Optimizing for Everyone using Cooperative Adaptive Control

Gregory J. Barlow, Ph.D.
Rapid Flow Technologies, Inc.
gjb@rapidflowtech.com

Chris Kennett
Green Signals Consulting Ltd
chris.kennett@greensignals.co.uk

Introduction

Urban areas present a challenging traffic signal control problem. Rather than simply optimizing for passenger vehicles following a single, dominant route, signals must serve many modes of travel, such as pedestrians, bicycles, transit, and freight, traveling in many conflicting directions. Urban areas also have many sources of uncertainty, with narrower roads that may be temporarily blocked by buses, parking vehicles, turning trucks, construction, or accidents, causing disruptions that can cause congestion to spill into other parts of the road network (Xie et al, 2014).

Around the world, different forms of adaptive traffic signal control have emerged and become widespread. Most of the existing adaptive approaches can be characterized as either *local* or *central* systems, each with its own merits and limitations. *Central* systems are efficient on networks that require progression and urban areas, but are slower to react to traffic dynamics and tend to employ more rigid coordination structures. *Local* systems may be very efficient in managing isolated intersections or small networks, but are not designed to optimize for more complex network traffic flows. In the recent past, the prevailing opinion of the transportation community, especially in the United States, was that adaptive systems that used real-time, local optimization could not work well for urban environments. With computing power at the intersection now sufficient to achieve both flexible, adaptive local control and efficient network coordination, powerful real-time, decentralized adaptive control now allows *local* traffic signals to cooperate and coordinate traffic across complex road networks, scaling to entire cities.

This paper will explore the Surtrac adaptive traffic signal control system, a real-time, decentralized adaptive system originally developed at Carnegie Mellon University and now commercialized by Rapid Flow Technologies. Surtrac combines perspectives and concepts from recent work in the field of multi-agent planning (Lesser, Decker, and Wagner 2004; Smith et al. 2007) with prior work in traffic theory (Barriere, Farges, and Henry 1986;

Mirchandani and Head 2001; Shelby 2004), and formulates traffic signal control as a decentralized, schedule-driven process. In brief, each intersection independently computes a schedule for servicing all currently approaching vehicles. This schedule is used locally to determine when to switch stages[1] and is recomputed in rolling horizon fashion every second. Network level coordination is achieved through the exchange of schedule information between neighboring intersections. At each decision point, the scheduled outflows from an intersection's immediate upstream neighbors are combined with directly sensed traffic inflows to provide an expanded look-ahead planning horizon. The Surtrac adaptive traffic signal system is unique in its ability to optimize traffic flows in complex, dynamic environments, where there are multiple competing dominant flows that change dynamically through the day.

First deployed in Pittsburgh, Pennsylvania in 2012, the Surtrac adaptive system has spread to 10 cities in the US and Canada, with interest and planned projects around the world, including the UK. Current deployments range from large installations in Pittsburgh (50 intersections), Atlanta, Georgia (24 intersections), Quincy, Massachusetts (21 intersections), and Portland, Maine (21 intersections) to small deployments in Needham, Massachusetts (2 intersections) and Halton Hills, Ontario (3 intersections).

Surtrac is now being used to optimize traffic signals in real-time for many modes of travel, such as pedestrians, cyclists, and transit vehicles, all within the context of overall traffic. In this paper, we will provide an overview of how the Surtrac adaptive system works, how Surtrac has been implemented already, how it is being used to optimize journeys for all road users, and how this approach can translate to the UK.

