

STABILITY AND IMPLEMENTATION OF A CYCLE-BASED MAX PRESSURE CONTROLLER FOR SIGNALIZED TRAFFIC NETWORKS

LEAH ANDERSON

Dept of Civil and Environmental Engineering
Univ. of California, Berkeley, Berkeley, CA, USA

THOMAS PUMIR

Dept. of Operations Research and Financial Engineering
Princeton University, Princeton, NJ, USA

DIMITRIOS TRIANTAFYLLOS

TSS-Transport Simulation Systems
Barcelona, Spain

ALEXANDRE M. BAYEN

Dept of Civil and Environmental Engineering
Dept. of Electrical Engineering and Computer Science
Univ. of California, Berkeley, Berkeley, CA, USA

(Communicated by Benedetto Piccoli)

ABSTRACT. Intelligent use of network capacity via responsive signal control will become increasingly essential as congestion increases on urban roadways. Existing adaptive control systems require lengthy location-specific tuning procedures or expensive central communications infrastructure. Previous theoretical work proposed the application of a max pressure controller to maximize network throughput in a distributed manner with minimal calibration. Yet this algorithm as originally formulated has unpractical hardware and safety constraints. We fundamentally alter the formulation of the max pressure controller to a setting where the actuation can only update once per multiple time steps of the modeled dynamics. This is motivated by the case of a traffic signal that can only update green splits based on observed link-counts once per “cycle time” of 60-120 seconds. Furthermore, we extend the domain of allowable actuations from a single signal phase to any convex combination of available signal phases to model intra-cycle signal changes dictated by pre-selected cycle green splits. We show that this extended max pressure controller will stabilize a vertical queueing network given restrictions on admissible demand flows that are slightly stronger than those suggested in the original formulation of max pressure. We ultimately apply our cycle-based extension of max pressure to a simulation of an existing arterial network and provide comparison to the control policy that is currently deployed at the modeled location.

2010 *Mathematics Subject Classification.* Primary: 90B15, 93C99; Secondary: 90B20.

Key words and phrases. Adaptive signal control, vertical queueing network, network stability.

This work was funded by the California Department of Transportation under the Connected Corridors program.

1. Introduction. This article investigates the design and stability of decentralized controller for vertical queueing networks. In a *vertical queueing network*, agents traveling across the network are stored in *point queues* that do not inhabit a “horizontal” position along the length of a network edge, but instead are considered to be stored in “vertical” stacks at each node. Such models are inherited from fields such as supply chain management or internet routing, but are also representative of signalized urban traffic networks and have recently regained popularity for use in applications of arterial traffic control [6, 1, 17, 20]. The concept of a *stabilizing* network controller, or one which ensures that the mean length of all queues in the network remain bounded, is relevant to many applications such as communications networks [8, 11, 12] or industrial systems [7, 4, 10].

In the present work we examine a vertical queueing model in which only a finite set of non-conflicting flow movements (or *phases*) can be permitted to proceed simultaneously across each network node. Phase actuation is dictated by a controller, such as a traffic light signaling specific combinations of green lights. Specifically, we consider application of a *maximum pressure* controller.

Max pressure is a distributed network control policy derived from the concept of a “back-pressure” or “MaxWeight” controller, which was first studied in the context of routing packets through a multi-hop communications network [16]. It has since been introduced to many other networked applications including process scheduling [5], manufacturing [15], wireless networks [3] and general stochastic networks [14]. The idea was applied to road traffic management more recently by [18] as well as [19]. The concept is intuitive: at each controlled network node, priority is given to the control phase which will be able to service the most demand given knowledge of both available upstream demand and the subsequent feasibility of downstream queues. It is a particularly attractive concept for control of a signalized urban traffic network because it can be operated in a distributed manner on local controller hardware (such as locally-managed traffic lights) but still provides theoretical guarantees on network-wide performance. Therefore unlike existing adaptive signal control systems such as SCOOT [9] or SCATS [13], max pressure does not require communications between each node or central control infrastructure. Also, max pressure is a universal algorithm which does not have to be optimized and manually tuned for a given network geometry or expected flows. In fact it operates with no a-priori knowledge of demand beyond a basic requirement of serviceability. This presents a huge benefit over most existing traffic control systems which require a timely and expensive re-timing process in the event of changes in demand patterns.

The original adaptation of this controller in [18], however, does not fully consider the practical limitations on the rate of queue measurement and signal actuation in vehicle traffic networks. For example, it has no bound on the rate of signal switches which may occur relative to the rate of modeled queue formation and dissipation. In implementation, a traffic signal incurs a penalty upon every change in actuation in the form of capacity loss due to “intersection clearance time”: a 2-3 second period where all movements are given a red light in order to allow traffic from the previous phase to clear the intersection before possibly conflicting traffic can be permitted to enter. Max pressure also lacks the ability to synchronize adjacent signals in a network by constraining the actuation periods of critical phases to fixed relative offsets. This feature is valued by traffic managers who wish to promote continuity of flow and limit vehicle stops on a preferred throughway. Furthermore, a standard max pressure implementation provides no explicit lower bound on the service rate

of queues on minor approaches where demand may be very low relative to the main direction.

These limitations motivate a new extension of the max pressure algorithm which bounds signal switches and can maintain timed cyclical behaviors for signal coordination and queue service equity. While a similar concept was suggested in [18], the current work further extends a simple proportional phase controller to allow model dynamics to explicitly act at a faster rate than the controller update period. We then extend the stability proof of [18] to prove that our *cycle-based max pressure* (Cb-MP) controller still provides the desired guarantee of queue stability with a penalty to the theoretical bound on queue lengths due to the decreased rate of controller update.

The remainder of this article is organized as follows: Sections 2-3 describe the modeling framework and standard max pressure controller from [18]; Section 4 formulates an extended cycle-based max pressure controller; Section 5 proves that this extended controller stabilizes a vertical queueing network; finally, Section 6 presents numerical results provided by this controller using a microscopic traffic simulation running in the Aimsun platform.

2. Network model dynamics. We consider a network of arterial roads with infinite storage capacity, modeled topologically as a graph with road segments (or *links*) being graph edges and intersections being nodes. We require that each link in the network has an exit path, or a continuous set of connected links on which vehicles can travel from the original link to eventually exit the network. This would in fact be the case in a practical road network.

An individual link $l \in \mathcal{L}$ can be either at the entry of the network ($l \in \mathcal{L}_{\text{ent}}$) or in the interior of the network ($l \in \mathcal{L} \setminus \mathcal{L}_{\text{ent}}$). The inflow on network entry links is defined entirely by a random *network demand* flow. We assume the expectation of this randomized boundary demand is known. The inflow on all other links (referred to as link demand) depends on observable queues on upstream links and the known set of constraints imposed by physical characteristics of the modeled network.

