

Back-pressure traffic signal control with unknown routing rates

Jean Gregoire* Emilio Frazzoli** Arnaud de La Fortelle***
Tichakorn Wongpiromsarn****

* *Mines ParisTech, Paris, France*
jean.gregoire@mines-paristech.fr.

** *Massachusetts Institute of Technology, Boston, USA*
frazzoli@mit.edu

*** *Mines ParisTech, Paris, France; Inria Paris-Rocquencourt, France*
arnaud.de_la_fortelle@mines-paristech.fr

**** *Thailand Center of Excellence for Life Sciences, Thailand*
tichakorn@tcels.or.th

Abstract: The control of a network of signalized intersections is considered. Previous works proposed a feedback control belonging to the family of the so-called back-pressure controls that ensures provably maximum stability given pre-specified routing probabilities. However, this optimal back-pressure controller (BP*) requires routing rates and a measure of the number of vehicles queuing at a node for each possible routing decision. It is an idealistic assumption for our application since vehicles (going straight, turning left/right) are all gathered in the same lane apart from the proximity of the intersection and cameras can only give estimations of the aggregated queue length. In this paper, we present a back-pressure traffic signal controller (BP) that does not require routing rates, it requires only aggregated queue lengths estimation (without direction information) and loop detectors at the stop line for each possible direction. A theoretical result on the Lyapunov drift in heavy load conditions under BP control is provided and tends to indicate that BP should have good stability properties. Simulations confirm this and show that BP stabilizes the queuing network in a significant part of the capacity region.

Keywords: road traffic, traffic lights, traffic control, transportation control, queuing theory, back-pressure, network control.

1. INTRODUCTION

In today's metropolitan transportation networks, traffic is regulated by traffic light signals which alternate the right-of-way of users (e.g., cars, public transport, pedestrians). Congestion is a major problem resulting in a loss of utility for all users due to delayed travel times over the network Shepherd (1992). That is why it is of high interest to find a control policy that can stabilize a network of signalized intersections under the largest possible arrival rates. Under traffic light control, a particular set of feasible simultaneous movements, called a phase, is decided for a period of time Papageorgiou et al. (2003). Controlling a traffic light consists of designing rules to decide which phase to apply over time. Pre-timed policies activate phases according to a time-periodic pre-defined schedule, and the signal settings can be fixed by optimization, assuming within-day static demand Casetta et al. (2006); Miller (1963); Gartner et al. (1975). They are not efficient under changing arrival rates which require adaptive control. Many major cities currently employ adaptive traffic signal control systems including SCOOT Hunt et al. (1982), SCATS Lowrie (1990), PROLYN Henry et al. (1984), RHODES Mirchandani and Head (2001), OPAC Gartner (1983) or TUC Diakaki et al. (2002). These systems update some control variables of a configurable

pre-timed policy on middle term, based on traffic measures. Control variables may include phases, splits, cycle times and offsets Papageorgiou et al. (2003). More recently, feedback control algorithms that ensure maximum stability have been proposed both under deterministic arrivals Varaiya (2013), and stochastic arrivals Varaiya (2009); Wongpiromsarn et al. (2012). These algorithms are based on the so-called back-pressure control presented in the seminal paper Tassiulas and Ephremides (1992) for applications in wireless communication networks and require real-time queues estimation. An optimal back-pressure traffic signal controller (BP*) is presented in Wongpiromsarn et al. (2012) and Varaiya (2009). They are defined under different modelling assumptions but they are algorithmically equivalent. The key benefit of back-pressure control is that it can be completely distributed over intersections, i.e., it requires only local information and it is of $\mathcal{O}(1)$ complexity. However, the strong assumptions of the model in Varaiya (2009) (and also implicitly in Wongpiromsarn et al. (2012)) is that controllers require routing rates and a measure of the number of vehicles queuing at every node of the network for each possible routing decision. However, in reality, apart from the proximity of the intersection, vehicles (going straight, turning left, turning right, etc.) are all gathered, and it is difficult to estimate the number of vehicles queuing for each di-

rection (see Figure 1). Cameras can give good estimations of the total number of vehicles queuing at a given node, but not the direction of vehicles. However, it is feasible to detect if there are some vehicles (or no vehicle) that want to go to a given destination, if we assume the existence of dedicated lanes from the proximity of the intersection with loop detectors at the stop line.