Surtrac Adaptive Traffic Signal Control

In the Surtrac system, the traffic signal control problem is formulated as a decentralized, schedule-driven process. Each intersection is controlled independently by a local scheduler, that computes a stage schedule that minimizes the total delay for travellers moving through the intersection in a real-time manner 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 formulating the intersection control problem as a single machine scheduling problem. Key to this formulation is an aggregate representation of traffic flows as inflows. Each inflow includes a sequence of jobs, where a job contains

[1] An attempt has been made to use UK-centric traffic signal control terminology in this paper as much as possible. Existing Surtrac deployments in the US and Canada all work with North American standard phase-based ring barrier control systems, which differ in many ways from the stage-based control used in the UK. At the optimization level, however, Surtrac is effectively stage-based.

vehicles in close proximity that have the right of way in a given stage, 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 stage 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 a 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. 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 similar to that previously utilized by PRODYN-D (Barriere et. al, 1986). 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 neighboring intersections that are very close together. To this end, additional coordination mechanisms are incorporated into Surtrac for handling specific nontrivial mis-coordination situations. Spillback that blocks the progress of traffic flow from an upstream intersection is a common inefficiency in urban road networks. Surtrac keeps internal models of queues and all other vehicles inside the network and also uses exit detectors to help detect spillback. In a truly urban network with tight spacing, managing queues is integral to the traffic signal control problem. Surtrac has multiple mechanisms to avoid excess queuing, primarily using soft pressure from a downstream intersection so an upstream intersection avoids pushing vehicles onto queues that cannot accommodate them (Hu and Smith, 2017; Hu and Smith, 2018). Additionally, it uses advance detection at edges of a network with high demand where coordination is necessary or excessive queues might form. Surtrac has been deployed in very congested urban areas, including Pittsburgh, Pennsylvania and Atlanta, Georgia, where these issues arrive frequently.

Deployment Results

Surtrac was first deployed in 2012 at 9 intersections in Pittsburgh, Pennsylvania, showing substantial improvements on many traffic metrics, including emissions and travel time. Prior to the installation of Surtrac, these intersections used conventional coordinated, fully actuated control, with timings installed in 2011. Although the pilot road network size is not large, it has several interesting characteristics that make efficient traffic control challenging. Unlike many arterial road networks optimized by adaptive signal control, this network is an urban grid. Three major roads cross in a triangle, with the dominant traffic flows changing throughout the day. The network also has several very short road segments – as short as 90 feet – connecting intersections, posing a challenge for effective signal coordination. Finally, like most urban environments, events like transit operations, commercial deliveries, construction activity, and pedestrian traffic cause uncertainties and disruptions.

Evaluating this original pilot network showed substantial improvement throughout the day. Performance was compared directly with the existing actuated-coordinated system on the 12 highest volume routes over 4 periods of the day. Overall, the average vehicle travel time was reduced by 26%, number of stops by 31%, wait time at an intersection by 41%, and emissions by 21%. Further details of these results may be found in (Smith et al, 2013). Based on the success of this first deployment, the size of the Surtrac network has grown to 50 total intersections, with another 150 in planning. Similar performance results have been seen at all stages of deployment.

In Portland, Maine, the Surtrac system has shown significant reduction in congestion, with a 20% reduction in delay and a 16% reduction in travel time through the initial Morrill's Corner network of three intersections at the busiest junction in the state of Maine. The success of this deployment on Forest Avenue has led to two subsequent expansions of the network which are taking place in 2019, adding 18 more Surtrac-controlled intersections in Portland for a total network size of 21 intersections, primarily focused on Forest Avenue and Franklin Street. Future expansions in Portland beyond 2019 are also being considered.

The Surtrac system is currently being deployed in the city of Quincy, Massachusetts across 21 intersections, including three HAWK or pedestrian-only intersections. The majority of the intersections are currently operational, and though quantified results are not yet available, qualitative feedback to date has been strong. This deployment is on a complex network with many competing and intersecting roadways located in Quincy Center as well as on Southern Artery, Hannon Parkway, and Burgin Parkway. Critical to this deployment was Surtrac's ability to incorporate passive pedestrian detection into its optimization planning.

Optimizing for Complex Networks

Surtrac was designed for complex, dynamic, urban networks. One of the key differences in the way Surtrac optimizes for efficient traffic signal control is that it optimizes stage durations and coordination second-by-second, while also planning over a long time horizon. This allows Surtrac to both maintain efficient progression across a network and to respond immediately to sources of uncertainty, like double-parked vehicles, blocked lanes, queues that are slow to clear, or other events that can leave other traffic signal control systems in a state of congestion.

