



HAL
open science

Back-pressure traffic signal control with unknown routing rates

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

► **To cite this version:**

Jean Gregoire, Emilio Frazzoli, Arnaud de La Fortelle, Tichakorn Wongpiromsarn. Back-pressure traffic signal control with unknown routing rates. 2013. hal-00905055v1

HAL Id: hal-00905055

<https://hal.science/hal-00905055v1>

Preprint submitted on 15 Nov 2013 (v1), last revised 5 May 2014 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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

**** *Ministry of Science and Technology, Thailand*
tichakorn@gmail.com

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*) suffers from two limitations. Controllers are assumed to have knowledge of routing rates; and more importantly, to have access to the number of vehicles queuing at a node for each possible routing decision. However, it is an idealistic assumption for our application since vehicles (going straight, turning left/right, etc.) 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 the measurement of queue lengths on the dedicated lanes from the proximity of the intersections. This is a more realistic requirement compared to state-of-the-art back-pressure traffic signal control. Optimal stability is proved for BP*. 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 Cascetta et al. (2006); Miller (1963); Gartner et al. (1975). They are not efficient under changing arrival rates which require adaptive control. Many major cities cur-

rently 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 Tassioulas 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 are assumed to have a perfect knowledge of routing rates and to have access to 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 direction (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.




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. Note that most of the time, in standard intersections, dedicated lanes are required only for vehicles turning left because the phase giving the right-of-way to vehicles going straight also gives the right-of-way to vehicles turning right. Note also that the positioning of vehicles in the right dedicated lane is not always respected which can significantly affect the traffic through the intersection.

By contrast, the back-pressure control (BP) proposed in this paper requires such vehicle detectors from the proximity of the intersection 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 and 4 are mainly expository: Section 3 presents the notion of capacity region and Section 4 describes BP* highlighting its stability-optimality. The contributions of the paper are presented in Section 5 and 6. Section 5 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 6 confirm this and show that BP stabilizes the network in a significant part of the capacity region. Section 7 concludes the paper and opens perspectives.

2. MODEL

2.1 The time slot

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 .

2.2 Queuing network topology

The network of intersections is modelled as a directed graph of nodes $(N_a)_{a \in \mathcal{N}}$ and links $(L_j)_{j \in \mathcal{L}}$. The graph is referred as the network graph. Nodes represent lanes with queuing vehicles, and links enable transfers from node to node. This is a standard queuing network model.

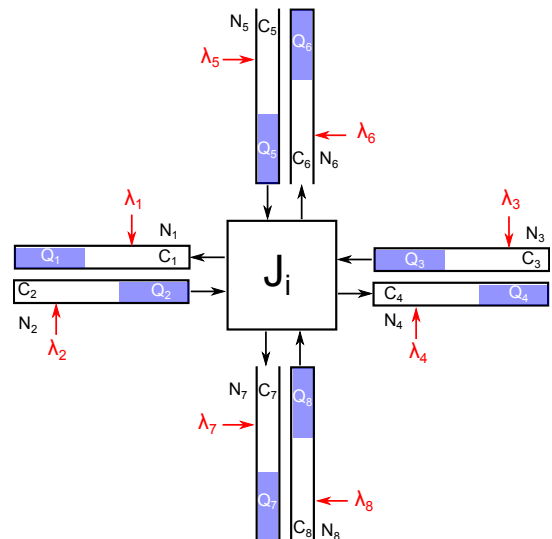


Fig. 2. A junction with 4 incoming nodes and 4 outgoing nodes which corresponds to the intersection depicted in Figure 3.

A key point for our application in traffic signal control is that it is a multiple queues one server queuing network. Every signalized intersection is modelled as a server managing a 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 and we assume that at every junction, there exists a link from any input

to any output. If the link does not exist physically, the flow through this link will be constrained to zero.

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. Due to routing of vehicles, 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. (2013b) for an adaptation of back-pressure traffic signal control in the context of finite capacities).

2.3 Phase-based control

At every time slot t , servers work, resulting in vehicles transfers. It is convenient to consider the service matrix μ defined below:

Definition 1. (Service rate, Flow). For all $a, b \in \mathcal{N}$,

- the service rate μ_{ab} represents the transmission rate offered by servers to transfer vehicles from N_a to N_b , i.e. the maximum number of vehicles transferred from N_a to N_b during the next time slot;
- the endogenous flow variable f_{ab} represents the actual number of vehicles leaving N_a and entering N_b .

