Traffic Network Control from Temporal Logic Specifications

Samuel Coogan, Ebru Aydin Gol, Murat Arcak, and Calin Belta

Abstract—We propose a framework for generating a signal control policy for a traffic network of signalized intersections to accomplish control objectives expressible using linear temporal logic. By applying techniques from model checking and formal methods, we obtain a correct-by-construction controller that is guaranteed to satisfy complex specifications. To apply these tools, we identify and exploit structural properties particular to traffic networks that allow for efficient computation of a finite state abstraction. In particular, traffic networks exhibit a *componentwise monotonicity* property which allows reach set computations that scale linearly with the dimension of the continuous state space.

I. INTRODUCTION

State-of-the-art approaches to coordinated control of signalized intersections often focus on limited objectives such as maximizing throughput [1] or maintaining stability of network queues [2], [3]; see [4] for a review of the literature. However, traffic networks are a natural domain for a much richer class of control objectives that are expressible using *linear temporal logic (LTL)* [5], [6]. LTL *formulae* allow control objectives such as "actuate traffic flows such that throughput is always greater than C_1" where C_1 is a threshold throughput, or such that "traffic link queues are always less than C_2" where C_2 is a threshold queue length. LTL formulae also allow more complex objectives such as "infinitely often, the queue length on road ℓ should reach 0," "anytime link ℓ becomes congested, it eventually becomes uncongested," or any combination of these conditions. As these examples suggest, many objectives that are difficult or impossible to address using standard control theoretic techniques are easily expressed in LTL.

In this paper, we propose a technique for synthesizing a signal control policy for a traffic network such that the network satisfies a given control objective expressed using LTL. The synthesized policy is a finite-memory, state feedback controller that is provably correct, that is, guaranteed to result in a closed loop system that satisfies the control objective.

Recent approaches to control synthesis from LTL specifications such as [7]–[21] allow automatic development of correct-by-construction control laws; however, despite these promising developments, scalability concerns prevent direct application of existing results to large traffic networks.

To overcome these scalability limitations, we identify and exploit *componentwise monotonicity* [22] properties inherent in flow networks such as traffic networks. These properties allow efficient computation of bounds on the one-step reachable set from a rectangular box of initial conditions, which in turn allows efficient computation of a finite state abstraction of the dynamics, thereby mitigating a crucial bottleneck in the control synthesis process. A related approach to abstractions of monotone systems is suggested in [23], however the componentwise monotonicity properties exploited in this work are much more general and encompass monotone systems as a special case. The present paper builds on our preliminary work in [24] by defining componentwise monotonicity and identifying it as the enabling property for efficient abstraction.

This paper is organized as follows: Section II gives necessary preliminaries. Section III presents the model for signalized networks, and Section IV establishes the problem formulation. Section V identifies componentwise monotonicity properties of the traffic networks, and Section VI presents scalable algorithms that rely on these properties to construct a finite state representation of the traffic network. Section VII describes the controller synthesis approach, and discusses the computation requirements of our method. We present a case study in Section VIII and conclude our work in Section IX.

II. PRELIMINARIES

The set $\mathcal{I} \subseteq \mathbb{R}^n$ is a *box* if it is the cartesian product of intervals, or equivalently, $\mathcal{I}$ is a box if there exists $x, y \in \mathbb{R}^n$ such that $\mathcal{I} = \prod_{i=1}^n \{z \in \mathbb{R} \mid x_i \prec_i^1 z \prec_i^2 y_i\}$ where $\prec_i^1, \prec_i^2 \in \{<, \leq\}$ and x_i, y_i denote the ith coordinate of x and y, respectively. Defining $\prec^1 \triangleq \{\prec_i^1\}_{i=1}^n$ and $\prec^2 \triangleq \{\prec_i^2\}_{i=1}^n$, we may write $\mathcal{I} = \{z \in \mathbb{R}^n \mid x \prec^1 z \prec^2 y\}$. The vector x is the *lower corner* of $\mathcal{I}$, and likewise y is the *upper corner*.

When applied to vectors, $<$, $\leq$, $>$, and $\geq$ are interpreted elementwise. The notation $\mathbf{0}$ denotes the all-zeros vector where the dimension is clear from context. We denote closure of a set Y by $\mathbf{cl}(Y)$. Given an index set $\mathcal{L}$ and a set of values $x_\ell \in \mathbb{R}$ for $\ell \in \mathcal{L}$, $\{x_\ell\}_{\ell \in \mathcal{L}}$ denotes the collection of x_ℓ, $\ell \in \mathcal{L}$, but we also interpret $x = \{x_\ell\}_{\ell \in \mathcal{L}}$ as an element of $\mathbb{R}^{|\mathcal{L}|}$.