Because Surtrac is designed for complex, dynamic urban road networks, it works well with tightly spaced grids. Surtrac was first deployed on a network in Pittsburgh where the largest distance between intersections was 500 feet and the shortest was just 90 feet. In addition to the general design of Surtrac, optimizing traffic second-by-second for real-time responsiveness to dynamic urban traffic, Surtrac also has supplementary mechanisms to achieve coordination in complex networks and in congested conditions. Projected outflows of vehicles from neighboring intersections make it possible to plan over a longer time window to achieve coordination. At times of high saturation, when queues are inevitable, Surtrac uses a form of soft pressure to avoid overfilling downstream queues.

Multimodal Demands

Part of designing Surtrac for urban networks is considering all the modes of travel in that road network. As detection allows, Surtrac optimizes for all modes, including pedestrians, cyclists, transit, and freight. For example, Surtrac monitors pedestrian and bicycle detection as available (often only push-button actuations), includes it in the optimization process, and if desired, prioritizes it. If density information about pedestrians from passive pedestrian detection is available, Surtrac can use that in the same way as vehicle information in the optimization. Bicycles are similar, but often easier to detect with normal vehicle detection systems, particularly if bicycles have separated facilities.

At most of the intersections Surtrac operates, at least one pedestrian phase is on recall, and at many, all are. Pedestrian phases can certainly limit the flexibility of adaptive traffic signal control, but we have always had to operate under those constraints, and Surtrac factors them into the optimization to minimize the impact on the overall system.

Surtrac currently operates in the field with concurrent pedestrian phases, exclusive pedestrian phases, and overlapping pedestrian phases. Pedestrian demands (either from push-buttons or passive pedestrian detection) and pedestrian constraints (e.g., guaranteed maximum pedestrian waiting times and pedestrian minimums) are incorporated into the Surtrac optimization process, not just layered on top of it.