Each link in the network model can contain multiple *queues*, each of which correspond to separate *movements*: all vehicles in a single queue are intending to advance onto the same subsequent link. We describe the dynamics of network queues as a discrete time dynamical model using the following notation:

- A *movement* (l, m) distinguishes an intention to travel from link l to immediately downstream link m ;
- A *queue* $x(l, m)(t)$ is the number of vehicles on link l waiting to enter link m at time step t , and $X(t)$ is the set (vector or matrix) of all the queue lengths on the network at time step t ;
- A *link capacity* $C(l, m)(t)$ is the number of vehicles that can travel from link l to link m in time step t given maximum demand on queue $x(l, m)$. To model heterogenous driver behaviors we assume that $C(l, m)(t)$ is a random variable, but it has known expectation $c(l, m)$.
- Random variable $R(l, m)(t) = \frac{x(l, m)(t)}{\sum_k x(l, k)(t)}$ is the *turn ratio* at time t . R is assumed to have known expectation $r(l, m) \in [0, 1]$.
- The *demand vector* d of dimension $|\mathcal{L}_{\text{ent}}|$ specifies demands at network entry links;
- The *flow vector* f of dimension $|\mathcal{L}|$ denotes flows on all links of the network such that f_l is the flow in link l .

The expected flow through an entry link is defined to be equal to the expected demand on that link. Because we consider fixed turn ratios, the flow on an internal link is a weighted sum of the flows on the upstream neighboring links (with weights equal to the relevant expected turn proportions). Hence there is necessarily a linear relationship between the expected link flow f and the boundary demand d :

$$f = dP \quad (1)$$

where the (possibly non-unique) matrix P depends only on observed routing proportions within the network.

Intersection signal controller. A road intersection is modeled as a node in our framework. Controllers (traffic signals) are placed at every node to limit the set of queues permitted to discharge at any given time. A set of movements that can be simultaneously actuated without flow conflicts is called a *phase*. Each permissible phase for a given intersection can be represented as a binary control matrix S that is defined as follows:

$$S(l, m) = \begin{cases} 1 & \text{if movement } (l, m) \text{ is discharged in phase} \\ 0 & \text{otherwise} \end{cases}. \quad (2)$$

We denote the known, finite set of permissible control matrices for a single node as U . Note that U can be different for each node, as its construction depends on the number and physical arrangement of node entrances and egresses.

At each modeled time step t , a single control matrix $S(t) \in U$ encodes which single non-conflicting set of queues approaching the intersection are permitted to discharge during that time step. However in this work we consider a *relaxed controller* which operates on a finite set of contiguous time steps. This relaxed controller is defined as a matrix $\tilde{S}(t)$, which is populated with elements $\tilde{S}(l, m) \in [0, 1]$ that each represent *the fraction of the relevant set of time steps that are allocated to the movement $\{l, m\}$* . Such a relaxed controller must be able to be expressed as a convex combination of all $S \in U$,

$$\tilde{S} = \sum_{S \in U} \lambda_S S, \text{ with } \lambda_S \in [0, 1] \text{ and } \sum_{S \in U} \lambda_S = 1. \quad (3)$$

This work will demonstrate how to optimally select λ_S for a given set of time steps $(t + 1, t + 2, \dots, t + \tau)$ based only on the state of the network queues at the previous time period $(t - \tau, t - (\tau - 1), \dots, t)$.

Queue dynamics. The evolution of network queue lengths $X(t)$ can be seen as a Markov chain: the state at time $(t + 1)$ is a function of only the state at time t and external demand d ,

$$X(t + 1) = F(X(t), d). \quad (4)$$

Define $[a \wedge b] := \min\{a, b\}$. To describe queue dynamics explicitly, we must make a distinction between entry links and internal links: if $l \in \mathcal{L}_{\text{ent}}$,

$$x(l, m)(t + 1) = x(l, m)(t) + d_l(t + 1)R(l, m)(t + 1) - [C(l, m)(t + 1)S(l, m)(t + 1) \wedge x(l, m)(t)] \quad (5)$$

and if $l \in \mathcal{L} \setminus \mathcal{L}_{\text{ent}}$,

$$\begin{aligned} x(l, m)(t + 1) &= x(l, m)(t) \\ &+ \sum_k [C(k, l)(t + 1)S(k, l)(t + 1) \wedge x(k, l)(t)]R(l, m)(t + 1) \\ &- [C(l, m)(t + 1)S(l, m)(t + 1) \wedge x(l, m)(t)]. \end{aligned} \tag{6}$$

Demand feasibility. We focus on networks for which the boundary inflow demands are *feasible*—that is, the network is servicing a distribution of inflows for which it is possible to find a controller that allows *in average* more departures than arrivals at each link. Note that a stabilizing controller, defined as one that ensures that the mean length of all queues remains bounded, would not be possible by definition if demands were not feasible.

Define $\mathcal{C}(U)$ to be the convex hull of the set of permissible control matrices U . The following properties are then shown in [18]:

Property 1. A matrix M is in $\mathcal{C}(U)$ iff \exists a sequence of control matrices $\bar{S} = \{S(1), S(2), \dots, S(t), \dots | S(\cdot) \in U\}$ such that $\forall(l, m)$

$$M(l, m) = \liminf_T \frac{1}{T} \sum_{t=1}^T S(l, m)(t). \tag{7}$$

The element $M(l, m)$ in (7) can be interpreted as the long-term average proportion of intersection capacity given to movement (l, m) . Hence define $M_{\bar{S}}$ to be the long-term control proportion matrix constructed as in (7) using a specific control sequence $\bar{S} = \{S(1), S(2), \dots, S(t), \dots\}$.

Property 2. A demand d is feasible if and only if $\exists M_{\bar{S}} \in \mathcal{C}(U)$ and $\varepsilon > 0$ such that

$$c(l, m)M_{\bar{S}}(l, m) > f_l r(l, m) + \varepsilon \tag{8}$$

where $f = dP$ as in (1).

Define D^0 to be the set of all average demand vectors d that satisfy (8) and are therefore feasible.

A network is *stable* if the following quantity is bounded:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}\{|X(t)|_1\} \tag{9}$$

where $|X|_1 = \sum_{l,m} |x(l, m)|$ and the network state evolves according to dynamics under state dynamics (5)-(6).

3. Max pressure controller. Consider a weight assigned to each queue (l, m) as a function of all network queue lengths X :

$$w(l, m)(X(t)) = x(l, m)(t) - \sum_{p \in \text{Out}(m)} r(m, p)x(m, p)(t) \tag{10}$$

where $\text{Out}(m)$ is the set of all links receiving flow from link m . The *pressure* $\gamma(S)$ that is potentially alleviated by a control action S at time step t is defined as follows:

$$\gamma(S)(X(t)) = \sum_{l,m} c(l, m)w(l, m)(X(t))S(l, m)(t). \tag{11}$$

At each time step t , the standard max pressure controller $u^*(X(t))$ explicitly chooses the phase $S^* \in U$ that maximizes $\gamma(S)(X(t))$:

$$S^*(t) = u^*(X(t)) = \arg \max\{\gamma(S)(X(t)) | S \in U\}. \quad (12)$$

Reference [18] shows the following stability result for the standard max pressure controller:

Theorem 3.1. *The max pressure control u^* is stabilizing, i.e. the quantity in equation (9) remains bounded, whenever the average demand vector $d = \{d_i\}$ is within the set of feasible demands D^0 .*

We omit further explanation of the above theorem in this paper for brevity.

