these *strengthening strategies*, we first summarize the SURTRAC approach to real-time adaptive traffic signal control and the initial pilot deployment.

SURTRAC - SCALABLE URBAN TRAFFIC CONTROL

In SURTRAC, the traffic signal control problem is formulated as a decentralized, schedule-driven process (*16, 17*). Each intersection is controlled independently by a local scheduler, which main-

tains a phase schedule that minimizes the total delay for vehicles traveling through the intersection and continually makes decisions to update the schedule according to a rolling horizon. The intersection scheduler also communicates outflow information implied by its current schedule to its neighbors, to extend visibility of incoming traffic and achieve network level coordination.

At the individual intersection level, the ability to consider real-time (second-by-second) variability of traffic flows is made tractable by a novel formulation of online planning as a single machine scheduling problem (*16*). Key to this formulation is an aggregate representation of traffic flows as *inflows*. Each inflow includes a sequence of jobs, where a job contains vehicles in close proximity that have the right of way in a given phase, over a limited prediction horizon. Each job can be represented as a triple, i.e., *<vehicle count, arrival time, departure time>*. These job sequences preserve the non-uniform nature of real-time flows while providing a more efficient *scheduling search space* than traditional time-tick based search space formulations. The scheduling problem is to construct an optimal sequence of all jobs that preserves the ordering of jobs along each inflow. A given sequence dictates the order in which jobs will pass through the intersection and can be associated with an expected phase schedule that fully clears the jobs in the shortest possible time, subject to basic timing and safety constraints. The optimal sequence (schedule) is the one that incurs minimal delay for all vehicles. This scheduling problem is solved using a dynamic programming process.

When operating within an urban road network, any local intersection control strategy without sufficiently long prediction horizon is susceptible to myopic decisions that look good locally but not globally. To reduce this possibility, network level coordination mechanisms are layered over SURTRAC's basic schedule-driven intersection control strategy.

As a basic protocol, each intersection sends its projected *outflows* to its direct neighbors (*17*). Given an intersection schedule, projected outflows to all exit roads are derived from models of current inflows and recent turning proportions at the intersection. Intuitively, the outflows of an intersection's upstream neighbors become its predicted non-local inflows. The joint local and non-local inflows essentially increase the look-ahead horizon of an intersection, and due to a chaining effect, a sufficiently long horizon extension can incorporate non-local impacts from indirect upstream neighbors. This basic coordination protocol is quite similar to that previously utilized in (*19*). One difference is that we assume asynchronous coordination, so that temporary communication failures can be mostly ignored. Furthermore, the optimistic assumption that is made is that direct and indirect neighbors are trying to follow their schedules. Normally, the optimization capability of the base intersection control approach results in schedules that are quite stable, given enough jobs in the local observation and large jobs (platoons) in the local and non-local observation. It is also the case that minor changes in the schedules of neighbors can often be absorbed, if there is sufficient slack time between successive jobs.

In practice, circumstances can cause schedules to change, in which cases mis-coordination can occur, especially for neighbor intersections that are very close together. To this end, additional coordination mechanisms are incorporated into SURTRAC for handling specific nontrivial mis-coordination situations. One common inefficiency is caused by spillback that blocks the progress of traffic flow from an upstream intersection. The basic coordination protocol is augmented with a mechanism that acts to detect and prevent unnecessary spillback in advance of its occurrence by accelerating phase changes. Another source of mis-coordination is the tendency for the schedules of coordinating neighbors to oscillate due to small inconsistencies, which is handled by a second mechanism. Further description of these additional coordination mechanisms can be found in (*17*).

THE EAST LIBERTY DEPLOYMENT SITE

The SURTRAC system was installed on a nine-intersection road network in the East Liberty neighborhood of Pittsburgh, PA, (intersections **A**-**I** shown in Fig. 1) for initial field testing in June 2012. Although the total network size is not large, this road network has several interesting characteristics. First, in contrast to the arterial settings that are typically studied in most traditional systems, this network has more of a grid-like character. It contains a triangle where three major streets (Penn Circle, Penn Avenue, and Highland Avenue) cross, with changing traffic flows throughout the day. The network also consists of compact road segments. The lengths of roads between intersections range from 90 to 500 feet with an average of 272 feet, imposing a nontrivial challenge for achieving effective coordination in a decentralized traffic control system. Finally, like most urban environments, there are a range of uncertainties and disruptions stemming from bus movements, commercial deliveries, construction projects and pedestrian traffic.

Detection Capabilities

Several information sources provide inputs to SURTRAC. Both phase and pedestrian call status are extracted continually from each intersection controller. For any given intersection, phase start and end time points are used to ensure synchronization of the local scheduler, and are also communicated to neighboring intersections. The pedestrian call status is only used locally. A given pedestrian call ped_i corresponds to phase i in the controller, and remains pending until the pedestrian walk time is serviced in phase i.

Detector groups (realized as video camera sensing zones) are placed on both entry and exit road segments for reporting two fundamental measures: traffic counts and occupancy of vehicles. For each entry link, a group of *stop-bar detectors* is placed near the intersection, and a group of *advance detectors* is placed sufficiently far away from the intersection; For each exit link, a group of *exit detectors* is placed near the intersection. The advance detectors associated with a given intersection are typically the exit detectors of its upstream neighbors, if the intersection is not a boundary node in the overall adaptive network. Further details can be found in (*18*).

Since video detection is employed throughout the pilot test site, all of these detectors can be defined with no additional cost than would be required to support the detection requirement for vehicle-actuated logic and other traffic-responsive control techniques.

Modeling Inflows

Effective operation of the SURTRAC scheduler at a given intersection requires an accurate view of local vehicle inflows.