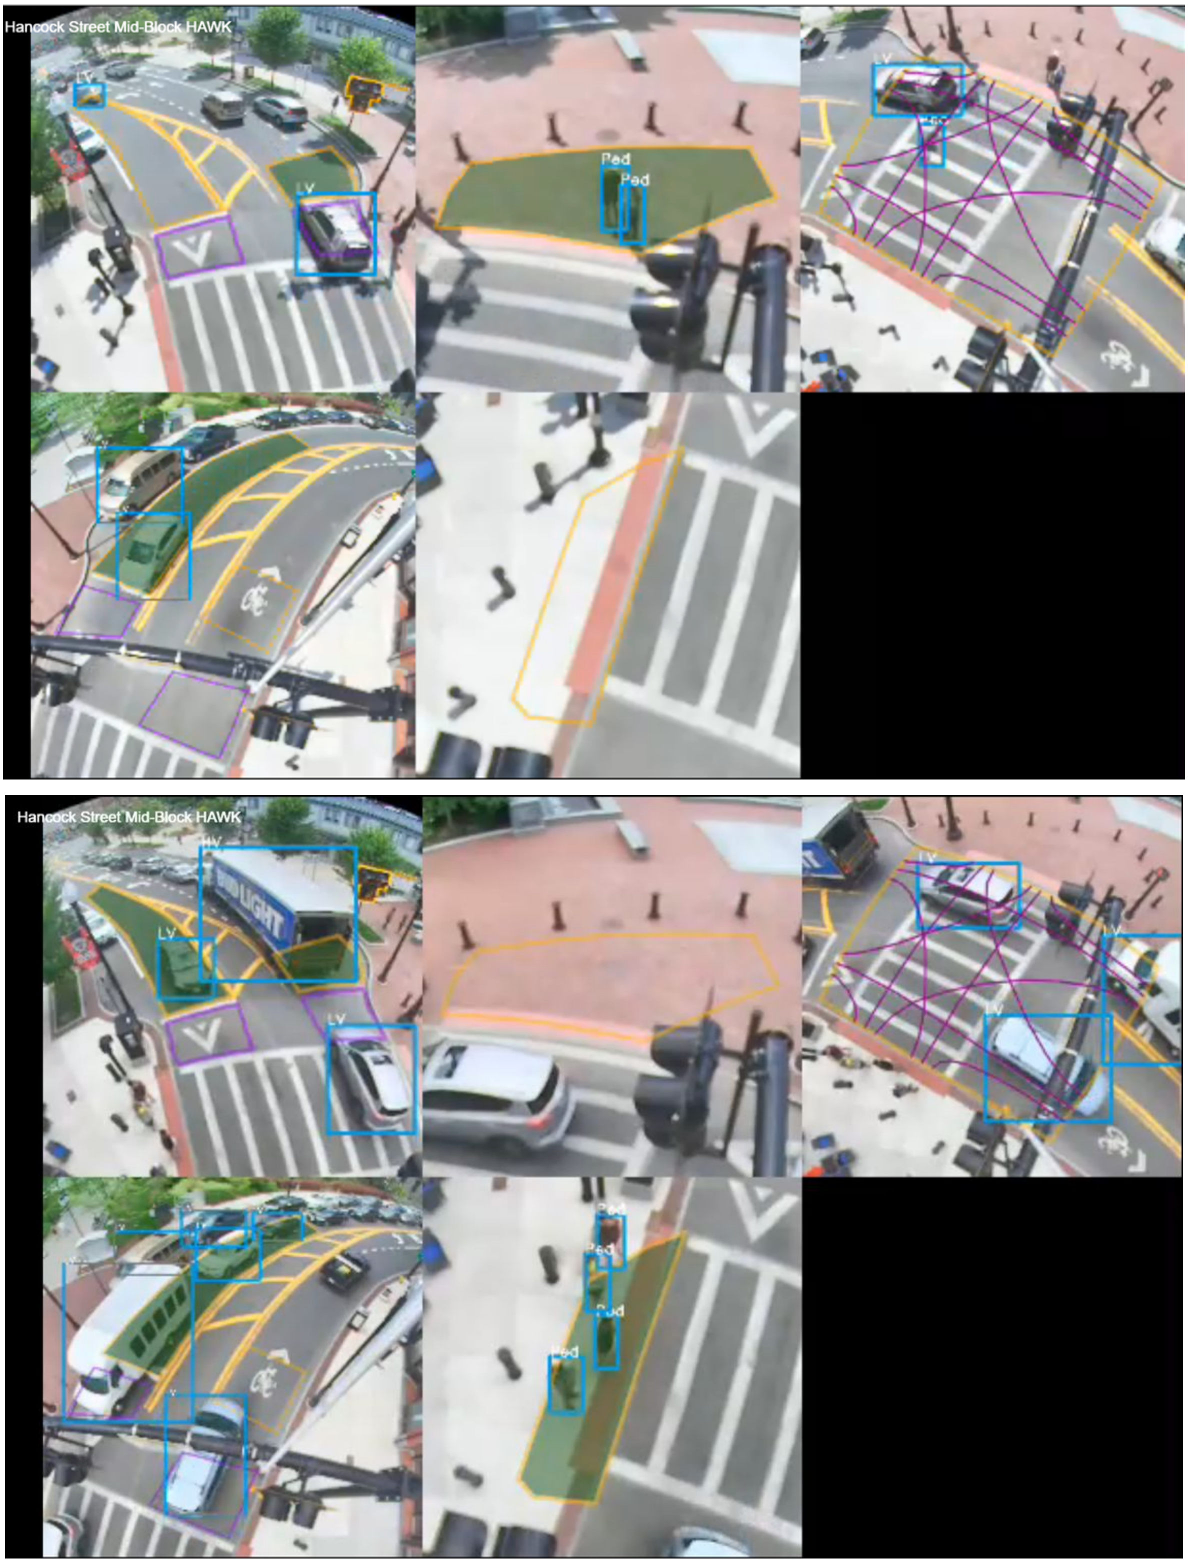

Figure 1: Sample Miovision detection images, showing passive pedestrian detection

In our deployment in Quincy, Massachusetts, Surtrac uses the Miovision SmartSense video detection system for passive pedestrian detection in addition to vehicle detection. This system uses sophisticated computer vision technology on a single-point camera to track and categorize travelers, including pedestrians and vehicles. In this deployment, the number of pedestrians waiting for service is sensed and incorporated directly into Surtrac's optimization. Particularly at the pedestrian-only intersections (called HAWK beacons in the US), pedestrian demand outweighs vehicle demand at certain times of day, so by incorporating pedestrians directly into the optimization, Surtrac can reduce pedestrian wait time while also considering all modes. Several examples of the detection system are shown in Figure 1.