This theoretical guarantee is one of the many attractive qualities of max pressure for controlling vehicular traffic in urban road networks. Yet the controller as originally formulated is not practical for application on a signalized traffic network for three reasons:

- a) it does not account for capacity reductions (lost time) due to excessive signal switching,
- b) it cannot enforce coordination between subsequent intersections for purposes of maximizing flow continuity, and
- c) it does not provide guarantees that low-demand queues will be served within a finite time period.

4. Cycle-based max pressure (Cb-MP). The above limitations motivate our extension of the standard immediate feedback max pressure control algorithm. In the following section, we define a new *cycle-based max pressure* (Cb-MP) controller which bounds the number of signal switches per fixed time period, provides capacity for standard signal coordination methods, and can easily guarantee a minimum service rate for all intersection phases.

For safety reasons, an intersection controller cannot switch signal phase actuation immediately. Instead, it must incorporate a pause of $R \approx 2.5$ seconds in which all signal phases have a red light. This *clearance time* allows all vehicles in the previously actuated phase to clear the intersection before a conflicting phase can be permitted to use the intersection. In the standard formulation of max pressure, the controller chooses an appropriate action based feedback received at every time step of the modeled dynamics. To accurately capture queuing behaviors observed on arterial roadways, a model would need to operate with a time discretization of $\Delta t < 10$ seconds. A signal switch at every time step could therefore result in more than 25% loss of intersection service capacity, which is not considered in the theoretical examination presented in [18].

In this work, we explicitly specify the number of signal switches that occur in a fixed number of model time steps using the familiar concept of a *signal cycle*. As typical with modern traffic signals Cb-MP rotates through all available signal phases within a known time period. We define *cycle time* τ as a predefined number of model time steps and require that each controller phase S must be green for some proportion $\lambda_S \geq \kappa_S$ of the τ steps, where the *minimum green splits* $\kappa_S \in (0, 1) \forall S \in U$ are parameters selected by a network manager to enforce equity in movement actuation.

The selection of a cycle time τ intuitively affects intersection capacity. Our proof of network stability in the following sections relies on the fact that road links are *undersaturated*: that is, the expected demand is served (on average) within a signal

cycle. To avoid link saturation, we pose the following convex optimization problem (extended from that in [2]) to determine minimum constrained feasible actuation time Λ^* :

$$\begin{aligned} \Lambda^* &= \min_{\lambda=\{\lambda_S\}} \sum_{S \in U} \lambda_S \\ \text{subject to} \quad &\lambda_S \geq \kappa_S \quad \forall S \in U \\ &f_l r(l, m) < \sum_S \lambda_S c(l, m) S(l, m) \end{aligned} \quad (13)$$

where $\kappa_S \in [0, 1] \quad \forall S \in U$ and $\sum_S \kappa_S < 1$. If $\Lambda^* > 1$, the demand is not feasible under the set of $\{\kappa_S\}$ for any cycle length. If $\Lambda^* < 1$, then we can define a cycle length for which flow is admissible without link saturation. However, this cycle length τ must be significantly greater than Λ^* to account for clearance times. If we define $L = \lceil (\frac{R}{\Delta t} \cdot |U|) \rceil$ to be the total number of *lost time steps* per cycle, a feasible cycle length τ must satisfy the following condition:

$$\tau > \frac{L}{1 - \Lambda^*}. \quad (14)$$

Given an appropriate τ which satisfies (14), the cycle-based max pressure controller is a relaxed control matrix that is constructed as follows:

$$\tilde{S}^*(t) = u^{c^*}(X(t)) = \sum_{S \in U} \lambda_S^* S, \quad (15)$$

$$\text{where } \{\lambda_S^*\} = \arg \max_{\lambda_1, \dots, \lambda_{|U|}} \sum_{S \in U} \lambda_S \gamma(S)(X(\lfloor t/\tau \rfloor)) \quad (16)$$

$$\text{subject to } \lambda_S \geq \kappa_S, \quad \sum_S \lambda_S \leq 1 - \frac{L}{\tau}.$$

At time step $t = n\tau$ for integer n , the controller u^{c^*} uses feedback measurements $x(t)$ to select a relaxed control matrix \tilde{S}^* with components λ_S^* that satisfy (16). This relaxed controller is then applied for the subsequent τ time steps $\{t, t+1, \dots, t+\tau-1\}$ before the controller is updated.

Note that this controller is modeled such that all phases in an intersection are simultaneously actuated at some proportion of their maximum flow capacity. This is not possible in practice, as many phases will have to make conflicting use of the same intersection resources. Hence individual phases S will have to be actuated in series, with each having a duration corresponding to a number of “time units” that are equal to cycle proportions ($\lambda_S \tau \cdot \Delta t$). Feedback measurements will then be a measure of “average” cycle queue length acquired over a set of measurements spanning the previous cycle. If the solution for (16) is not unique, one of the optimal solutions is chosen at random using a uniform probability distribution (or according to some practical actuation priority criteria chosen by the network manager). Because it can be implemented such that phases occur in a predictable order, a controller running Cb-MP can be synchronized with neighboring controllers to enforce a “green-wave progression” as is standard practice in existing traffic signal control design.

The Cb-MP controller formulated in (15)-(16) is fundamentally different from the standard max pressure formulation in [18] in two ways: first, it only updates the controller once every signal cycle (or τ model time steps); second, it applies a relaxed phase actuation (which is some convex combination of standard phase actuations). In the next section, we address how each of these modifications impact the resulting network dynamics, and ultimately show that the application of Cb-MP

yields a similar stability guarantee to that shown by Varaiya for the standard max pressure controller given slightly weaker conditions on demand flow.

5. Stability of Cb-MP. Define \mathbb{C}_κ as the set of convex combinations of control matrices with coefficients larger than κ :

$$\mathbb{C}_\kappa = \left\{ \sum_S \lambda_S S \mid \lambda_S > \kappa_S \forall S \in U \right\}. \quad (17)$$

Also define a set of *undersaturated* admissible demands D_κ with elements d such that $f = dP$ and

$$fir(l, m) < c(l, m)\tilde{S}(l, m). \quad (18)$$

This condition (also seen in (13)) ensures that a demand $d \in D_\kappa$ can in average be served *within a single cycle* by a relaxed control matrix that maintains a specified minimum time allocation for each phase.

Theorem 5.1. *The cycle-based max pressure controller defined in (15)-(16) stabilizes a network whenever the demand is within a set of feasible undersaturated demands D^κ .*

The remainder of this section proves Theorem 5.1 by finding a bound on (9) given a cycle-based max pressure controller. The structure of this proof is as follows:

- 1) First, we introduce the concept of a τ -updated controller and we show that switching control only once every τ time steps does not impact the set of feasible demands.
- 2) Then we explain the mathematical structure used in [18] to derive a bound for the expected network state (9).
- 3) Next we define an intermediate “relaxed max pressure” formulation to demonstrate the impacts of expanding the domain of control actions to relaxed controllers which are convex combinations of allowable phase matrices.
- 4) We then demonstrate the intra-cycle queue dynamics given a τ -updated relaxed controller.
- 5) We combine the above steps to show that queue stability holds given a Cb-MP controller with both relaxed actuation and τ -updating.
- 6) Finally, we compare the our Cb-MP queue length bounds to those originally derived in [18] to illustrate the increase due to cycle-updating.