Fig. 1. Dedicated lanes for turning vehicles. The dedicated lanes are indicated by road markings when vehicles approach the intersection. Apart from the proximity of the intersection, vehicles are all gathered.

The back-pressure control proposed in this paper (BP) requires such loop detectors and an estimation of the total number of vehicles queuing at each node (gathering all possible directions). It does not assume any knowledge of routing rates. We evaluate the performance of BP with regards to the optimal BP* control. The contribution of the paper is to provide a back-pressure traffic signal controller based on more realistic assumptions on the available measurements than state-of-the-art back-pressure traffic signal control and to show in simulations that stability is conserved in a significant part of the capacity region.

The paper is organized as follows. Section 2 describes the queuing network model. Sections 3 is mainly expository: it describes BP* highlighting its stability-optimality. The contributions of the paper are presented in Section 4 and 5. Section 4 exhibits BP and a theoretical result on the Lyapunov drift that tends to indicate that it should have good stability properties. The simulations of Section 5 confirm this and show that BP stabilizes the network in a significant part of the capacity region. Section 6 concludes the paper and opens perspectives.

2. MODEL

As standard in queuing network control, time is slotted, and each time slot maps to a certain period of time during which a control is applied. It is convenient to use a fixed pre-defined time slot length, whose size corresponds to the minimal duration of a phase. When the time slot size is fixed, the traffic signal control problem consists of computing at the beginning of each time slot t the phase to apply during slot t . The network of intersections is modelled as a directed graph of nodes $(N_a)_{a \in \mathcal{N}}$ and links $(L_j)_{j \in \mathcal{L}}$. Nodes represent lanes with queuing vehicles, and links enable transfers from node to node: this is a standard queuing network model.

It is a multiple queues one server queuing network. Every signalized intersection is modelled as a server managing a

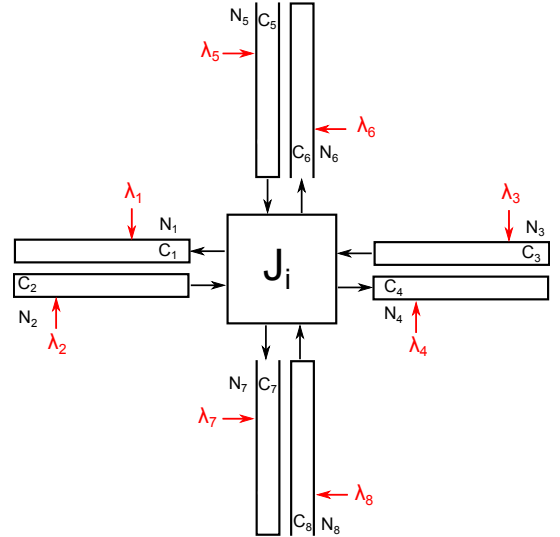


Fig. 2. A junction with 4 incoming nodes and 4 outgoing nodes which corresponds to the intersection depicted in Figure 3.

junction which consists of set of links. Junctions $(J_i)_{i \in \mathcal{J}}$ are supposed to form a partition of links. For every junction J , $\mathcal{I}(J)$ and $\mathcal{O}(J)$ denote respectively the inputs and the outputs of J . Inputs (resp. outputs) of junction J are nodes N such that there exists a link $L \in J$ pointing from (resp. to) N . The reader should consider the introduction of junctions in the model as an overlay of the queuing network model. For the sake of simplicity, we do not represent links in the queuing network representation of Figure 2.

Every server maintains an internal queue for every input/output, and server work enables to transfer vehicles from an input to an output of the junction. The internal queue at node N_a is a vector Q_a and $Q_{ab}(t)$ denotes the number of vehicles in the queue of node N_a entering N_b upon leaving N_a . The aggregated queue length $Q_a(t) = \sum_b Q_{ab}(t)$ denotes the total number of vehicles at node N_a considering all possible routings after exiting N_a . In this paper, queues are supposed to have infinite capacities: there is no blocking (see Gregoire et al. (2013a) for an adaptation of back-pressure traffic signal control in the context of finite capacities).