To obtain inflows, we rely on a set of local *link flow profiles*. Each link flow profile describes the state of *queuing* and *arriving* vehicles along a particular link in a high-resolution prediction horizon. The main challenge is to achieve accurate short-term estimation through the limited detection available in real-world settings (*1, 4*).

Link flow profiles are computed via an input-output technique (*3*), using advance and stop-bar detectors on each entry link. On each road, all vehicles recorded by advance detectors are assumed to be moving at constant free-flow speed (v_f) when they are not stopped, and the predicted arrival time of each vehicle is shifted by the horizon L/v_f, where L is the location of the advance detectors. Actual vehicle departures are recorded by stop-bar detectors. From time t_0 to t_1, each arriving vehicle is shifted by $(t_1 - t_0)$, and the queue q is adjusted by the difference of the number of predicted arrivals and the number of departures (*20*). For a given lane, a queue is discharged in
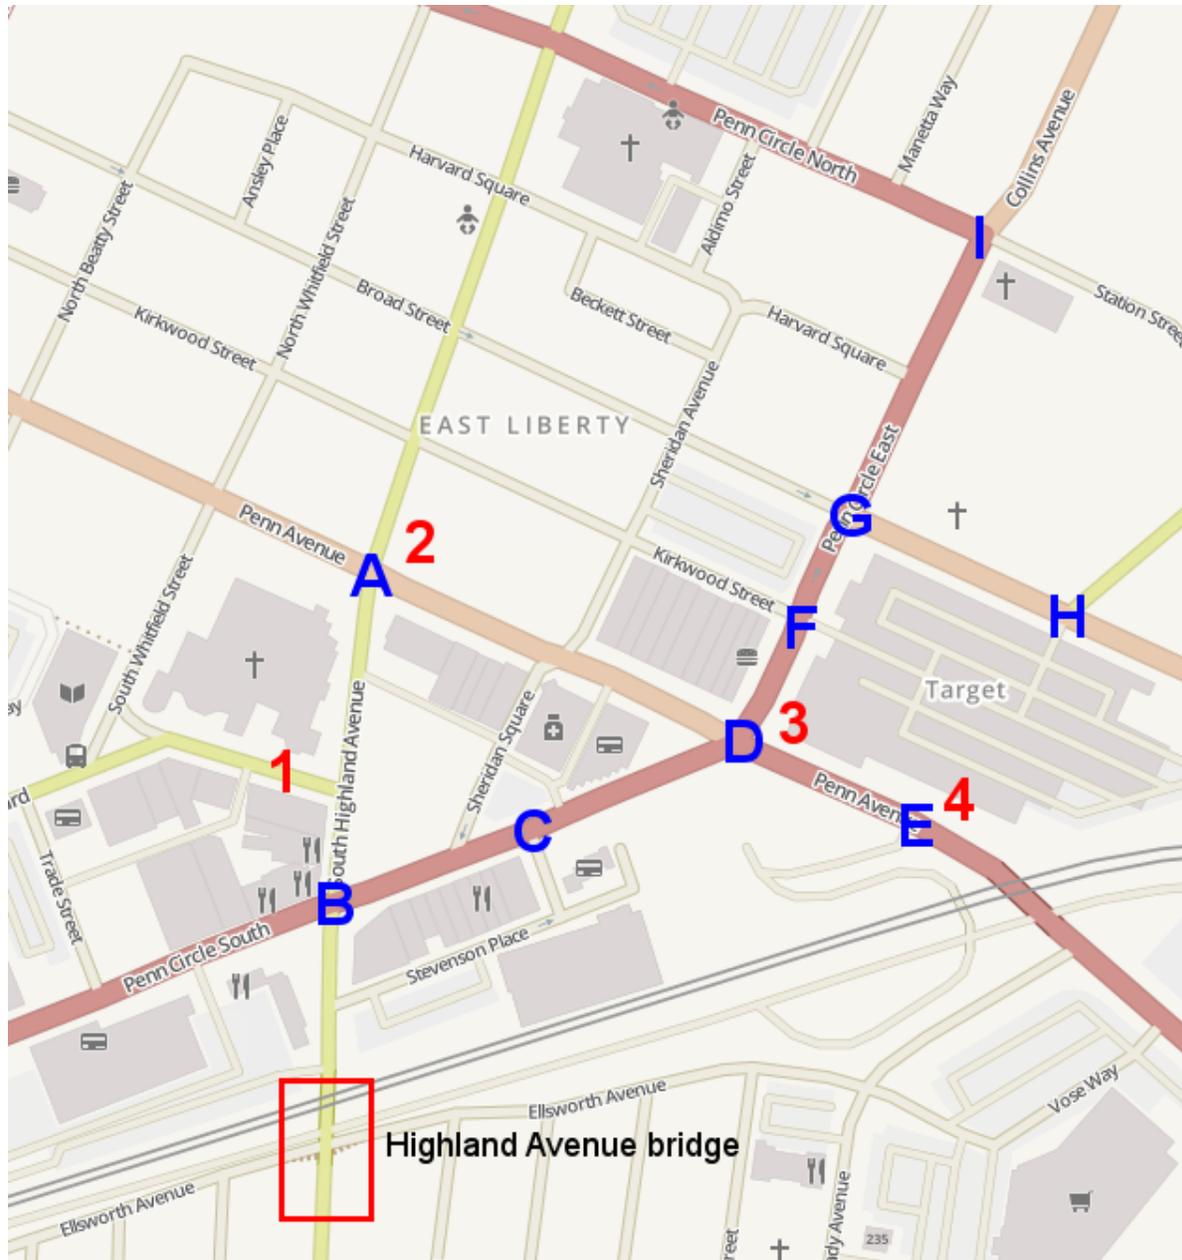
FIGURE 1 : The East Liberty pilot test site with nine signalized intersections (**A-I**). In urban environments, there are various issues faced by a real-time adaptive traffic signal control system. As a basic requirement, the system should be robust to significant changes in flow patterns over time (e.g., road closures, such as the closing of the bridge on South Highland Avenue from March to October 2013). In this paper, several issues are considered (examples noted on the map): (1) queue management for unaccounted traffic due to detection errors and mid-block traffic flows that are not covered by limited detection (detectors are placed near the intersections); (2) disruption management when the flow is temporarily blocked during green periods (e.g., bus stops with long dwell time periods near the stop bar); (3) flexible minor flow management to achieve more effective and stable operations; and (4) more active pedestrian flow management to limit the average wait time for pedestrians, without interrupting the progressing of major vehicular flows.