5.1. Properties of a τ -updated controller. Suppose that we impose that the control actuation $S^*(t)$ can only be updated every τ model time steps. A resulting τ -updated control sequence is composed of a single control matrix repeated for τ time steps of the model dynamics:

$$S(n\tau + 1) = S(n\tau + 2) = \dots = S((n + 1)\tau). \quad (19)$$

In the Appendix of this paper, we prove that the set of demands that can be accommodated using τ -updated control sequences is the same set of feasible flows as in (8). This equivalence becomes intuitive when one considers that our definition of feasible flows considers only the long-term (more precisely, infinite-term) average of both demand and service rates, and any infinite control sequence with limited admissible phases can be re-arranged to form a τ -updated sequence for some τ .

As will be shown in Section 5.6, occasional updating will also lead to an increased bound on queue lengths relative to the standard max pressure setting.

5.2. Formulating a queue bound. Our ultimate goal is to derive a bound for the average expected queue state (9). The approach taken in this work follows that of [18]: we bound the incremental model-step queue increase $|X(t+1) - X(t)|$ and then recursively compute a bound on average queue lengths $\sum_{t \in [0, T]} \mathbb{E}\{|X(t)|\}$ for an arbitrary time horizon T .

Begin by considering the expectation of the following function of queue state perturbation conditioned on the past queue state:

$$\begin{aligned} |X(t+1)|^2 - |X(t)|^2 &= |X(t) + \delta(t)|^2 - |X(t)|^2 \\ &= 2X(t)^T \delta(t) + |\delta(t)|^2 \\ &= 2\alpha(t) + \beta(t) \end{aligned} \quad (20)$$

with $\delta(t) = X(t+1) - X(t)$, $\alpha(t) = X(t)^T \delta(t)$, and $\beta(t) = |\delta(t)|^2$. We continue by addressing bounds on β and α separately.

First we consider $\beta(t) = |\delta(t)|^2$.

Lemma 5.2.

$$\beta(t) = |\delta(t)|^2 \leq NB^2 \quad (21)$$

where $B = \max\{\bar{C}(l, m), \sum_k \bar{C}(k, l), \bar{d}(l, m)\}$, N is the number of queues in the network, $\bar{C}(l, m)$ is the maximum value of the random service rate $C(l, m)(t)$, and $\bar{d}(l, m)$ is the maximum value of random demand $d(l, m)$.

The proof of Lemma 5.2 is exactly the same as presented in [18] and will therefore not be repeated here. Note that because these bounds hold for any arbitrary $S(l, m)(t) \in [0, 1]$, this bound on β is trivially extended to any convex combination of control matrices; hence it is still valid in our extension.

Now we examine a bound on $\alpha(t) = X(t)^T \delta(t)$. Again following [18], we define additional sub-terms:

$$\begin{aligned} \mathbb{E}\{\alpha(t)|X(t)\} &= \sum_{l \in \mathcal{L}, m} w(l, m)(t) \cdot \left[f_{lr}(l, m) \right. \\ &\quad \left. - \mathbb{E}\left\{ [C(l, m)(t+1)S(l, m)(t) \wedge x(l, m)(t)] | X(t) \right\} \right] \\ &= \alpha_1(t) + \alpha_2(t) \end{aligned} \quad (22)$$

with

$$\alpha_1(t) = \sum_{l \in \mathcal{L}, m} [f_{lr}(l, m) - c(l, m)S(l, m)(t)] w(l, m)(t) \quad (23)$$

$$\begin{aligned} \alpha_2(t) &= \sum_{l \in \mathcal{L}, m} S(l, m)(t) w(l, m)(t) \cdot \left[c(l, m) - \right. \\ &\quad \left. \mathbb{E}\left\{ [C(l, m)(t+1) \wedge x(l, m)(t)] | X(t) \right\} \right]. \end{aligned} \quad (24)$$

Lemma 5.3. For all l, m, t ,

$$\alpha_2(t) \leq \sum_{l \in \mathcal{L}, m} c(l, m) \bar{C}(l, m). \quad (25)$$

The proof of Lemma 5.3 again directly follows that presented in [18]; an extension from a binary controller $S \in \{0, 1\}$ to a relaxed controller $\tilde{S} \in [0, 1]$ is trivial.

In fact, the extension made here only affects the $\alpha_1(t)$ term. To demonstrate a bound on $\alpha_1(t)$ given application of a cycle-based max pressure controller u^{c*} ,

we first examine the stability of a standard max pressure controller using relaxed controllers with minimum phase proportion constraints and a stricter limitation on network demands. We will then show that a τ -updated cycle-based max pressure controller also stabilizes a network, but results in an increase in queue length bounds that is proportional to cycle length τ .

5.3. Effect of a relaxed controller. Define an intermediate “relaxed max pressure” policy in which relaxed controllers are applied at the standard max pressure update rate (once per time step of the model dynamics). This situation was suggested in [18] to simulate enforcing minimum phase proportions in a cycle formulation of max pressure. Yet this proposal unrealistically models “cycle” updates at the same rate as the model of queueing and discharging behaviors (hence the introduction of the τ -updated formulation in this work). Nonetheless, we use this intermediate formulation to demonstrate that queue stability is still achieved upon use of a relaxed controller.

Lemma 5.4. *If a “relaxed” max pressure control policy \tilde{S}^* is updated and applied at each time step t and the demand d is in the set of feasible undersaturated demands D^* , then there exists an $\varepsilon > 0$, $\eta > 0$ such that*

$$\alpha_1(t) \leq -\varepsilon\eta|X(t)|. \quad (26)$$

Proof of Lemma 5.4. Consider the relaxed max pressure control matrix \tilde{S}^* defined in (15) for $\tau = 1$. By construction, $\forall \tilde{S} \in \mathbb{C}_\kappa$

$$\sum_{l,m} c(l,m)w(l,m)(X(t))\tilde{S}(l,m) \leq \sum_{l,m} c(l,m)w(l,m)(X(t))\tilde{S}^*(l,m)$$

with equality only if $\tilde{S} = \tilde{S}^*$. Thus $\forall (\tilde{S} \in \mathbb{C}_\kappa) \neq \tilde{S}^*$,

$$\begin{aligned} \sum_{l,m} [f_l r(l,m) - c(l,m)\tilde{S}^*(l,m)(t)]w(l,m)(X(t)) \\ < \sum_{l,m} [f_l r(l,m) - c(l,m)\tilde{S}(l,m)]w(l,m)(X(t)). \end{aligned} \quad (27)$$

If the demand flow is admissible according to (18), then $\exists \hat{S} \in \mathbb{C}_\kappa$ and some small $\varepsilon > 0$ such that

$$c(l,m)\hat{S}(l,m) = \begin{cases} f_l r(l,m) + \varepsilon & \text{if } w(l,m)(X(t)) > 0 \\ 0 & \text{otherwise} \end{cases}.$$