At every time slot t , servers work, resulting in vehicles transfers. Under phase-based control, the transmission rate offered by servers are set by activating a given signal phase p_i at each junction J_i from a predefined finite set of feasible phases \mathcal{P}_i at every time slot t . Let $\mathcal{P} = \prod_{i \in \mathcal{J}} \mathcal{P}_i$ denote the set of feasible global phases. Each global phase $p = (p_i)_{i \in \mathcal{J}} \in \mathcal{P}$ results in a different service matrix $\mu(p)$ where $\mu_{ab}(p)$ represents the transmission rate offered by servers to transfer vehicles from N_a to N_b in a time slot when phase p is activated. The transmission rate is assumed to be binary, $\mu_{ab}(p) \in \{0, s_{ab}\}$: it is zero or it equals the saturation rate s_{ab} . Only the vehicles which are at a node at the beginning of time slot t can be transferred from that node to another node during slot t . Figure 3 depicts the 4 typical phases of a 4 inputs/4 outputs junction.




Fig. 3. A typical set of feasible phases at a junction. For example, supposing that service rates equal 0 or 1, the non zero service rates for phase (a) are μ_{31} , μ_{36} , μ_{24} and μ_{27}

Exogenous arrivals occur at every node of the network. Let $A_a(t)$ denote the number of vehicles that exogenously arrive at node N_a during slot t . The arrival process $A_a(t)$ is assumed to be rate convergent with long-term arrival rate $\lambda_a \geq 0$. When a quantity of vehicles arrives at node $N_a \in \mathcal{I}(J_i)$ during slot t , exogenously and endogenously, it is split and added into queues Q_{ab} , $b \in \mathcal{O}(J_i)$. The routing process is exogenous and assumed to be rate convergent with ratios r_{ab} with $\sum_b r_{ab} \leq 1$ (see the supplementary material Gregoire et al. (2013b) for more details). Exits are modelled by assuming that the routing matrix is non-conservative. $1 - \sum_b r_{ab}$ represents the exit rate of vehicles entering node N_a , it is the ratio of vehicles directly removed from the network when entering node N_a , i.e. not added to any queue Q_{ab} . Note that the only variable that is controlled is the activated phase at every time slot t , denoted by $p(t)$, and yielding a service matrix $\mu(p(t))$ during slot t .

3. BP* CONTROLLER

3.1 The controller

In the following, we expose BP* signal control. It is an extension of the algorithm proposed in Varaiya (2009) where internal/exit links are not differentiated, because exits may occur at any link of the network. It is quite equivalent to the back-pressure controller of Wongpiromsarn et al. (2012), assuming the nodes carry direction information. Loosely speaking, the idea of back-pressure control is to compute pressure at every node based on node occupancy and to open flows which have a high upstream pressure and a low downstream pressure, like opening a tap.

Algorithm 1 defines BP* control. At every junction i , for each phase $p \in \mathcal{P}_i$, the weighted sum $\sum_{a,b} W_{ab}(t)\mu_{ab}(p)$ is computed. $W_{ab}(t)$, the weight associated to transfers from N_a to N_b , is the difference between the upstream

Algorithm 1 BP* control

Require:

Queues lengths matrix $Q(t)$,
 Pressure functions $P_{ab}(Q_{ab})$ for all $a, b \in \mathcal{N}$,
 Routing matrix r .

function BP*

```

5:   for  $i \in \mathcal{J}$  do
       for  $a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)$  do
            $\Pi_{ab}(t) \leftarrow P_{ab}[Q_{ab}(t)]$ 
       end for
       for  $a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)$  do
10:           $W_{ab}(t) \leftarrow \max(\Pi_{ab}(t) - \sum_c r_{bc}\Pi_{bc}(t), 0)$ 
       end for
        $p_i^*(t) \leftarrow \arg \max_{p_i \in \mathcal{P}_i} \sum_{a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)} W_{ab}(t)\mu_{ab}(p_i)$ 
       end for
       return Phase  $p^*(t)$  to apply in time slot  $t$ 
15: end function

```

pressure $\Pi_{ab}(t)$ and the weighted downstream pressure $\sum_c r_{bc}\Pi_{bc}(t)$. BP* consists of selecting the phase that maximizes the weighted sum. Moreover, we assume that in case of equality the selected phase $p^*(t)$ always satisfies $\mu_{ab}(p^*(t)) = 0$ if $W_{ab}(t) = 0$.