a green phase at the *saturation headway* (*shw*) after the *start-up lost time*. These model parameters are estimated from historical flow data.

Once equipped with link flow profiles, local inflows can be obtained by applying a road-to-phase mapping (*17*), based on turning movement proportions at the intersection.

This technique cannot detect queues beyond the advance detectors (*1, 21*). This does not cause a problem in our case, however, since non-local inflows are communicated from upstream neighbors in the road network.

System Performance and Robustness
A formal evaluation of the initial SURTRAC deployment, consisting of comparison of a series of "before" and "after" drive through runs, was carried out shortly after installation, and shown to yield significant performance improvement over the coordinated-actuated timing plans that pre-existed at the pilot site (e.g., 25% reduction in average travel times, 40% reduction in average wait times) (*18*). SURTRAC has been running continuously at the test site since its initial installment, and continues to provide a live testbed for further development (see below). In addition, high-resolution field data provides significant support for performance analysis (*22*).

Over the course of SURTRAC's ongoing deployment in East Liberty, traffic flow patterns have changed significantly. Between March 4 and October 23, 2013, the bridge on South Highland Avenue (location shown in Fig. 1) was closed for replacement, essentially cutting off one of the network's three major flows for more than six months. A large portion of traffic (including all bus lines) on Highland Avenue was then forced to pass through intersection **D**, which is a pivotal intersection that services for most vehicles in this road network (even before the event).

As evidence of the inherent robustness of the adaptive approach, Table 1 shows the flow difference for movements at intersection **D** before and after the closing of Highland bridge, both averaged over six weeks of collected data. Each movement is defined by the names of three intersections. As shown in the table, the total throughput in each day increased by 17.8% to 20.7% after the bridge closing. Individual flows CDE, EDC, EDA, ADE, EDF, FDE increased heavily (ranging from 14.9% to 48.0%), whereas CDF, ADF, and FDC actually decreased proportionally. After the change, CDE became the biggest movement, although EDF and EDA (which previously dominated) remained heavy. Although we have not been able to measure the effect that this change in flow patterns would have had on the pre-existing coordinated-actuated timing plans, we suspect that it would have resulted in a significant "aging" problem.

In addition, Table 1 provides basic flow information for discussions in the following sections, since intersection **D** is a pivotal intersection for most vehicles, in this urban road network.

STRENGTHENING STRATEGIES
As we have observed operations of the SURTRAC deployment over the past year, we have detected specific sub-optimal behaviors due to various uncertainties, disruptive events and multi-modal demands. This has led us to investigate a set of strengthening strategies and incorporate them into the baseline system. These strategies are summarized and analyzed in the subsections below.

Queue Management
For any intersection entry link, the queue is a hidden state that changes over time. Given various forms of sensing uncertainty (e.g., detection errors or hidden flows contributed by mid-block roads), the estimated queue at any point can deviate quite significantly from the actual queue. To

TABLE 1 : Flow information at intersection **D** before and after the closing of Highland bridge. Most vehicles traveling in this road network pass through this intersection. Between March 4 and October 23, 2013, the bridge on South Highland Avenue (location shown in Fig. 1) was closed for replacement, essentially cutting off one of the network's three major flows for more than six months. The road closure forced a large portion of traffic (including all bus lines) on Highland Avenue to pass through intersection **D**. Such significant changes in flow patterns often impose a challenge for the robustness of traffic signal control systems.

(a) Vehicle counts of all movements before the bridge closing (1/21/2013 — 3/3/2013)

	EDF	EDA	EDC	FDC	FDE	ADE	ADF	CDE	CDF	Total
Mon	3,814	4,003	1,627	2,191	868	3,936	799	3,638	3,294	24,171
Tus	4,289	4,232	1,731	2,337	888	4,163	795	3,721	3,376	25,530
Wed	4,191	4,190	1,758	2,316	940	4,335	782	3,839	3,522	25,872
Thu	4,359	4,556	1,871	2,351	920	4,513	816	3,947	3,436	26,769
Fri	4,405	4,740	1,883	2,577	996	4,640	869	4,313	3,696	28,119
Sat	4,102	4,226	1,836	2,398	965	4,110	902	3,971	3,273	25,782
Sun	3,281	3,259	1,486	1,939	827	3,236	690	3,031	2,636	20,385
avg	4,063	4,172	1,742	2,301	915	4,133	807	3,780	3,319	25,232

(b) Vehicle counts of all movements after the bridge closing (3/10/2013 — 4/13/2013)

	EDF	EDA	EDC	FDC	FDE	ADE	ADF	CDE	CDF	Total
Mon	4,715	4,922	2,112	2,223	995	4,889	762	5,294	3,099	29,011
Tus	4,907	5,265	2,285	2,204	1,045	5,129	739	5,563	3,221	30,357
Wed	4,982	5,327	2,287	2,258	1,079	5,129	750	5,636	3,233	30,681
Thu	5,041	5,575	2,381	2,336	1,085	5,476	773	5,926	3,280	31,873
Fri	5,386	5,931	2,501	2,457	1,134	5,528	779	6,161	3,260	33,137
Sat	4,952	5,236	2,298	2,348	1,106	5,184	848	5,943	3,115	31,029
Sun	3,945	3,995	1,975	1,854	912	4,134	676	4,631	2,490	24,612
avg	4,847	5,179	2,263	2,240	1,051	5,067	761	5,593	3,100	30,100