Surtrac supports the use of transit signal priority (TSP) systems in multiple ways. Conventional TSP requests (e.g. green extension or red truncation) may be processed through Surtrac, and due to Surtrac's modular nature, it is possible to build an interface directly with a TSP system already in place. Alternatively, rather than having fixed priority requests, Surtrac can use the outputs from a TSP system to incorporate transit vehicles into the optimization process to achieve what we call transit signal optimization (TSO). This gives Surtrac the ability to more flexibly optimize for transit travel (e.g. clearing queues so a transit vehicle can arrive at a stop, planning around a dwell time model so transit vehicles load on red, etc.) Transit vehicles may also be given variable priority based on agency preferences, such as a value based on the fullness or lateness of a bus. We have been developing these capabilities in cooperation with the transit agency in Pittsburgh, Pennsylvania.

Connected Travelers

The way Surtrac uses detection information is ideally suited for a future where vehicles (and pedestrians and cyclists) are connected to infrastructure and provide rich data about where they are now and where they would like to go. We've been working with dedicated short-range communication (DSRC) connected vehicle (CV) technology since 2015 in our CV testbed in Pittsburgh, and we have also integrated with CV technology in Atlanta. Though the future path of CAV operations is still unclear, we believe Surtrac is the ideal adaptive system for the long bridge period from now (with essentially no CAVs on the road) to a fully connected future. Surtrac is agnostic to the transport mechanism that will be used in the future (e.g., DSRC, 5G, C-V2X, etc), so we've built integrations that work for both DSRC and cellular. We think about the near-term use of CV technology in a few ways.

Early deployments of CV technology in the US focus almost entirely on signal phase and timing (SPaT), real-time information from the intersection on the current state of the traffic signals and expected near-term behavior. Surtrac generates SPaT output data both in Pittsburgh and Atlanta for use with DSRC. Recently, Rapid Flow has also developed systems to allow users to receive these data using smartphones.

We've also been working on ways to extend this technology to pedestrians. In a Carnegie Mellon University led project funded by the US Department of Transportation ATTRI program that Rapid Flow is involved in, we are also working to bring signal data to pedestrians via an accessible smartphone app. Particularly aimed at blind and disabled pedestrians, the app allows users with mobility difficulties to request additional time to cross an intersection to make crossings safer and more efficient.

We have been working on other low penetration applications that can improve efficiency. In simulation, we've shown that when vehicles share their short-term route information (e.g., over the next two intersections), their delay can be reduced a further 20-30% over Surtrac control, without incurring delay for other vehicles (Hawkes, 2016). This works because the uncertainty in the optimization process can be significantly reduced by knowing where vehicles plan to go. This approach is effective from the first vehicle that opts-in, but as more vehicles participate, even the vehicles not participating realize a benefit. This approach would be especially useful for autonomous vehicles (AVs), which already know their routes and could easily share them, but could be accomplished with any traveler. The first field pilot of this approach was completed in 2019 with one of the large AV companies testing in Pittsburgh, and subsequent field pilots are planned with a major international last mile freight provider and a major automotive OEM. Full-scale deployment in a city is planned over the next 18 months.

Implementation Details

Each Surtrac intersection is controlled locally by a software agent running on an embedded computer (Surtrac processor) located in the traffic cabinet for the intersection. The agent manages the control of the traffic signals and all of the vehicle detectors located at that intersection. Surtac is capable of interfacing with any type of detection (either directly or through the traffic signal controller), and detection at a single intersection can be heterogeneous. Agents at neighboring intersections communicate both projected outflows and raw sensor data from advance detectors with their direct neighbors. By distributing control, the system can be more robust to network latency and failure, detector malfunctions, and other sources of uncertainty.