Therefore,

$$\begin{aligned} \alpha_1(t) &= \sum_{l,m} [f_l r(l,m) - c(l,m)\tilde{S}^*(l,m)(t)]w(l,m)(X(t)) \\ &< \sum_{l,m} [f_l r(l,m) - c(l,m)\hat{S}(l,m)(t)]w(l,m)(X(t)) \\ &= -\varepsilon \sum_{l \in \mathcal{L}, m} \max\{w(l,m)(X(t)), 0\} \\ &\quad + \sum_{l \in \mathcal{L}, m} f_l r(l,m) \min\{w(l,m)(X(t)), 0\}. \end{aligned} \quad (28)$$

We assume that by our choice of \hat{S} , $f_l r(l,m) > \varepsilon$. Hence $\alpha_1(t) < -\varepsilon \sum_{l,m} w(l,m)(t)$. Given the linearity of (10) and the known properties of $r(l,m)(t)$, it can be show

that $\sum_{l,m} w(l,m)(t) \geq \eta|X(t)|$ for some $\eta > 0$. This completes the derivation of (26). \square

For ease of notation, now combine (21), (25) and (26) to obtain the following expression given application of the “relaxed max pressure” controller:

$$\begin{aligned} \mathbb{E}\left\{|X(t+1)|^2 - |X(t)|^2|X(t)\right\} &= \mathbb{E}\left\{2\alpha(t) + \beta(t)\right\} \\ &< -2\varepsilon\eta|X(t)| + 2 \sum_{l \in \mathcal{L}, m} [c(l,m)\bar{C}(l,m)] + NB^2 \end{aligned} \quad (29)$$

where N and B are as in (21).

5.4. Intra-cycle queue bound. Next we establish a bound on queue growth in a single time step between controller updates.

Lemma 5.5. *Assuming a cycle-based max pressure controller with an cycle steps τ beginning at time t , the following bound on state perturbation holds for all steps since update $p \in [0, \tau - 1]$:*

$$\begin{aligned} \mathbb{E}\left\{|X(t+p+1)|^2 - |X(t+p)|^2|X(t) \dots X(t+p)\right\} \\ < Y + h(p) - 2\varepsilon\eta|X(t+p)| \end{aligned} \quad (30)$$

$$\text{for } Y = 2 \sum_{l,m} c(l,m)\bar{C}(l,m) + NB^2 \quad \text{and} \quad (31)$$

$$h(p) = 2pNB \left(\varepsilon\eta + \sum_{l,m} [f_{lr}(l,m) + c(l,m)] \right). \quad (32)$$

Proof of Lemma 5.5. As in Lemmas 5.2-5.4 above, begin by dividing the argument of (30) into three parts: $|X(t+p+1)|^2 - |X(t+p)|^2 = 2(\alpha_1(t+p) + \alpha_2(t+p)) + \beta(t+p)$, where β , α_1 and α_2 are quantities that depend on the controller applied at $(t+p)$:

$$\beta(t+p) = |X(t+p+1) - X(t+p)|^2 \quad (33)$$

$$\alpha_1(t+p) = w(l,m)(X(t+p)) \cdot \sum_{l,m} \left(f_{lr}(l,m) - c(l,m)S(l,m)(t) \right) \quad (34)$$

$$\begin{aligned} \alpha_2(t+p) &= w(l,m)(X(t+p)) \cdot \sum_{l,m} \left(c(l,m)S(l,m)(t) \right. \\ &\quad \left. - \mathbb{E}\left\{[C(l,m)(t+p+1) \wedge x(l,m)(t+p)]|X(t+p)\right\} \right) \end{aligned} \quad (35)$$

Bounds on the expectations of $\beta(\cdot)$ and $\alpha_2(\cdot)$ were previously established for any binary or relaxed control matrix in (21) and (25), respectively. Thus we already know that:

$$\begin{aligned} \mathbb{E}\left\{|X(t+p+1)|^2 - |X(t+p)|^2|X(t) \dots X(t+p-1)\right\} \\ < 2 \sum_{l,m} c(l,m)\bar{C}(l,m) + NB^2 + \mathbb{E}\left\{2\alpha_1(t+p)\right\}. \end{aligned} \quad (36)$$

The remainder of the bound proposed in (30) originates from the $2\alpha_1(t+p)$ term, which is directly dependent on the explicit form of the controller S . Rewrite $2 \cdot \alpha_1$

from (34) as follows:

$$\begin{aligned}
& 2 \sum_{l,m} w(l,m)(X(t+p))[f_l r(l,m) - c(l,m)S(l,m)(t)] \\
&= 2 \sum_{l,m} w(l,m)(X(t))[f_l r(l,m) - c(l,m)S(l,m)(t)] \\
&\quad + 2 \sum_{l,m} \left\{ w(l,m)(X(t+p) - X(t)) \cdot [f_l r(l,m) - c(l,m)S(l,m)(t)] \right\} \\
&= \xi_1(t,p,S) + \xi_2(t,p,S)
\end{aligned}$$

for

$$\xi_1(t,S) = 2 \sum_{l,m} w(l,m)(X(t))[f_l r(l,m) - c(l,m)S(l,m)(t)]$$

and

$$\xi_2(t,p,S) = 2 \sum_{l,m} \left\{ w(l,m)(X(t+p) - X(t)) \cdot [f_l r(l,m) - c(l,m)S(l,m)(t)] \right\}.$$

By Lemma 5.4 we know that $\xi_1(t,S) < -2\varepsilon\eta|X(t)|$. Because $|X(t)| = |X(t+p) - (X(t+p) - X(t))| > |X(t+p)| - |X(t+p) - X(t)|$, we find that

$$\begin{aligned}
\xi_1(t,p,S) &< -2\varepsilon\eta(|X(t+p)| - |X(t+p) - X(t)|) \\
&< -2\varepsilon\eta|X(t+p)| + 2\varepsilon\eta \sum_{i=1}^p |X(t+i) - X(t+i-1)| \\
&= -2\varepsilon\eta|X(t+p)| + 2\varepsilon\eta \sum_{i=1}^p |\delta(t+i-1)|. \tag{37}
\end{aligned}$$

So by (37) and (21),

$$\begin{aligned}
\xi_1(t,S) &< 2\varepsilon\eta p \sum_{l,m} \max \left\{ \bar{C}(l,m), \sum_k \bar{C}(k,l), \bar{d}(l,m) \right\} - 2\varepsilon\eta|X(t+p)| \\
&= 2\varepsilon\eta \cdot (pNB - |X(t+p)|). \tag{38}
\end{aligned}$$

To bound ξ_2 , we study the term

$$\begin{aligned}
& w(l,m)(X(t+p)) - w(l,m)(X(t)) \\
&= \sum_{n=1}^p w(l,m)(X(t+n)) - w(l,m)(X(t+n-1)) \\
&= \sum_{n=1}^p \left\{ x(l,m)(t+n) - x(l,m)(t+n-1) \right. \\
&\quad \left. - \sum_{s \in \text{Out}(m)} [x(m,s)(t+n) - x(m,s)(t+n-1)] \cdot r(m,s) \right\} \\
&= \sum_{n=1}^p w(l,m)(\delta(t+n-1)) \tag{39}
\end{aligned}$$