Since the number of vehicles transferred is less than or equal to the transmission rate offered by the servers, the following inequality holds:

$$f_{ab} \leq \mu_{ab}. \quad (1)$$

Only the vehicles which are currently at a node at the beginning of time slot t can be transferred from that node to another node during slot t .

The definition below applies to a rate matrix g of any kind, hence both to the service matrix μ and and the flow matrix f .

Definition 2. (Input rate, Output rate). Given a matrix g , for all $a \in \mathcal{N}$, the input rate g_a^{in} and the output rate g_a^{out} with regards to N_a are defined as follows:

$$g_a^{\text{in}} = \sum_b g_{ba}, \quad (2)$$

$$g_a^{\text{out}} = \sum_c g_{ac}. \quad (3)$$

Under phase-based control, service rates 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 . Each global phase $p = (p_i)_{i \in \mathcal{J}} \in \mathcal{P}$ results in a different service $\mu(p)$ where $\mathcal{P} = \prod_{i \in \mathcal{J}} \mathcal{P}_i$ denotes the set of feasible global phases.

Assumption 1. (Phase-controlled service). If $N_a \in \mathcal{I}(J_i)$ and $N_b \in \mathcal{O}(J_i)$, the service rate μ_{ab} satisfies:

$$\mu_{ab} \in \{\mu_{ab}(p_i) : p_i \in \mathcal{P}_i\}. \quad (4)$$

The service matrix μ satisfies:

$$\mu \in \{\mu(p) : p \in \mathcal{P}\} \quad (5)$$




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}

The abuse of notation in the above assumption is justified by the fact that the service rate depends only on the applied phase p , i.e., can also be considered as a function of p . Figure 3 depicts the 4 typical phases of a 4 inputs/4 outputs junction.

In this paper, for the sake of simplicity, we do not take into account any exogenous variable z which would affect the flow matrix associated to each phase, yielding a service matrix $\mu(z, p)$. However, as proved in previous works Neely et al. (2005), back-pressure properties can be extended to this case. As a result, we assume that for each phase the service rate from one node to another node is zero or equals the saturation rate:

Assumption 2. (Binary service rates). For all $a, b \in \mathcal{N}$, there exists s_{ab} , the saturation rate from N_a to N_b , such that for all $p \in \mathcal{P}$, $\mu_{ab}(p) \in \{0, s_{ab}\}$.

Finally, we make the following assumption which enables to switch off the service from N_a to N_b independently from other service rates:

Assumption 3. (Service rates independence). For all phase $p \in \mathcal{P}$ and for all $a, b \in \mathcal{N}$, there exists a phase $\tilde{p}_1 \in \mathcal{P}$ such that $\mu_{ab}(\tilde{p}_1) = 0$ and for all $(c, d) \neq (a, b)$, $\mu_{cd}(\tilde{p}_1) = \mu_{cd}(p)$.

2.4 Exogenous arrivals and routing model

We assume that there are exogenous arrivals 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 .

Definition 3. (Rate convergent process). A process $X(t)$ is rate convergent with rate x if:

- $\lim_{t \rightarrow +\infty} \frac{1}{T} \sum_{t=0}^{T-1} X(t) = x$
- For any $\delta > 0$, there exists an interval size T such that for any initial time t_0 and regardless of past history,

the following condition holds: $|\mathbb{E}\{\frac{1}{T} \sum_{t=0}^{T-1} X(t)\} - x| \leq \delta$

Assumption 4. (Rate convergent arrival process). For all $a \in \mathcal{N}$, the process $A_a(t)$ is rate convergent with rate $\lambda_a \geq 0$. For all t , $A_a(t)$ is independent from $\{Q(\tau)\}_{\tau \leq t}$.

Under rate convergence assumption, λ_a represents the long-term arrival rate at node N_a .

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} , $b \in \mathcal{O}(J_i)$, according to an exogenous routing process $R_{ab}^{(t)}$, defined for all $a, b \in \mathcal{N}$.