(c) Percentage increase in vehicle counts after the bridge closing

	EDF	EDA	EDC	FDC	FDE	ADE	ADF	CDE	CDF	Total
Mon	23.6%	23.0%	29.8%	1.4%	14.7%	24.2%	-4.6%	45.5%	-5.9%	20.0%
Tus	14.4%	24.4%	32.0%	-5.7%	17.6%	23.2%	-6.9%	49.5%	-4.6%	18.9%
Wed	18.9%	27.1%	30.1%	-2.5%	14.8%	18.3%	-4.1%	46.8%	-8.2%	18.6%
Thu	15.7%	22.4%	27.3%	-0.6%	17.9%	21.4%	-5.3%	50.1%	-4.5%	19.1%
Fri	22.3%	25.1%	32.8%	-4.7%	13.9%	19.1%	-10.3%	42.9%	-11.8%	17.8%
Sat	20.7%	23.9%	25.1%	-2.1%	14.6%	26.2%	-6.0%	49.6%	-4.8%	20.4%
Sun	20.2%	22.6%	32.9%	-4.4%	10.3%	27.8%	-2.1%	52.8%	-5.5%	20.7%
avg	19.3%	24.1%	29.9%	-2.7%	14.9%	22.6%	-5.7%	48.0%	-6.6%	19.3%

cope with this circumstance and improve the accuracy of queue estimation, we have developed and integrated a more sophisticated queue management strategy.

We first define some basic states that can be obtained from stop-bar and advance detectors (*3, 20, 21, 23*). The *queue spillback* state QS at each link n is recognized if the occupancy ratio on any advance detectors is nearly 100% for t_{QS} seconds. Let $l_{n,i}$ be the lanes of each entry link n serviced in phase i. The *queue blocking* and *queue clearance* states $QB_{n,i}$ and $QC_{n,i}$ are respectively obtained if the occupancy ratios of all stop-bar detectors on $l_{n,i}$ are nearly 100% for t_{QB} seconds and 0% for t_{QC} seconds. By default, $t_{QS} = t_{QB} = 3 \cdot shw$, and $t_{QC} = 2 \cdot shw$.

With these state definitions in mind, the basic queue length calibration is realized as follows. For each entry link, the predicted queue size q is a typical hidden state. A basic adjustment is $q = \max(0, q)$, and $q = \min(N_{SC}, q)$, where N_{SC} is the link storage capacity. The queue size is also adjusted from the measured states: $q = N_{SC}$ if QS and QB flags are identified, as the link is not serviced; and $q = 0$ if QC is obtained, and the link is serviced. The adjustment by spillback detection is especially important for video detection under high traffic volumes, where the count accuracy may deteriorate substantially due to the difficulty in sensing gaps between vehicles.

Some properties can be obtained from these basic adjustments. If q is over-estimated, an ideal full clearance will adjust q to 0. For each intersection, it costs an extra t_{QC} seconds. For its neighbors, the estimation error imposes some uncertainty on weights and execution durations of jobs in non-local inflows, and might cause some disturbance of the coordination between intersections. From the viewpoint of robust scheduling, any over-estimation might be considered as a buffer insertion (*24*), and since SURTRAC with the rolling horizon scheme is essentially doing on-line rescheduling continuously, it can effectively respond to such unexpected dynamics and provide good stability guarantees. If q is under-estimated, however, significant delay can occur from residual queues, and these residual queues will not be seen in subsequent cycles. The situation can become significantly worse if the queue starts to spill back to upstream intersections. Thus under-estimation should normally be avoided, although arriving vehicles in the look-ahead horizon will sometimes alleviate most negative impacts on pure queue clearance if the current phase is further extended to accommodate arriving platoons.

For each entry link n, we have the link flow rates fr_n^{arr} and fr_n^{dep} respectively from the groups of advance and stop-bar detectors. Each average flow rate is updated every 300 seconds. The *arrival/departure ratio* is defined as $ADRatio_n = fr_n^{arr} / fr_n^{dep}$.

The *arrival-adjusting* strategy is used to account for general arrival inaccuracy. As in (*3*), we assume that the group of stop-bar detectors can yield a reasonably accurate estimation of departing vehicles. If $ADRatio < 1$, some arriving vehicles are missed, and the numbers of queuing and arriving vehicles are under-estimated. Thus, when vehicles are detected at the advance detectors, the count is divided by $ADRatio$ to reclaim those missing vehicles in the link arrival profile.