3.2 Optimal stability

The following theorem states that under linear pressure functions with strictly positive slope, BP* as defined by Algorithm 1 is optimal in terms of stability, i.e. stabilizes the network for all arrivals rates that can be stabilized considering all possible control strategies. It is an extension of the results of Varaiya (2009), because vehicles can enter/exit the network at any node, there is no distinction between exit nodes and internal nodes. Moreover, in contrast with Varaiya (2009); Wongpiromsarn et al. (2012), pressure functions are just assumed to be linear with strictly positive slope in this paper: $P_{ab}(Q_{ab}) = \theta_{ab}Q_{ab}$, $\theta_{ab} > 0$.

Theorem 1. (Back-pressure optimality). Assuming that pressure functions are linear with strictly positive slopes, BP* as defined by Algorithm 1 is stability-optimal.

Proof. Due to space limitations, the full proof is not provided in this paper and is available in the supplementary material Gregoire et al. (2013b). Stability is proved using the Lyapunov function $V(t) = \mathcal{V}(Q(t)) = \sum_{a,b} \theta_{ab}Q_{ab}(t)^2$. The existence of $B, \eta > 0$ such that:

$$\mathbb{E}\{V(t+1) - V(t)|Q(t)\} \leq B - \eta \sum_{a,b} Q_{ab}(t), \quad (1)$$

enables to conclude stability for the queuing network using the sufficient condition proved in Neely (2003).

4. BP CONTROLLER

4.1 The controller

Back-pressure control proposed in Section 3 requires complete knowledge of the queues lengths matrix $Q(t)$ and the routing rates. For our application, a complete knowledge of $Q(t)$ is not realistic because dedicated lanes for turning

vehicles are only from the proximity of the junction. Farther, all vehicles are gathered and the controller does not have access to the direction of every vehicle in the absence of vehicle-to-infrastructure communications. That is why we propose in the present paper a controller that uses only the aggregated queues lengths $Q_a(t) = \sum_b Q_{ab}(t)$, i.e. a queue length without direction information. It is defined by Algorithm 2. It computes the phase to apply at every time slot without requiring neither routing rates nor complete knowledge of queues lengths matrix $Q(t)$ and takes as inputs the aggregated queues lengths $Q_a(t) = \sum_b Q_{ab}(t)$. However, it still requires vehicle detectors variables $d_{ab}(t) \in [0, 1]$ defined below:

$$d_{ab}(t) = \min(Q_{ab}(t)/s_{ab}, 1) \quad (2)$$

The variable $d_{ab}(t)$ can be measured by loop detectors positioned at dedicated lanes.

Algorithm 2 BP control

Require:

Queues lengths $Q_a(t)$,
Pressure functions $P_a(Q_a)$,
Loop detectors variables $d_{ab}(t)$.

function BP

```

5:   for  $i \in \mathcal{J}$  do
      for  $a \in \mathcal{I}(J_i) \cup \mathcal{O}(J_i)$  do
           $\Pi_a(t) \leftarrow P_a[Q_a(t)]$ 
      end for
      for  $a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)$  do
10:       $W_{ab}(t) \leftarrow d_{ab}(t) \max(\Pi_a(t) - \Pi_b(t), 0)$ 
      end for
       $p_i^*(t) \leftarrow \arg \max_{p_i \in \mathcal{P}_i} \sum_{a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)} W_{ab}(t) \mu_{ab}(p_i)$ 
    end for
    return Phase  $p^*(t)$  to apply in time slot  $t$ 
15: end function

```

Algorithm 2 defines BP control. Note that for transfers from N_a to N_b , the upstream pressure is now $\Pi_a(t)$ and the downstream pressure is $\Pi_b(t)$: individual queue pressures $\Pi_{ab}(t)$ are not required. Moreover, the difference between the upstream pressure and the downstream pressure is multiplied by $d_{ab}(t)$ to form $W_{ab}(t)$. Hence, if at time slot t , there is no vehicle at N_a going to N_b , the weight $W_{ab}(t)$ associated to transfers from N_a to N_b vanishes.