By (21) and the fact that $w(\cdot)$ is linear,

$$|w(l,m)(\delta(t+n-1))| < NB. \tag{40}$$

Plugging (40) back into the definition of ξ_2 , we obtain

$$\begin{aligned}
\xi_2(t, p, S) &= 2 \left(\sum_{l, m} [f_{lr}(l, m) - c(l, m)S(l, m)(t)] \cdot \sum_{n=1}^p w(l, m)(\delta(t+n-1)) \right) \\
&< 2 \sum_{n=1}^p \sum_{l, m} [f_{lr}(l, m) - c(l, m)S(l, m)(t)] \cdot \sum_{u, v} \max \left\{ \bar{C}(u, v), \sum_k \bar{C}(k, u), \bar{d}(u, v) \right\} \\
&= 2NB \sum_{n=1}^p \sum_{l, m} [f_{lr}(l, m) - c(l, m)S(l, m)(t)]. \tag{41}
\end{aligned}$$

Also note that

$$\left| \sum_{n=1}^p \sum_{l, m} [f_{lr}(l, m) - c(l, m)S(l, m)(t)] \right| < p \sum_{l, m} [f_{lr}(l, m) + c(l, m)] \tag{42}$$

so (41) becomes

$$\xi_2(t, p, S) < 2NBp \cdot \left(\sum_{l, m} [f_{lr}(l, m) + c(l, m)] \right). \tag{43}$$

Substituting (38) and (43) into (36) yields (30). \square

5.5. Network queue bound. Given Lemmas 5.2-5.5, we show that for a time t within any number K of τ -updated cycles, the following quantity is bounded:

$$\begin{aligned}
&\sum_{t=1}^{K\tau} \mathbb{E} \left\{ |X(t+1)|^2 - |X(t)|^2 |X(t) \right\} \\
&= \sum_{t=1}^{K-1} \sum_{p=0}^{\tau-1} \mathbb{E} \left\{ |X(t+p+1)|^2 - |X(t+p)|^2 |X(t+p) \right\} \\
&< \sum_{t=1}^{K-1} \sum_{p=0}^{\tau-1} (Y + h(p) - 2\varepsilon\eta |X(t+p)|) \\
&< -2\varepsilon\eta \sum_{t=1}^{K\tau} |X(t)| + (K-1) \left(\tau Y + \sum_{p=0}^{\tau-1} h(p) \right) \tag{44}
\end{aligned}$$

which, when taking the expectation, yields

$$\mathbb{E} \left\{ |X(K\tau+1)|^2 - |X(1)|^2 \right\} < -2\varepsilon\eta \sum_{t=1}^{K\tau} \mathbb{E} \left\{ |X(t)| \right\} + (K-1) (\tau Y + \sum_{p=0}^{\tau-1} h(p)).$$

Rearranging gives

$$\begin{aligned}
&\frac{1}{K\tau} \sum_{t=1}^{K\tau} \mathbb{E} \left\{ |X(t)| \right\} \\
&< \frac{1}{2\varepsilon\eta K\tau} \mathbb{E} \left\{ |X(1)|^2 - |X(K\tau+1)|^2 \right\} + \frac{\tau-1}{2\varepsilon\eta K\tau} \left(\sum_{p=0}^{\tau-1} h(p) + \tau Y \right) \\
&< \frac{1}{2\varepsilon\eta K\tau} \mathbb{E} \left\{ |X(1)|^2 \right\} + \frac{1}{2\varepsilon\eta\tau} \left(\sum_{p=0}^{\tau-1} h(p) + \tau Y \right). \tag{45}
\end{aligned}$$

By (9), the bound

$$2\varepsilon\eta\frac{1}{K\tau}\sum_{t=1}^{K\tau}\mathbb{E}\{|X(t)|\} < \frac{1}{K\tau}\mathbb{E}\{|X(1)|^2\} + \frac{1}{\tau}\sum_{p=0}^{\tau-1}h(p) + Y \quad (46)$$

establishes that the cycle-based max pressure controller $u^{c^*}(X(t))$ defined in (15) will stabilize a vertical queueing network with dynamics $X(t)$ as in (5)-(6).

5.6. Increase in queue bounds. The following bound on network queue state for a standard max pressure controller is derived in [18] (for Y as in (31)):

$$2\varepsilon\eta\frac{1}{T}\sum_{t=1}^T\mathbb{E}\{|X(t)|\} < \frac{1}{T}\mathbb{E}\{|X(1)|^2\} + Y. \quad (47)$$

Notice that by comparison between (47) and (46), the bound on the long-term sum of expected network queues in cycle-based max pressure is larger by a term that increases linearly in cycle length τ :

$$\frac{1}{2\varepsilon\eta\tau}\sum_{p=0}^{\tau-1}h(p) = (\tau-1)NB\left(1 + \frac{1}{\varepsilon\eta}\sum_{l,m}[f_l r(l,m) + c(l,m)]\right). \quad (48)$$

6. Numerical implementation. To demonstrate the effectiveness of our algorithm, a cycle-based max pressure controller was implemented on a network of 11 signalized intersections modeled in the Aimsun, a commonly-used microsimulation platform. While microsimulation dynamics do not precisely represent the queuing behaviors represented above, the following results are provided as a valuable proof-of-concept of this controller in an environment that is actually more realistic and applicable than the vertical queueing model described in this work.

The specific model used in this work was generated as part of the I-15 Integrated Corridor Management project undertaken by the San Diego Association of Governments (SANDAG) in San Diego, CA. Demand and other model parameters are calibrated to match the morning peak period (5:00 AM to 10:00 AM).




FIGURE 1. The chosen network was calibrated to represent realistic demands and physical parameters observed on a stretch of Black Mountain Road near the I-15 freeway in San Diego, California.

This section of road is currently controlled using an offset-optimized actuated-coordinated control scheme. Under this system, each signal operates with a fixed cycle time of 100 seconds and a fixed phase ordering, but uses instantaneous feedback of intersection vehicle approaches to adjust cycle green splits (effectively λ_i)

within fixed minimum and maximum green time constraints per cycle. This control algorithm is coded and calibrated in the SANDAG Aimsun model to represent realistic conditions and was therefore deemed an appropriate benchmark for performance comparison to Cb-MP.

Six variations of Cb-MP were implemented. First, a version with a cycle length of 100 seconds and minimum green time constraints of 10 seconds for each available signal phase was used to closely match the operational constraints of the existing fully-actuated controller. The relative offsets for the southbound coordination phases in this implementation were the same as those used in the actuated-coordinated system. We then ran variations which extended the cycle time for Cb-MP to 120, 140, 160, 180, and 200 seconds to demonstrate the effect of increased cycle time τ on observed queue lengths.