In the East Liberty pilot test network, link arrival/depart ratios can be influenced by different uncertainties, e.g., on-street parking (on all three major streets), hidden flows from/to mid-block side streets (Baum Boulevard for Highland Avenue, Sheridan Avenue for Penn Avenue and Penn Circle), and detection errors. Fig. 2a gives the actual hourly *ADRatio* data for one day (from 6am to 21pm, July 2, 2012) for all entry links of the intersection of Penn Avenue and Penn Circle (the central intersection **D** of the pilot test site). This data shows that link arrival/depart ratios varied greatly, where the fluctuations indicate nontrivial uncertainty.
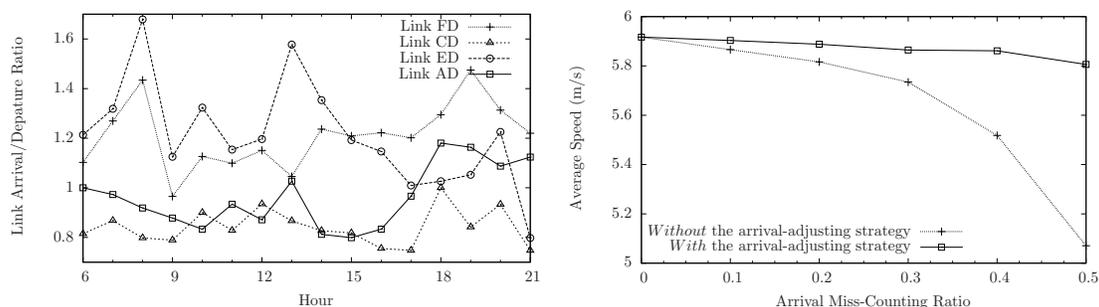
(a) Measured link arrival/departure ratios at **D**   (b) Results for the arrival-adjusting strategy

FIGURE 2 : Arrival-adjusting strategy for inaccuracy in arrival/departure counts

To assess the performance impact of this queue management strategy, we utilize a microscopic simulation model of the East Liberty network, implemented in the Simulation of Urban Mobility (SUMO) traffic simulator [1]. For these experiments and other simulation experiments following later in the paper, the saturation headway is assumed to be $shw = 2$ seconds/vehicle, which is approximated from field observations. Fig. 2b gives the average vehicle speed results in simulation with and without the arrival-adjusting strategy on different mis-counting ratios at the advance detectors, using the field flow data from the AM rush period. For each instance, we calculate the mean of 100 independent runs. Without the arrival-adjusting strategy, the performance drops significantly for increasing ratios of missing counts. With the strengthening strategy, the performance can be restored to near optimal. In the field implementation, moving average is used to quickly catch up with real-time fluctuations, as shown in Fig. 2a.

Disruption Management

In an urban road network, flow disruptions happen on some links from time to time. A flow disruption is defined as a situation where the flow is fully blocked for a time period during a green period. A typical example is from the bus transit system, as buses often share the right of way with general traffic. Due to long dwell times of buses (*11*) when picking up or dropping off passengers, bus stops often significantly reduce the capacity of an intersection, and can have a major impact on vehicle delay. The congestion caused by a bus may also impose unexpected delays on subsequent buses. A flow disruption can also have other sources, e.g., spillback from a downstream intersection.

The *basic disruption management* strategy is realized when constructing the inflows for the local scheduler by imposing an earliest start time $est_{n,i} \geq 0$ for each disrupted flow on link n in phase i (essentially enforcing a delay of $est_{n,i}$). Using this strategy, the SURTRAC scheduler provides more accurate results when the disruption persists longer than the time $est_{n,i}$. Thus, there is no need to define a number of decision rules to handle this disruption (*9*). The larger the value of $est_{n,i}$, the more likely it is that the new schedule will switch to serving the next phase early, though if there are many more vehicles in the current flow, the schedule will remain in the current phase. If the phase is switched, this strategy might impose some delay on the currently disrupted flow (if the phase cannot switch back before the disruption is over). However, in this case the scheduler might service the current flow longer in the next cycle, as other flows will have largely been cleared. This
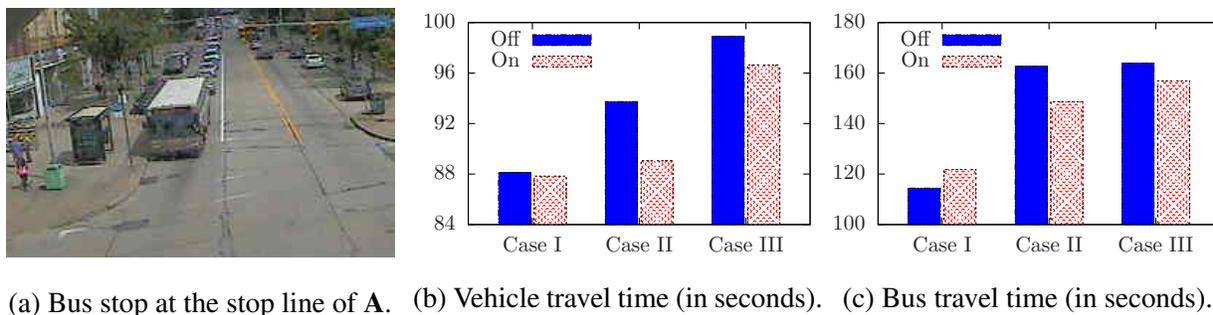
---

[1]http://www.sumo-sim.org

(a) Bus stop at the stop line of **A**.   (b) Vehicle travel time (in seconds).   (c) Bus travel time (in seconds).

FIGURE 3 : Disruption of bus stops at intersection **A** and results for vehicle and bus travel times, with and without the basic disruption management strategy.

strategy is expected to both reduce overall congestion and to reduce the delay for buses in the flow as well. In contrast, transit signal priority, which often causes vehicle delay (*7*), might also suffer from congestion, as buses are sharing the flow.

Several bus transit lines move through the pilot test site on Penn Avenue. As shown in Fig. 3a, there is a bus stop at the stop line of the link **DA**, and buses dwelling at the stop often cause disruptions on this link. For a link n in phase i, a flow disruption is identified if the following condition is observed: the queue blocking state $QB_{n,i}$ is on, and no vehicle departs from the link for t_{NM} seconds during the green time. Figs. 3b and 3c give simulation results for average vehicle and bus travel times (in seconds), with and without use of the basic disruption management strategy. The vehicle travel times are averaged over all vehicles in the network. Here the bus frequency is assumed to be 20 buses per hour, which is quite close to actual bus traffic along the link **DA**. Cases I and II have the same flow condition, with bus dwell times of 20 and 40 seconds, respectively. In case III, the flow is increased by 10%, and the bus dwell time is set at 40 seconds. For the vehicle travel time, cases II and III are significantly improved, whereas case I is only slightly improved. For the bus travel time, Case II and III are also improved, but case I is worse. Thus, this strategy is more likely to improve both vehicle and bus flows if the bus dwell time is long. Note that $t_{NM} = 6 \cdot shw = 12s$ is required to recognize the stopped bus in these experiments, suggesting that the performance can be further improved if recognition time can be reduced. In addition, accurate bus dwell-time estimation (*11*) is useful for field implementation.