4.2 Behaviour of the Lyapunov drift under heavy load conditions

Let us consider the Lyapunov function $\mathcal{V}(Q)$ and its evolution through time $V(t)$ defined below:

$$V(t) = \mathcal{V}(Q(t)) = \sum_a \theta_a Q_a(t)^2 = \sum_a \theta_a \left(\sum_b Q_{ab}(t) \right)^2 \quad (3)$$

Let us define heavy load conditions at time slot t as states of the network such that if the right-of-way is given to any individual queue, it can be emptied at saturation flow, i.e. there are enough vehicles in the individual queue to ensure saturation:

$$\forall a, b \in \mathcal{N}, Q_{ab}(t) \geq s_{ab} \quad (4)$$

The following theorem proves that under heavy load conditions the Lyapunov drift respects the sufficient condition for network stability if $\lambda + \epsilon \in \Lambda_r$, for sufficiently large ϵ .

Theorem 2. (Lyapunov drift under heavy load conditions). Assume $\lambda + \epsilon \in \Lambda_r$, BP control as defined in Algorithm 2 is applied and the network is in heavy load conditions, then there exists $B, \eta > 0$ such that :

$$\mathbb{E}\{V(t+1) - V(t) \mid Q(t)\} \leq B - \eta \sum_a Q_a(t) \quad (5)$$

for sufficiently large ϵ .

Proof. Due to space limitations, the full proof is not provided in this paper and is available in the supplementary material Gregoire et al. (2013b).

The above theorem tends to indicate that the network should have good stability properties because the condition for stability is verified in heavy load conditions for λ sufficiently interior to the capacity region. Unfortunately it does not enable to conclude that the network is stable in a significant part of the capacity region. Indeed, heavy load conditions can not be guaranteed at all time, and when an individual queue Q_{ab} is below the saturation flow s_{ab} , it is a constraint for the emptying of Q_a , that can unstabilize the queuing network. Hence, the characterization of the stability region of the queuing network under BP control with the modelling assumptions presented in Section 2 is still a challenging problem. That is why we propose to implement the two back-pressure controllers and to compare their behaviour. The results of the simulations are presented in the next section.

5. SIMULATIONS

5.1 The simulation platform

The model and the algorithms presented in this paper have been implemented into a simulator coded in Java. It simulates a grid network and every junction of the grid has 4 inputs, 4 outputs, and 4 feasible phases as depicted in Figure 3. The height and the width are parameters. Every individual flow allowed by phases of Figure 3 equals 10 (it is the saturation rate). Vehicles are generated at each node N_a at an arrival rate λ_a that can be set as desired. The arrival process generates individual arrivals as well as batches of 10 vehicles. The routing ratios are fixed at the beginning of the simulation.

5.2 Behaviour of the two back-pressure controllers

Simulations have been carried out for a 21×21 square grid network (see Figure 4). First of all, we present simulations results in the case of a network that has been configured with the same arrival rates and routing rates at every node of the network.

Simulation results for a particular network and particular arrival/routing rates The numerical results of Figure 5 correspond to the following parameters. Turn left probability when a vehicle enters a node: 0.2; turn right probability




Fig. 4. The 21×21 grid network used for the presented simulations.

when a vehicle enters a node: 0.2; go straight probability when a vehicle enters a node: 0.5; exit probability when a vehicle enters a node: 0.1; probability of a batch: 0.05; pressure functions $P_a(Q_a) = Q_a$ and $P_{ab}(Q_{ab}) = Q_{ab}$ ($\theta_a = \theta_{ab} = 1$); vehicles are generated at every node with the same arrival rate $\lambda > 0$ that can be set as desired at the beginning of the simulation.