To compare performance, we calculated vehicle service rates, average delay, average number of stops and stopped time, and mean and maximum queue lengths that were modeled using each controller. These metrics were only calculated for vehicles and links corresponding to the southbound direction on Black Mountain Road as well as the short connections to the I-15 freeway on westbound Mira Mesa Blvd and eastbound Mirimar Rd. This pathway simulates a viable “freeway-alternative” in the congested direction during the morning peak period. During implementation, the Cb-MP algorithm most often chose to give actuation priority to this high-demand Southbound direction, as expected.

The comparison of network vehicle counts in Figure 2 suggests that Cb-MP is able to service approximately the same volumes as the optimized actuated-coordinated control when cycle times were comparable. The higher cycle length Cb-MP implementations are omitted from this plot for clarity; these controllers resulted in higher variations of vehicle service between 5-minute periods but ultimately only reduced total service rates slightly.

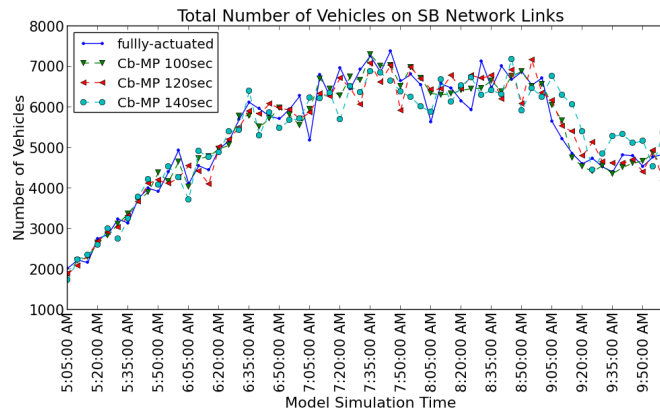


FIGURE 2. Cb-MP demonstrated service rates that are consistent with a fully-actuated control system for similar cycle lengths.

Yet distinct differences between the fully-actuated and Cb-MP controllers were observed in measurements of delays. Figure 3 compares the average vehicle delay given fully-actuated and max pressure control with the same cycle length. It is apparent that while the fully-actuated controller produces less delay when demand is far below network capacity, Cb-MP outperforms the existing controller given

consistently high demand; it imposes less delay with a noticeably smaller variance. This may not be surprising given known deficiencies in actuated controllers, however it is important to point out that this implementation of max pressure produces very promising network delays with almost no controller parameters that require tuning.




FIGURE 3. Cb-MP outperforms the actuated controller given high demand in terms of vehicle delay.

Despite maintaining the same relative offsets for actuation of the main (coordination) direction as the fully actuated controller, Cb-MP induced slightly more stops during a vehicle’s south-bound journey across the network. Again, this may be expected given the stop-minimizing design objectives of the fully-actuated system: the small but consistent differences in average vehicle-stops shown in Figure 4 are likely caused by the on-demand service extensions provided for low density “back-of-queue” arrivals in the fully-actuated system. Notice that the average vehicle stopped time is actually lower in Cb-MP than with the fully actuated system during peak demand, which is consistent with the estimates of total delay demonstrated in Figure 3.

The higher cycle length Cb-MP implementations are again omitted from Figures 3-4 for clarity, yet it is important to note that higher cycle lengths predictably led the longer stops and more delay, as vehicles which encountered a red light would have to wait longer for the cycle to reach their desired green phase. This increase in wait time also corresponds to the predicted larger queues.

Figure 5 demonstrates the increase in mean queue lengths with increase in cycle length τ . While the linear increase in maximum queue length derived in (48) is not explicitly depicted in the observations, these results appear to remain consistent with such an upper bound.

7. Conclusion. The numerical implementation of cycle-based max pressure in Section 6 suggests a promising alternative for signal control in periods of high demand where the performance of existing actuated controllers is known to deteriorate. It is an intuitive and scalable control algorithm that is appealing because it maintains theoretical network-wide performance guarantees without the need for centralized communication or control centers. The cyclical operation of Cb-MP can still maintain the flow-progression benefits obtained from existing offset optimization algorithms along a prioritized route, as demonstrated by the fact that the average number of vehicle-stops on the southbound route only increases slightly using Cb-MP

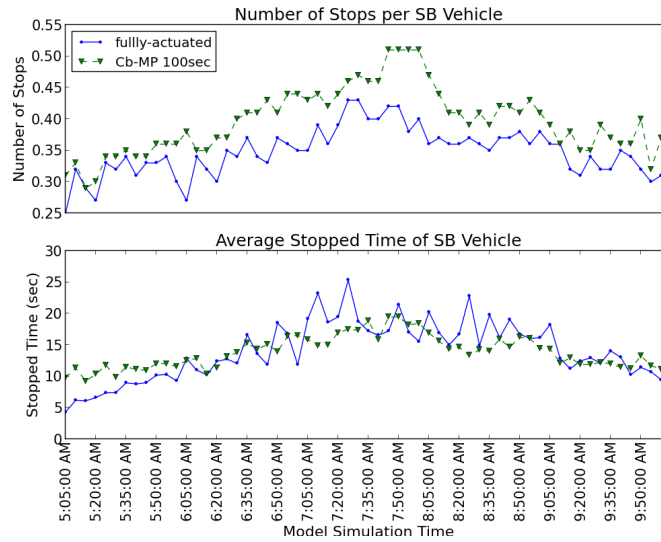


FIGURE 4. While Cb-MP caused more vehicle stop events, stoppage times were similar to those observed using the actuated controller.




FIGURE 5. Observed queues increase with cycle length using Cb-MP control.

over an implementation of the optimized actuated system. Because the cycle splits are more predictable in Cb-MP than in current actuated-coordinated algorithms,

it may even be possible to further optimize progression on multiple (conflicting) routes with additional linear constraints on cycle splits in (16).

Furthermore, Cb-MP is a widely-applicable algorithm which requires significantly less tuning and site-specific adjustment than the typical fully-actuated control system. For example, the timing parameters for a fully-actuation system deployed on networks such as the SANDAG site referenced above are often a result of many hours of both model-based and heuristic optimization procedures for a specific network geometry and expected demand. Yet in our implementation, the generalized Cb-MP algorithm with arbitrary reasonable minimum green parameters achieved approximately equal performance without requiring any knowledge of location or demand.

By addressing the practical problem of lost capacity due to frequent switching, our extension of the proofs of [18] brings the concept of max pressure closer to a realistic implementation. Yet it is important to point out that existing sensing infrastructure does not typically provide the capabilities necessary to accurately measure approaching link-counts, nor does it provide any measurement of downstream link state.

Future work should also address limitations of the unrealistic assumptions inherent in the vertical queueing model framework, such as the concept of infinite link buffer size. In practice, over-saturation could cause unmodeled instabilities in a traffic network if expected queue bounds exceed physical road storage.

One could also consider multiple ways in which the performance of Cb-MP could benefit from heuristic modification, such as enforcing a maximum value of the green split λ_l as a function of queue measurement x_l to prevent wasted green in a given cycle. Our extended proof of stability should hold for any additional linear constraints on λ in (16) (which maintain the concept of a relaxed, τ -updated controller) given that there exists any controller satisfying these constraints for which the network demand is feasible according to (18).