Flow Management
Flow management is relevant to different types of flows. For vehicles, there are *minor* and *major flows*. Major flows are expected to be continually serviced if other flows have a low vehicle count or if this major flow is not expected to be stopped by the downstream intersection. For a minor flow, there is a higher probability that no vehicle needs to be serviced, or that the flow can be serviced permissively. In urban environments, the impact of *pedestrian flow* must also be considered.

Minor Flow Management
Due to conservative constraints for the initial pilot deployment, the originally deployed SURTRAC intersection scheduler (*16*) assumed that a cycle consists of a fixed phase sequence, in which some phases might be (partially) dedicated to minor flows regardless of current traffic volumes. This constraint can add significant delay, especially when the yellow and red clearance time between

phases is factored in. Most controllers, (e.g., the model 170 that the City of Pittsburgh uses), support more flexible phase sequences and phase skipping. Taking advantage of this fact, a *hybrid cycle* approach to managing phase durations was developed, wherein major phases are controlled (as before) by the local scheduler, and intervening minor phases consist of a set of sub-phases (corresponding to different flow combinations) that are chosen directly by an actuated mode. For minor flows that have permissive right of way in some other phases, actuation might only be triggered if there are sufficient vehicles (e.g., two vehicles) waiting. Although the mode of flexible phase sequences might bring potential efficiency at the intersection level, the coordination between neighbors can be weakened since actual phase sequences are now less predictable. To account for actuated control of minor phases, the local scheduler predicts the most likely choice (based on recent history) and uses this prediction to better estimate planned outflows in real time.

At the initial pilot test site, intersection **D** (Fig. 4a) is the only intersection that currently utilizes the hybrid cycle strategy, since the other eight intersections have only two phases. In subsequent deployments, this strategy has been widely used. For convenience, a phase is described as a possible combination of the movement groups (as shown in Fig. 4b) defined in the controller. At this intersection, the two major phases (2+6) and (4+8) service Penn Circle and Penn Avenue, respectively. For the minor phase after the phase (2+6), the possible combined choices are (3+7), (3+8), (4+7), or "skip", where the movement groups 3, 4, 7, and 8 in the controller correspond respectively to the movements EDC, ADC+ADE, ADF, and EDA+EDF. Both movement groups 3 and 7 service minor flows, but movement group 7 is also permissive during the major phase (4+8).

Figures 4d and 4c give the phase and vehicle statistics averaged over one week. Based on Table 1b, flow EDA was 6.8 times of flow ADF on average. During the minor phase, ADF and EDA are respectively serviced if (3+7) and (3+8) are chosen, otherwise they are delayed to be serviced during (4+8). Obviously, the flow efficiency will be improved as more vehicles are serviced during this minor phase. As shown in Fig. 4c, (3+8) was selected much more often than (3+7), and this corresponds to the ratio between the counts for EDA and ADF in Fig. 4d. Although service for major flows (e.g., EDA) is generally more efficient than for minor flows (e.g. ADF), sufficiently long queues for any minor flow still need to be cleared in order to prevent potential risk of interferring with through traffic. As shown in Fig. 4c, (3+7) were selected slightly more during the PM peak period to accommodate increased traffic in ADF. This queue for ADF failed to clear, on average, only 1.14 times per day (or in 0.11% of cycles). Figure 4c shows that the number of cycles per hour decreases during peak periods as the SURTRAC scheduler adaptively generates longer phase lengths for major flows. During non-peak periods, this minor phase was often skipped to further increase the efficiency.

Pedestrian Flow Management

In urban environments, pedestrian flow often competes with vehicle flow, as is the case at the pilot site. When a pedestrian button is pressed, the minimal green time of the corresponding phase will be replaced by the significantly longer pedestrian walk time (typically increasing the minimum by 20 seconds or longer). From a vehicle-centric focus, this can lead to disturbances in the coordination of major flows between intersections, particularly when links are short (as they are in East Liberty). From the viewpoint of pedestrians and emerging multi-modal urban policy, pedestrian wait times should be bounded to be reasonably short (*12, 14*).
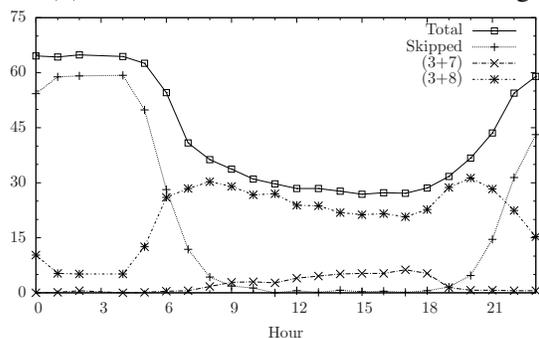
To give more active attention to pedestrian traffic, a *vehicle-pedestrian mixed coordination* protocol has been defined and incorporated. This protocol can be seen as a form of network driven
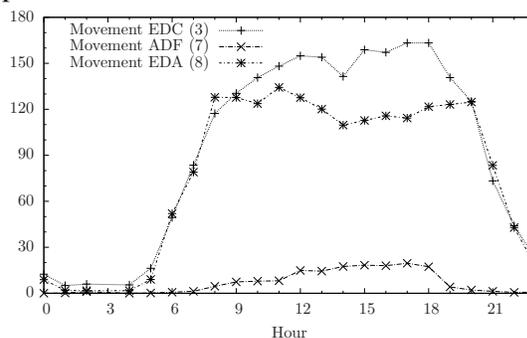
(a) The intersection **D** at Penn Avenue and Penn Circle