Experiments are carried out at height different arrival rates: $\lambda = 0.4, 0.5, 0.6, 0.65, 0.7, 0.75, 0.8$ and 0.9 vehicles per time slot. Figure 5 depicts the global queue of the network over time, i.e. $\sum_a Q_a(t) = \sum_{a,b} Q_{ab}(t)$, for the height arrival rates, under BP* control and under BP control. One can observe in Figure 5 that under BP* control, the queuing network is stabilized for $\lambda \leq 0.7$ and gets unstable from $\lambda = 0.75$. Under BP control, it is stabilized for $\lambda \leq 0.65$ and gets unstable from $\lambda = 0.7$. First of all, it proves that as expected, BP control is not stability-optimal. However, in the particular setting of this experiment, (uniform arrivals/routing rates and grid network), the performance of BP and BP* are very close, and the optimality gap is around $0.05/0.7 \simeq 10\%$, i.e. a performance of 90%.

However, such a uniform network is not realistic and the results of the next paragraph try to evaluate the performance of BP with regards to BP* with less specific routing/arrival parameters.

Evaluation of BP with regards to BP on several samples of parameters* In the following simulations, the routing/arrival process parameters are not uniform over nodes any more. 10 samples of parameters have been generated. For each sample, the routing/arrival rates are generated as follows. For each direction (straight, left, right), (uniformly) random values between 0 and 1 are generated, say y_s, y_l, y_r ; a (uniformly) random value between 0 and 0.1 is generated for exits, say y_ω ; and the routing rates are set by normalization of the generated real values, i.e. for the left direction for example, the routing rate is $y_l/(y_s + y_l + y_r + y_\omega)$. The arrivals rates are set by generating a (uniformly) random value between 0 and 1 for every node, say λ_a^0 . At the beginning of the simulation, a parametrizable scaling value x enables to fix the actual arrival rate of the current simulation: $\lambda_a = x\lambda_a^0$, where x has the same value over nodes. The value of 0.1 for the scale of exits is quite arbitrary and, loosely speaking, fixes




Fig. 5. Evolution of the global queue of the network over time for height arrival rates. Comparison of the behaviour of the network under BP/BP* control.

the averaged number of travelled nodes before exiting the network.

Note that the routing rates and the values λ_a^0 are fixed for a given sample. However, the value of λ_a depends on the value of x set at the beginning of the simulation. The parameter x enables to define a performance for BP with regards to BP* for a given sample. We let x vary and we observe the maximum value of x such that the network is stable under BP versus BP* (say x_{\max}^* for BP* and x_{\max}^{BP} for BP). We define the performance of BP with regards to BP*, or more shortly the performance of BP (because BP* is optimal), as follows:

$$\text{performance}(\text{BP}) = x_{\max}^{\text{BP}}/x_{\max}^* \quad (6)$$

As for previously presented simulations, the probability of a batch is 0.05 and the pressure functions are linear with slope 1. Figure 6 depicts the performance obtained for the 10 samples, the average performance and the standard deviation. The average performance is around 80%, i.e. the optimality gap is about 20%. The simulation results prove that the performance of BP is affected by the routing/arrival rates. Hence, the distribution (over samples) of the performance would be different for a different distribution of routing/arrival rates. Nevertheless, in the particular setting of the experiment, the average optimality gap of 20% seems again a low price to pay with regards to the

much more realistic assumptions on the measurements available to compute the control.




Fig. 6. Performance distribution for ten samples. The point above the axis represents the average performance over samples and the horizontal bar is the standard deviation.

However, these promising results can not be extended to any kind of network of intersections and further simulations with a more general structure of network should be carried out to confirm the closeness of performance. We are currently implementing our algorithms in a traffic simulator in order to test the performance of BP control with real traffic data of the city of Singapore.