Appendix A. Feasible flows with a τ -updated controller.

Lemma A.1. *All flows which satisfy Property 2 given a controller u updated at every model time step will also satisfy Property 2 with a τ -updated controller for some τ .*

Proof of Lemma A.1. Given the set admissible phases U , define:

- \mathcal{U} is the set of control sequences with distinct elements
 $\{S(1), S(2) \dots S(t) \dots | S(\cdot) \in U\}$,
- \mathcal{U}_τ is the set of τ -updated control sequences
 $\{S(1), S(1), \dots, S(\tau + 1), S(\tau + 1), \dots, S(n\tau + 1), S(n\tau + 1), \dots | S(\cdot) \in U\}$,

Also define the following sets of *long-term control proportion matrices*, which are similar to the formulation in (7):

$$M_{\mathcal{U}} = \left\{ \liminf_T \frac{1}{T} \sum_{t=1}^T S(t) \mid \{S(1), S(2), \dots, S(t), \dots\} \in \mathcal{U} \right\}$$

$$M_{\mathcal{U}_\tau} = \left\{ \liminf_T \frac{1}{T} \sum_{t=1}^T S(t) \cdot \left\{ S(1), S(1), \dots, S(\tau + 1), \right. \right.$$

$$\left. \left. S(\tau + 1), \dots \right\} \in \mathcal{U}_\tau \right\}$$

By Property 2, a demand d is only feasible if there exists a control sequence \bar{S} such that the corresponding long-term control proportion matrix $M_{\bar{S}}$ satisfies (8). Here we show $M_{\mathcal{U}} = M_{\mathcal{U}_\tau}$, and therefore any flows that are admissible given an unrestricted controller in \mathcal{U} can also be accommodated using a τ -updated controller in \mathcal{U}_τ .

Obviously, $M_{\mathcal{U}_\tau} \subset M_{\mathcal{U}}$. To show equality, we must also demonstrate that $M_{\mathcal{U}} \subset M_{\mathcal{U}_\tau}$. Suppose there exists a control sequence $\hat{S} = \{S(1), S(2), \dots\} \in \mathcal{U}$. By definition,

$$\begin{aligned} M_{\hat{S}} &= \liminf_T \frac{1}{T} \sum_{t=1}^T S(t) = \liminf_T \frac{1}{\tau T} \sum_{t=1}^{\tau T} S'(t) \\ &\quad \text{where } S' = \{S(1), S(1), \dots, S(t), S(t), \dots\} \\ &= \liminf_T \frac{1}{T} \sum_{t=1}^T S'(t) \in M_{\mathcal{U}_\tau} \quad \implies M_{\mathcal{U}} \subset M_{\mathcal{U}_\tau} \end{aligned}$$

□

Acknowledgments. The authors would like to thank Pravin Varaiya for his insights on this work, as well as Brian Peterson and Joe Butler for their support in the numerical implementation.

REFERENCES

- [1] K. Aboudolas, M. Papageorgiou and E. Kosmatopoulos, [Store-and-forward based methods for the signal control problem in large-scale congested urban networks](#), *Transportation Research Part C: Emerging Technologies*, **17** (2009), 163–174.
- [2] R. E. Allsop, [Estimating the traffic capacity of a signalized road junction](#), *Transportation Research*, **6** (1972), 245–255.
- [3] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar and P. Whiting, [Scheduling in a queuing system with asynchronously varying service rates](#), *Probability in the Engineering and Informational Sciences*, **18** (2004), 191–217.
- [4] R. Brockett, [Stabilization of motor networks](#), in *Proceedings of the 34th IEEE Conference on Decision and Control*, 1995, 1484–1488.
- [5] J. G. Dai and W. Q. Lin, [Maximum pressure policies in stochastic processing networks](#), *Operations Research*, **53** (2005), 197–218.
- [6] E. J. Davison and U. Ozguner, [Decentralized control of traffic networks](#), *IEEE Transactions on Systems, Man, and Cybernetics*, **13** (1983), 476–487.
- [7] M. Egerstedt and Y. Wardi, [Multi-process control using queuing theory](#), in *Proceedings of the 41st IEEE Conference on Decision and Control*, 2002, 1991–1996.
- [8] P. Giaccone, E. Leonardi and D. Shah, [On the maximal throughput of networks with finite buffers and its application to buffered crossbars](#), in *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, 2005, 971–980.
- [9] P. B. Hunt, D. I. Robertson, R. D. Bretherton and R. I. Winton, *SCOOT - a Traffic Responsive Method of Coordinating Signals*, Transport and Road Research Laboratory, UK, 1981.
- [10] H. Ishii and B. A. Francis, [Stabilizing a linear system by switching control with dwell time](#), *IEEE Trans. Automat. Control*, **47** (2002), 1962–1973.
- [11] M. J. Neely, E. Modiano and C. E. Rohrs, [Dynamic power allocation and routing for time-varying wireless networks](#), *IEEE Journal on Selected Areas in Communications*, **23** (2005), 89–103.
- [12] M. Pajic, S. Sundaram and G. J. Pappas, [Stabilizability over Deterministic Relay Networks](#), in *Proceedings of the 52nd IEEE Conference on Decision and Control*, 2013.
- [13] A. G. Sims and K. W. Dobinson, [The Sydney coordinated adaptive traffic \(SCAT\) system philosophy and benefits](#), *IEEE Transactions on Vehicular Technology*, **29** (1980), 130–137.
- [14] A. L. Stolyar, [MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic](#), *Annals of Applied Probability*, **14** (2004), 1–53.

- [15] L. Tassiulas, [Adaptive back-pressure congestion control-based on local information](#), *IEEE Transactions on Automatic Control*, **40** (1995), 236–250.
- [16] L. Tassiulas and A. Ephremides, [Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks](#), *IEEE Transactions on Automatic Control*, **37** (1992), 1936–1948.
- [17] M. van den Berg, A. Hegyi, B. De Schutter and H. Hellendoorn, A macroscopic traffic flow model for integrated control of freeway and urban traffic networks, in *Proceedings of the 42nd IEEE Conference on Decision and Control*, 2003, 2774–2779.
- [18] P. Varaiya, [Max pressure control of a network of signalized intersections](#), *Transportation Research Part C: Emerging Technologies*, **36** (2013), 177–195.
- [19] T. Wongpiromsarn, T. Uthacharoenpong, Y. Wang, E. Frazzoli and D. Wang, [Distributed traffic signal control for maximum network throughput](#), in *Proceedings of the 15th International IEEE Conference on Intelligent Transportation Systems*, 2012, 588–595.
- [20] H. M. Zhang, Y. M. Nie and Z. S. Qian, Modelling network flow with and without link interactions: The cases of point queue, spatial queue and cell transmission model, *Transportmetrica B: Transport Dynamics*, **1** (2013), 33–51.

Received May 2016; revised February 2018.

E-mail address: leah.anderson@berkeley.edu

E-mail address: tpumir@princeton.edu

E-mail address: dimitris.triantafyllos@gmail.com

E-mail address: bayen@berkeley.edu