| 1 + 6 | 2 + 6 | 3 + 7 | 3 + 8 | 4 + 7 | 4 + 8 |
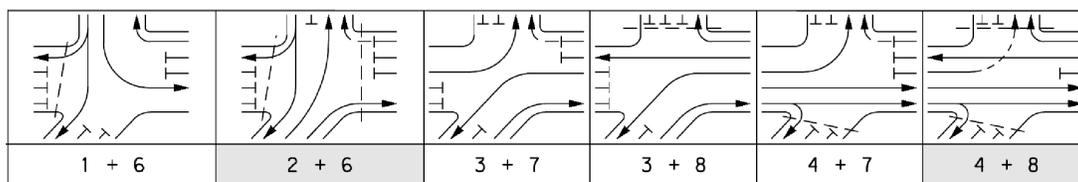
(b) Possible combinations of movement groups defined in the Model 170 controller at **D**

(c) Phase choice statistics

(d) Vehicle count statistics

FIGURE 4 : Hybrid cycle strategy at intersection **D**: the two major phases (2+6) and (4+8) are controlled by the local scheduler, and the minor phases between the major phases are chosen by an actuated mode. The minor phase between (2+6) and (4+8) has four possible choices, i.e. (3+7), (3+8), (4+7), or "skip". For this minor phase, average phase choice and vehicle count statistics are generated using collected data (3/17 — 3/22/2013). In (4c), the phase choice (4+7) is ignored here because the chance is nearly zero at this intersection. In (4d), the movement EDF in movement group 8 is not shown here, although the flow is heavy (as shown in Table 1), since this flow is permissive during many phases, as indicated in (4b).
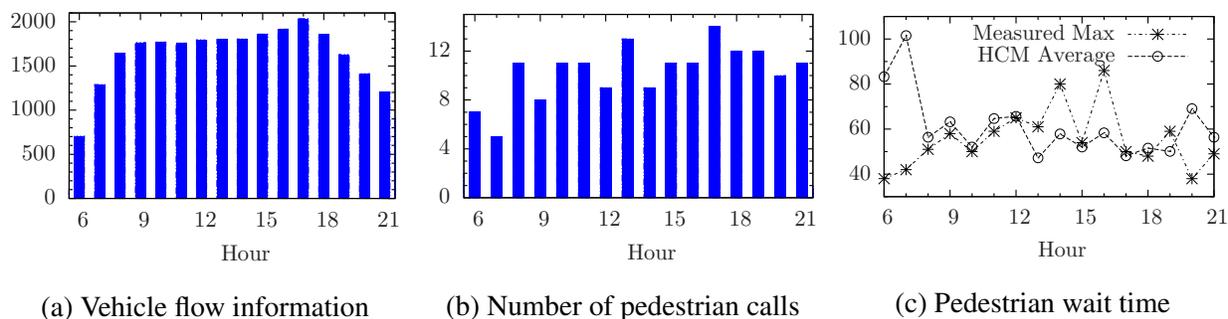
(a) Vehicle flow information   (b) Number of pedestrian calls   (c) Pedestrian wait time

FIGURE 5 : Vehicle flow and pedestrian information at intersection **E** (May-31-2013)

actuated control. Intersections signal the start of major phases to their upstream neighbors. If either the pedestrian call status *ped* is on or the number of waiting vehicles is larger than q_{TH} (q_{TH}=1 by default) for a side street, the phase shift at this upstream neighbor is triggered. The actual switch time point is the end time of the phase that is sending out the major flow toward the upstream intersection, offset by the free travel time on the link between the two intersections (hence ensuring that side street traffic is serviced with minimal disruption to major flows).

As shown in Fig. 4a, Intersection **E** is a neighbor to **D**. For this side street, the vehicle volume is extremely low, but many pedestrians cross Penn Avenue at this intersection. It is much safer than crossing at **D** since the right-turn movements at **D** (i.e., CDE and EDF) are often very heavy (as shown in Table 1b). The total switch-back time to service the pedestrian walk is 39s, obtained by summing the walk time (26s), both yellow times (3s each), and both red clearance times (3s and 4s).

As shown in Fig. 5, there are around 1800 vehicles per hour and around 10 pedestrian calls per hour for much of the day. The vehicle flow has a high priority if the objective is to reduce person delay. Fig. 5c gives a comparable result between our measured maximum pedestrian wait time and the average pedestrian delay calculated using the model (*14*) that is incorporated in the Highway Capacity Manual (HCM) 2010, where the former is the time difference between button actuation and actual walk, and the latter is a function of the cycle length and the effective green time. The actual average wait time might be much lower than the measured maximum, since many pedestrians arrive after the button is actuated.

CONCLUSIONS

In this paper, we have presented some techniques for strengthening the performance of the schedule-driven approach to adaptive traffic signal control implemented in the SURTRAC system. These strengthening strategies are aimed at handling real-world uncertainties and disruptive events, and were motivated by observations of the system in operation in an actual urban environment. The effectiveness of each strategy was evaluated in simulations and analyzed, using the data collected from the pilot site. The overall results demonstrate some initial success of using these techniques to enable more effective real-time adaptive traffic control in urban road networks, and further indicates the potential of decentralized, schedule-driven traffic control. The system has also been broadened to give more active attention to other transportation modes, specifically bus transit and pedestrian flows, in a dynamic, multimodal urban environment.