Surtrac requires lane by lane detection to provide both presence (occupancy) and pulse (event-based counting) detection functions. Figure 2 shows a typical detection setup for an approach. Stop bar and exit zones are required for all approaches, with the goal to detect vehicles as they enter and leave the intersection. For network edges, advance zones are configured to provide information about vehicles before they enter the network. Inside the network, exit zones are used to provide advance information to a downstream intersection. In most cases, advance and exit detection requires one detection zone per location, with that single zone performing both presence and pulse functions.

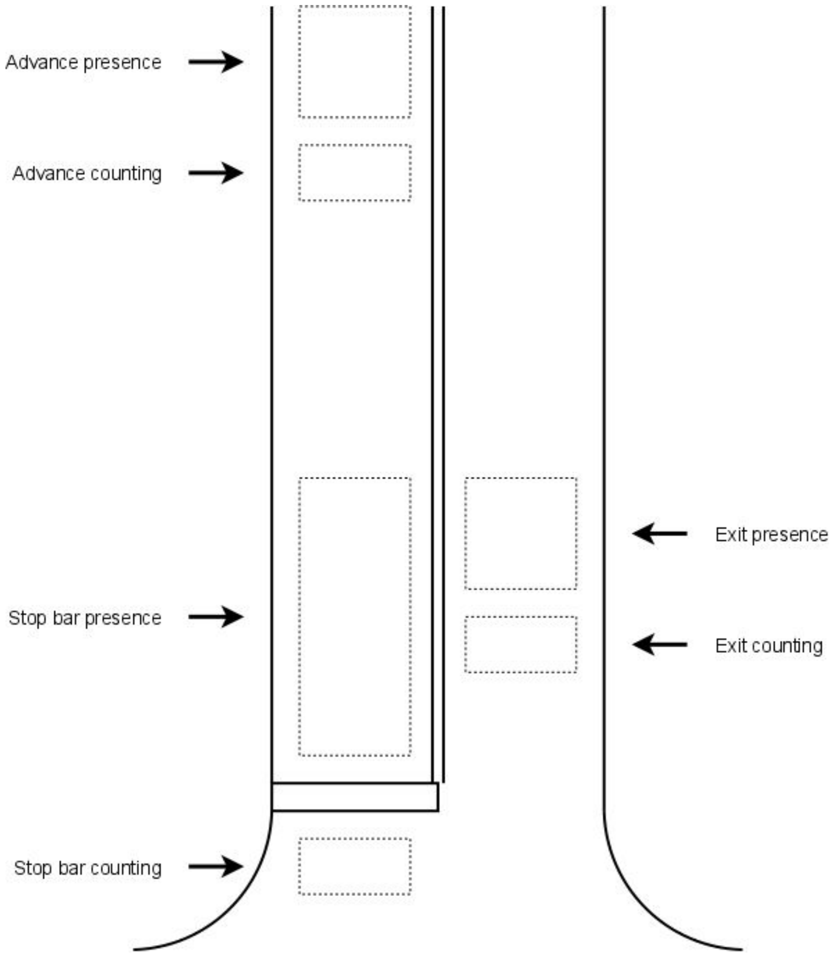

Figure 2: Surtrac detection requirements

Surtrac typically uses conventional video detection in the US. Surtrac is also using several advanced detection technologies in our US deployments. When a detection system performs event-based counting of vehicles using dedicated zones, Surtrac uses that instead of a simple detection zone for counting zones. This approach increases detection accuracy. Surtrac networks are typically designed with judicious use of advance vehicle detection, such as long-distance radar, on major entrances into the network.

In the North American market, there are several standards for traffic signal control equipment. Surtrac is designed as a modular system to most easily interface with multiple standards. Most North American traffic signal controllers now use a standard software API for communication, the National Transportation Communications for ITS (Intelligent Transportation Systems) Protocol (NTCIP). The majority of Surtrac's deployments use NTCIP for communication with and control over the traffic signal controller, though other types of controller communication are used in Pittsburgh, which uses a legacy controller type.