Assumption 5. (Rate convergent routing process). The arrival process and the routing process are independent, and for all t , $R_{ab}^{(t)}$ is independent from $\{Q(\tau)\}_{\tau \leq t}$. $R_{ab}^{(t)}$ takes an integer, returns an integer, and for $\bar{X} \in \mathbb{N}$, $\sum_b R_{ab}^{(t)}(X) \leq X$. For all process $X(t)$ such that for all t , $R_{ab}^{(t)}$ is independent from $\{X(\tau)\}_{\tau \leq t}$, there exists a rate $r_{ab} \geq 0$ for all $a, b \in \mathcal{N}$ such that $R_{ab}^{(t)}(X(t)) - r_{ab}X(t)$ is rate convergent with rate 0.

As a consequence of the above assumptions:

$$\sum_b r_{ab} \leq 1 \quad (6)$$

Exits are modelled by assuming that the routing matrix is non-conservative, i.e., $\sum_b r_{ab} \leq 1$. Note that $1 - \sum_b r_{ab}$ represents the exit rate of vehicles entering node N_a . One could consider an additional node ω representing the external world playing the role of sink of the exit flow from N_a at rate $r_{a\omega} = 1 - \sum_b r_{ab}$.

2.5 Network dynamics

The latter routing assumption closes our model, and the dynamics of the network is now fully described. Since service rates depend only on the phase applied at every junction, controlling the network consists of controlling the phase applied at every junction.

Definition 4. (Control). A control $p(t)$ for the queuing network returns the phase to apply at every junction during slot t . If $p(t)$ can be expressed as a function of $Q(t)$, $p(t)$ is a feedback control.

The network dynamics under control p follows:

$$Q_{ab}(t+1) = Q_{ab}(t) + R_{ab}^{(t)}(A_a(t) + f_a^{\text{in}}(t)) - f_{ab}(t) \quad (7)$$

$$Q_{ab}(t+1) \leq \max[0, Q_{ab}(t) - \mu_{ab}(p(t))] + R_{ab}^{(t)}[A_a(t) + \mu_a^{\text{in}}(p(t))] \quad (8)$$

An inequality holds instead of an equality because the number of vehicles transferred is less or equal to the transmission rate offered by the servers.

If $p(t)$ is a feedback control, i.e. a function of $Q(t)$, the process $Q(t)$ is a Markov chain with long-term stationary transition probabilities, which depend on the feedback control.

Note that feedback controllers may differ in the information required on $Q(t)$. For example, one feedback controller may require only aggregated queues, $p(t)$ being in fact a function of $(\sum_b Q_{ab}(t))_{a \in \mathcal{N}}$. By contrast, another feedback controller as BP of Section 4 may require the full knowledge of $Q(t)$, $p(t)$ being a function of all $(Q_{ab}(t))_{a, b \in \mathcal{N}}$. It is a key characteristic of the feedback controller because it determines the measurements that are assumed to be available.

3. CHARACTERIZATION OF THE CAPACITY REGION

3.1 Stability definition

A key property of queuing systems is stability, defined below:

Definition 5. (Stability). The queuing network is stable if each individual queue U satisfies:

$$\limsup_{T \rightarrow +\infty} \mathbb{E} \left\{ \frac{1}{T} \sum_{t=0}^{T-1} 1_{U(t) > V} \right\} \rightarrow 0 \text{ as } V \rightarrow +\infty \quad (9)$$

This definition of stability is standard and is applicable to networks with arbitrary inputs and control laws Neely (2003).

3.2 The capacity region

A remarkable property of queuing networks as modelled in Section 2 is that it is possible to define a capacity region which describes the set of arrivals rates vectors that can be stably handled by the network.

Definition 6. (Capacity region Neely (2003)). Given a routing matrix r , the capacity region Λ_r is the closure of the set of all arrival rate vectors λ that can be stabilized by some control.

The following theorem provides a characterization of the capacity region in our particular model:

Theorem 1. (Capacity region characterisation). Given a routing matrix r , the capacity region Λ_r is the set of arrival vectors λ such that there exists $g \in \Gamma$ satisfying:

$$\forall a, b \in \mathcal{N}, r_{ab}(\lambda_a + g_a^{\text{in}}) \leq g_{ab} \quad (10)$$

where Γ is the set of feasible long-term endogenous service rates, defined below:

$$\Gamma = \text{Convex_Hull}\{\mu(p) : p \in \mathcal{P}\} \quad (11)$$

Moreover,

- $\lambda \in \Lambda_r$ is a necessary condition for network stability, considering all possible controls (including those that have perfect knowledge of future events)
- $\lambda \in \text{int}(\Lambda_r)$ is a sufficient condition for the network to be stabilized by a control that does not have knowledge of future events.