There are several aspects of this work that warrant further study. First, it may be possible to

better manage flow disruptions due to bus traffic if provided with real-time information about bus arrivals and dwell time distribution. We are currently investigating the use of contemporary object recognition techniques from computer vision for real-time bus recognition. We are also interested in the application of multi-sensor data fusion techniques as a means for achieving more accurate flow prediction, and in the general use of machine learning techniques to continually self tune system parameters from collected performance data. A final direction of our current research aims toward more general, multi-modal traffic flow optimization, where the objective of intersection scheduling is reformulated to treat various modal flows in an integrated manner.

ACKNOWLEDGEMENTS

REFERENCES
[1] Head, L. Event-based short-term traffic flow prediction model. *Transportation Research Record*, Vol. 1510, 1995, pp. 45–52.

[2] Stevanovic, A. *Adaptive Traffic Control Systems: Domestic and Foreign State of Practice*. Tech. Rep. Synthesis 403, NCHRP, 2010.

[3] Sharma, A., D. Bullock, and J. Bonneson. Input-output and hybrid techniques for real-time prediction of delay and maximum queue length at signalized intersections. *Transportation Research Record*, Vol. 2035, 2007, pp. 69–80.

[4] Gentili, M. and P. Mirchandani. Locating sensors on traffic networks: Models, challenges and research opportunities. *Transportation Research Part C*, Vol. 24, 2012, pp. 227–255.

[5] Rhodes, A., D. M. Bullock, J. Sturdevant, Z. Clark, and D. G. Candey. Evaluation of the accuracy of stop bar video vehicle detection at signalized intersections. *Transportation Research Record*, Vol. 1925, 2005, pp. 134–145.

[6] Kwon, J., M. Mauch, and P. Varaiya. Components of congestion: Delay from incidents, special events, lane closures, weather, potential ramp metering gain, and excess demand. *Transportation Research Record*, Vol. 1959, 2006, pp. 84–91.

[7] Abdy, Z. and B. Hellinga. Analytical method for estimating the impact of transit signal priority on vehicle delay. *Journal of Transportation Engineering*, Vol. 137, 2010, pp. 589–600.

[8] Slavin, C., W. Feng, M. Figliozzi, and P. Koonce. A statistical study of the impacts of SCATS adaptive traffic signal control on traffic and transit performance. In *Transportation Research Board Annual Meeting*, 2013.

[9] Yagar, S. and B. Han. A procedure for real-time signal control that considers transit interference and priority. *Transportation Research Part B*, Vol. 28, 1994, pp. 315–331.

[10] Koehler, L. A. and W. Kraus Jr. Simultaneous control of traffic lights and bus departure for priority operation. *Transportation Research Part C*, Vol. 18, 2010, pp. 288–298.

[11] González, E. M., M. G. Romana, and O. M. Álvaro. Bus dwell-time model of main urban route stops. *Transportation Research Record*, Vol. 2274, 2012, pp. 126–134.

[12] Burden, D. and T. Litman. America needs complete streets. *ITE Journal*, Vol. 81, 2011, pp. 36–43.

[13] Kothuri, S., T. Reynolds, C. Monsere, and P. Koonce. Testing strategies to reduce pedestrian delay at signalized intersections: A pilot study in Portland, Oregon. In *Transportation Research Board Annual Meeting*, 2013.

[14] Petritsch, T., B. Landis, P. McLeod, H. Huang, S. Challa, and M. Guttenplan. Level-of-service model for pedestrians at signalized intersections. *Transportation Research Record*, Vol. 1939, 2005, pp. 53–62.

[15] Vedagiri, P. and R. Nagraj. Modeling pedestrian delay and level of service at signalized intersection crosswalks Under mixed traffic conditions. In *Transportation Research Board Annual Meeting*, 2013.

[16] Xie, X.-F., S. F. Smith, L. Lu, and G. J. Barlow. Schedule-driven intersection control. *Transportation Research Part C*, Vol. 24, 2012, pp. 168–189.

[17] Xie, X.-F., S. F. Smith, and G. J. Barlow. Schedule-driven coordination for real-time traffic network control. In *International Conference on Automated Planning and Scheduling*, 2012, pp. 323–331.

[18] Smith, S. F., G. J. Barlow, X.-F. Xie, and Z. B. Rubinstein. SURTRAC: Scalable urban traffic control. In *Transportation Research Board Annual Meeting*, 2013.

[19] Barriere, J. F., J. L. Farges, and J. J. Henry. Decentralization vs hierarchy in optimal traffic control. In *IFAC Control in Transportation Systems*, 1986.

[20] Mirchandani, P. and L. Head. A real-time traffic signal control system: Architecture, algorithms, and analysis. *Transportation Research Part C*, Vol. 9, 2001, pp. 415–432.

[21] Liu, H., X. Wu, W. Ma, and H. Hu. Real-time queue length estimation for congested signalized intersections. *Transportation Research Part C*, Vol. 17, 2009, pp. 412–427.

[22] Day, C. M., J. M. Ernst, T. M. Brennan, C.-S. Chou, A. M. Hainen, S. M. Remias, A. Nichols, B. D. Griggs, and D. M. Bullock. Performance measures for adaptive signal control. *Transportation Research Record*, Vol. 2311, 2012, pp. 1–15.

[23] Perez-Montesinos, J., M. Dixon, and M. Kyte. Detection of stop bar traffic flow state. *Transportation Research Record*, Vol. 2259, 2011, pp. 132–140.

[24] Herroelen, W. and R. Leus. Project scheduling under uncertainty: Survey and research potentials. *European Journal of Operational Research*, Vol. 165, 2005, pp. 289–306.