6. CONCLUSION AND PERSPECTIVES

The simulation results of this paper prove that BP is not optimal but tend to indicate that it stabilizes the queuing network in a significant part of the capacity region. The benefits of BP originate from the more realistic assumptions on queues measurements. Computing the phase to apply only requires aggregated queues lengths estimation that can be provided by cameras, and loop detectors at dedicated lanes. The optimality gap, around 20% in the particular setting of the experiments, seems a low price to pay for the benefits of relaxed assumptions on the available measurements. However, simulations have been conducted in a grid network, which is a particular structure, and with synthetic data which can strongly differ from real traffic data. To confirm the closeness of performance, simulations should be carried out in a more advanced traffic network simulator.

Finally, the emergence of vehicle-to-infrastructure communications opens avenues to enhance traffic signal control. The traffic signal controllers can have access, in particular, to the destination node of every vehicle. As a result, back-pressure control with a multiple-commodity queuing network model, as proposed in Neely (2003) in the context of wireless communication networks, should be investigated.

ACKNOWLEDGEMENTS

This work was supported in part by the Singapore National Research Foundation through the Future Urban Mobility Interdisciplinary Research Group at the Singapore-MIT Alliance for Research and Technology.

REFERENCES

Cascetta, E., Gallo, M., and Montella, B. (2006). Models and algorithms for the optimization of signal settings on urban networks with stochastic assignment models. *Annals of Operations Research*, 144(1), 301–328.

Diakaki, C., Papageorgiou, M., and Aboudolas, K. (2002). A multivariable regulator approach to traffic-responsive network-wide signal control. *Control Engineering Practice*, 10(2), 183–195.

Gartner, N.H. (1983). Opac: A demand-responsive strategy for traffic signal control. *Transportation Research Record*, (906).

Gartner, N.H., Little, J.D., and Gabbay, H. (1975). Optimization of traffic signal settings by mixed-integer linear programming part i: The network coordination problem. *Transportation Science*, 9(4), 321–343.

Gregoire, J., Frazzoli, E., de La Fortelle, A., and Wongpiromsarn, T. (2013a). Capacity-aware back-pressure traffic signal control. *arXiv preprint arXiv:1309.6484*.

Gregoire, J., Frazzoli, E., de La Fortelle, A., and Wongpiromsarn, T. (2013b). Supplementary material to: Back-pressure traffic signal control with unknown routing rates. *hal preprint hal-00905063*. URL <http://hal.archives-ouvertes.fr/hal-00905063>.

Henry, J.J., Farges, J.L., and Tuffal, J. (1984). The prodyn real time traffic algorithm. In *Proceedings of the 4th IFAC/IFORS Conference on Control in Transportation Systems*.

Hunt, P., Robertson, D., Bretherton, R., and Royle, M. (1982). The scoot on-line traffic signal optimisation technique. *Traffic Engineering & Control*, 23(4).

Lowrie, P. (1990). Scats, sydney co-ordinated adaptive traffic system: A traffic responsive method of controlling urban traffic. Technical report.

Miller, A.J. (1963). Settings for fixed-cycle traffic signals. *Operations Research*, 373–386.

Mirchandani, P. and Head, L. (2001). A real-time traffic signal control system: architecture, algorithms, and analysis. *Transportation Research Part C: Emerging Technologies*, 9(6), 415–432.

Neely, M.J. (2003). *Dynamic power allocation and routing for satellite and wireless networks with time varying channels*. Ph.D. thesis, LIDS, Massachusetts Institute of Technology.

Papageorgiou, M., Diakaki, C., Dinopoulou, V., Kotsialos, A., and Wang, Y. (2003). Review of road traffic control strategies. *Proceedings of the IEEE*, 91(12), 2043–2067.

Shepherd, S. (1992). A review of traffic signal control. Technical report.

Tassiulas, L. and Ephremides, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12), 1936–1948.

Varaiya, P. (2009). A universal feedback control policy for arbitrary networks of signalized intersections. Technical report.

Varaiya, P. (2013). The max-pressure controller for arbitrary networks of signalized intersections. In *Advances in Dynamic Network Modeling in Complex Transportation Systems*, 27–66. Springer.

Wongpiromsarn, T., Uthacharoenpong, T., Wang, Y., Frazzoli, E., and Wang, D. (2012). Distributed traffic signal control for maximum network throughput. In *Proceedings of the 15th international IEEE conference on intelligent transportation systems*, 588–595.