Surtrac in the UK

Over the past several years, Rapid Flow has been investigating the possibility of bringing Surtrac to the UK. Though most adaptive systems built for the US market focus on suburban arterial optimization, Surtrac is more consistent with the traffic flows in many UK cities and towns, with many of our US deployments exhibiting older road networks with narrower streets and more complex traffic demands. There are many differences in traffic signal control between the UK and North American markets, but the fundamental control problem solved by Surtrac remains the same. At the optimization level, Surtrac uses a stage-based approach consistent with UK control, so the challenges of adapting this approach to the UK lie primarily in use of detection and hardware compatibility.

Many of the existing traffic signal control systems in the UK have detection needs which match those of Surtrac, allowing existing inductive loops to be reused in many locations (see Figure 2). SCOOT loops are typically installed in ideal locations to serve as exit loops for Surtrac, and MOVA loops perform well as advance loops in Surtrac. VA loops and other stop bar loops satisfy Surtrac's need for detection at the stop line. Since Surtrac can take in many sources of detection, other technologies like passive pedestrian detection or connected travelers could be integrated as available.

A Surtrac-compatible outstation is currently in design for use as a Surtrac processor for the UK. Prototype hardware was tested in the field at a junction in the UK in February 2019. Detection testing showed the ability for Surtrac to use existing loops consistent with its detection requirements. Control has also been tested using an approach similar to that used by MOVA. Pilot deployments of Surtrac have been discussed with several UK cities, and installation for the first Surtrac system in the UK is expected to begin in 2020.

References

1. J. F. Berriere, J. L. Farges, and J. J. Henry, “Decentralization vs hierarchy in optimal traffic control,” IFAC Control in Transportation Systems, 1986.
2. Allen Hawkes, “Traffic Control with Connected Vehicle Routes in SURTRAC,” Master's Thesis, Tech. Report, CMU-RI-TR-16-20, Robotics Institute, Carnegie Mellon University, May, 2016
3. Hsu-Chieh Hu and S. F. Smith, “Coping with large traffic volumes in schedule-driven traffic signal control,” International Conference on Automated Planning and Scheduling, 2017.
4. Hsu-Chieh Hu and Stephen F Smith, “Bi-directional information exchange in decentralized schedule-driven traffic control,” International Conference on Autonomous Agents and Multiagent Systems, 2018.
5. Pitu Mirchandani and Larry Head, “A real-time traffic signal control system: Architecture, algorithms, and analysis,” Transportation Research Part C: Emerging Technologies 9(6), 2007, pp. 415-432.
6. Steven G. Shelby, “Single-intersection evaluation of real-time adaptive traffic signal control algorithms,” Transportation Research Record, 2007.
7. Stephen F. Smith, Anthony Gallagher, Terry Zimmerman, Laura Barbulescu, and Zachary B. Rubinstein, “Distributed management of flexible times schedules,” International Conference on Autonomous Agents and Multiagent Systems, 2007.
8. Stephen F. Smith, Gregory J. Barlow, Xiao-Feng Xie, and Zachary B. Rubinstein, “SURTRAC: Scalable Urban Traffic Control,” Transportation Research Board Annual Meeting, 2013.
9. Stephen F. Smith, Gregory J. Barlow, Xiao-Feng Xie, Zachary B. Rubinstein, “Smart urban signal networks: Initial application of the SURTRAC adaptive traffic signal control system,” International Conference on Automated Planning and Scheduling, 2013.
10. Xiao-Feng Xie, Stephen F. Smith, Liang Lu, and Gregory J. Barlow, “Schedule-driven intersection control,” Transportation Research: Part C, volume 24, 2012, pp. 168-189.
11. Xiao-Feng Xie, Stephen F. Smith, and Gregory J. Barlow, “Schedule-driven coordination for real-time traffic network control,” International Conference on Automated Planning and Scheduling, 2012.
12. Xiao-Feng Xie, Stephen F. Smith, Gregory J. Barlow, Tingwei Chen, “Coping with real-world challenges in real-time urban traffic control,” Transportation Research Board Annual Meeting, 2014.