Proof. Due to space limitations, the full proof is not provided in this paper and is available in the supplementary material Gregoire et al. (2013c). It is based on flow conservation that ensures the existence of rate-convergent sample paths $f_{ab}(t)$ verifying Equation 10 in the long-term.

4. BP* CONTROLLER

4.1 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 input pressure and a low output pressure, like opening a tap.

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 ,
 Phase selection policy ϕ in case of equality.

5: **function** BP*
 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
 10: **for** $a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)$ **do**
 $W_{ab}(t) \leftarrow \max(\Pi_{ab}(t) - \sum_c r_{bc} \Pi_{bc}(t), 0)$
 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
 15: **return** Phase $p^*(t)$ to apply in time slot t
end function

Algorithm 1 is a generic version of back-pressure. Indeed, it does not specify neither pressure functions, nor the policy ϕ that decides which phase to select at the arg max of Line 13 when an equality holds. For example, the policy ϕ^{random} consists of selecting randomly any phase which maximizes the weighted sum.

Remark 1. In Wongpiromsarn et al. (2012), assuming that nodes ("links" in the terminology of Wongpiromsarn et al. (2012)) carry direction information, they can be indexed by (a, b) . Then, the weight associated to a "movement" (as defined in the terminology of Wongpiromsarn et al. (2012)) from N_{ab} to N_{bc} for a phase p is $(Q_{ab} - Q_{bc})\xi(N_{ab}, N_{bc}, p)$, where $\xi(N_{ab}, N_{bc}, p)$ is the estimation of the flow of vehicles from N_{ab} to N_{bc} if phase p is activated as defined in Wongpiromsarn et al. (2012). If we assume that when the phase p is activated, the flow of vehicles from N_{ab} to N_{bc} is estimated to be $r_{bc}s_{ab}$, the weight associated to this "movement" is $r_{bc}s_{ab}(Q_{ab} - Q_{bc})$. Hence, the weight associated to the "movements" from N_{ab} to all N_{bc} is $\sum_c r_{bc}s_{ab}(Q_{ab} - Q_{bc})$. As a result, if $\sum_c r_{bc} = 1$, we obtain exactly $(Q_{ab} - \sum_c r_{bc}Q_{bc})s_{ab}$. This justifies the equivalence of the two controllers under the assumption that links carry direction information and the flow $\xi(N_{ab}, N_{bc}, p)$ is given by $r_{bc}s_{ab}$ for all ab, bc .

4.2 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. 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, pressure functions in Varaiya (2009); Wongpiromsarn et al. (2012) all have the same slope: $P_{ab}(Q_{ab}) = Q_{ab}$. In this paper, pressure functions are just assumed to be linear with strictly positive slope, but the slopes can be different: $P_{ab}(Q_{ab}) = \theta_{ab}Q_{ab}$, $\theta_{ab} > 0$. Note that the fact that different slopes θ_{ab} can be used while conserving the optimality guarantee has already been proved for back-pressure control with a multiple-commodity queuing network in the context of a wireless communication network Neely (2003). As noticed by the authors, this enables to give priority to some queues.

Theorem 2. (Back-pressure optimality). Assume that the phase selection policy ϕ in case of equality always privileges phases p such that for all $a, b \in \mathcal{P}$, $\mu_{ab}(p) = 0$ if $W_{ab} = 0$ and pressure functions are linear with strictly positive slopes. Then, 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. (2013c). 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:

$$\mathbb{E}\{V(t+1) - V(t)|Q(t)\} \leq B - \eta \sum_{a,b} Q_{ab}(t), \quad (12)$$

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

5. BP CONTROLLER

5.1 The controller

Back-pressure control proposed in Section 4 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:

$$d_{ab}(t) = \min(Q_{ab}(t)/s_{ab}, 1) \quad (13)$$

The variable $d_{ab}(t)$ is easier to measure than $Q_{ab}(t)$ because it only requires the knowledge of $Q_{ab}(t)$ in the range $[0, s_{ab}]$, i.e. from the proximity of the junction.

Algorithm 2 BP control

Require:

Queues lengths $Q_a(t)$,
Pressure functions $P_a(Q_a)$,
Vehicle detectors variables $d_{ab}(t)$,
Phase selection policy ϕ in case of equality.

5: **function** BP**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**10: **for** $a \in \mathcal{I}(J_i), b \in \mathcal{O}(J_i)$ **do** $W_{ab}(t) \leftarrow d_{ab}(t) \max(\Pi_a(t) - \Pi_b(t), 0)$ **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**15: **return** Phase $p^*(t)$ to apply in time slot t **end function**

5.2 Behaviour of the Lyapunov drift under heavy load conditions

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

$$V(t) = \mathcal{V}(Q(t)) = \sum_a \theta_a Q_a(t)^2 = \sum_a \theta_a \left(\sum_b Q_{ab}(t) \right)^2 \quad (14)$$

Let 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:

$$\forall a, b \in \mathcal{N}, Q_{ab}(t) \geq s_{ab} \quad (15)$$

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 3. (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 :

$$\mathbb{E}\{V(t+1) - V(t) \mid Q(t)\} \leq B - \eta \sum_a Q_a(t) \quad (16)$$

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. (2013c). The key point in the proof is that summing Equation 7 over all b gives:

$$Q_a(t+1) - Q_a(t) = \sum_b \left(R_{ab}^{(t)} (A_a(t) + f_a^{\text{in}}(t)) - f_{ab}(t) \right) \quad (17)$$

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.

6. SIMULATIONS

6.1 The simulation platform

The model and the algorithms presented in this paper have been implemented into a simulator coded in Java, which main specificities are described below:

- It simulates a grid network, as the one depicted in Figure 4. Every junction has 4 inputs, 4 outputs, and 4 feasible phases as depicted in Figure 3.
- Every individual flow allowed by phases of Figure 3 equals 10.
- 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 decisions at each junction are independent and identically distributed with fixed routing rates that are set as desired at the beginning of the simulation.




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

6.2 Behaviour of the two back-pressure controllers

Simulations have been carried out for the grid network of 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 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 λ 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 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




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.

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:

$$\text{performance}(\text{BP}) = x_{\max}^{\text{BP}}/x_{\max}^* \quad (18)$$

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.

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.




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.

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.

7. 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 vehicle detectors from the proximity of the intersections to detect the presence of vehicles in 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.

Further work should also focus on the analytical determination of the stability region of BP. The theoretical result provided in this paper tends to indicate that BP should have good stability properties, because the Lyapunov drift respects the condition for stability in heavy load conditions. However, this result does not prove the stability of the queuing network, even in a subset of the capacity region. Determining a set (ideally the maximal set) of arrival rates vectors that can be stabilized by BP control stays a challenging problem.

Finally, the emergence of intelligent vehicles opens avenues for vehicle-to-infrastructure communications to enhance traffic signal control. The works of Dresner and Stone (2008); Kowshik et al. (2011); Mehani and de La Fortelle (2007); Gregoire et al. (2013a) proved that cars automation enables much more precise vehicles coordination resulting in decreased travel times. However, these works do not study the question of stability under given arrival rates. With vehicle-to-infrastructure communication, the controllers can have access to much more information, and in particular, 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.
- Dresner, K. and Stone, P. (2008). A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research*, 31, 591–656.
- 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., Bonnabel, S., and De La Fortelle, A. (2013a). Priority-based coordination of robots. 29 pages.
- Gregoire, J., Frazzoli, E., de La Fortelle, A., and Wongpiromsarn, T. (2013b). Capacity-aware back-pressure traffic signal control.
- Gregoire, J., Frazzoli, E., de La Fortelle, A., and Wongpiromsarn, T. (2013c). Supplementary material to: Back-pressure traffic signal control with unknown routing rates.
- 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).
- Kowshik, H., Caveney, D., and Kumar, P. (2011). Provable systemwide safety in intelligent intersections. *IEEE Transactions on Vehicular Technology*, 60(3), 804–818. doi:10.1109/TVT.2011.2107584.
- Lowrie, P. (1990). Scats, sydney co-ordinated adaptive traffic system: A traffic responsive method of controlling urban traffic. Technical report.
- Mehani, O. and de La Fortelle, A. (2007). Trajectory planning in a crossroads for a fleet of driverless vehicles. In *Proceedings of the 11th international conference on Computer aided systems theory*, EUROCAST’07, 1159–1166. Springer-Verlag, Berlin, Heidelberg.
- 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.
- Neely, M.J., Modiano, E., and Rohrs, C.E. (2005). Dynamic power allocation and routing for time-varying wireless networks. *Selected Areas in Communications, IEEE Journal on*, 23(1), 89–103.
- 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.