A *transition system* is a tuple $\mathcal{T} = (Q, S, \rightarrow)$ where Q is a finite set of states, S is a finite set of actions, and $\rightarrow \subset Q \times S \times Q$ is a transition relation. We write $q \xrightarrow{s} q'$ instead of $(q, s, q') \in \rightarrow$. Note that all transition systems in this paper are *finite* [6]. The evolution of a transition system is described by $\rightarrow$. That is, a transition system is initialized in some state $q_0 \in Q$, and, given an action $s \in S$, the next state of the transition system is chosen nondeterministically from $\{q' \mid q \xrightarrow{s} q'\}$.
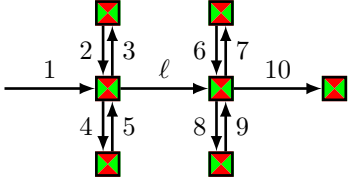
Fig. 1. A typical traffic network with 11 links and 7 signalized intersections. In the figure, $\mathcal{L}_\ell^{\text{down}} = \{\ell, 7, 8, 10\}$, $\mathcal{L}_\ell^{\text{up}} = \{1, 2, 5\}$, and $\mathcal{L}_\ell^{\text{adj}} = \{3, 4\}$. At each time step, a signal actuates a subset of upstream links.

III. SIGNALIZED NETWORK TRAFFIC MODEL

A signalized traffic network consists of a set $\mathcal{L}$ of *links* and a set $\mathcal{V}$ of *signalized intersections*. For $\ell \in \mathcal{L}$, let $\eta(\ell) \in \mathcal{V}$ denote the downstream intersection of link ℓ and let $\tau(\ell) \in \mathcal{V} \cup \emptyset$ denote the upstream intersection of link ℓ. A link ℓ with $\tau(\ell) = \emptyset$ serves as an entry-point into the network, and we assume $\eta(\ell) \neq \tau(\ell)$ for all $\ell \in \mathcal{L}$ (*i.e.*, no self-loops). Link $k \neq \ell$ is *upstream* of link ℓ if $\eta(k) = \tau(\ell)$, *downstream* of link ℓ if $\tau(k) = \eta(\ell)$, and *adjacent to* link ℓ if $\tau(k) = \tau(\ell)$. Roads exiting the traffic network are not modeled explicitly. For each $v \in \mathcal{V}$, define $\mathcal{L}_v^{\text{in}} = \{\ell \mid \eta(\ell) = v\}$, $\mathcal{L}_v^{\text{out}} = \{\ell \mid \tau(\ell) = v\}$ and for each $\ell \in \mathcal{L}$, define

$$\mathcal{L}_\ell^{\text{up}} = \{k \in \mathcal{L} \mid \eta(k) = \tau(\ell)\} \tag{1}$$
$$\mathcal{L}_\ell^{\text{down}} = \{k \in \mathcal{L} \mid \tau(k) = \eta(\ell)\} \cup \{\ell\} \tag{2}$$
$$\mathcal{L}_\ell^{\text{adj}} = \{k \in \mathcal{L} \mid \tau(k) = \tau(\ell)\} \setminus \{\ell\} \tag{3}$$

so that $\mathcal{L}_\ell^{\text{down}}$ includes link ℓ and the links downstream of link ℓ, and $\mathcal{L}_\ell^{\text{up}}$ and $\mathcal{L}_\ell^{\text{adj}}$ are the links upstream and adjacent to ℓ, respectively, see Fig. 1. We have $\mathcal{L}_\ell^{\text{down}} \cap \mathcal{L}_\ell^{\text{adj}} = \emptyset$ and $\mathcal{L}_\ell^{\text{up}} \cap \mathcal{L}_\ell^{\text{adj}} = \emptyset$, but note that it is possible for $\mathcal{L}_\ell^{\text{down}} \cap \mathcal{L}_\ell^{\text{up}} \neq \emptyset$, in particular, if there is a cycle of length two in the network. Let $\mathcal{L}_\ell^{\text{loc}} = \mathcal{L}_\ell^{\text{down}} \cup \mathcal{L}_\ell^{\text{up}} \cup \mathcal{L}_\ell^{\text{adj}}$ be links "local" to link ℓ.

Each link $\ell \in \mathcal{L}$ possesses a queue $x_\ell[t] \in [0, x_\ell^{\text{cap}}]$ representing the number of vehicles on link ℓ at time step $t \in \mathbb{N} \triangleq \{0, 1, 2, \dots\}$ where x_ℓ^{cap} is the capacity of link ℓ. We allow x_ℓ to be a continuous quantity, thus adopting a fluid-like model of traffic flow evolving in slotted time as in [1]–[3].

Movement of vehicles among link queues is governed by mass-conservation laws and the state of the signalized intersections. A link is said to be *actuated* if outgoing flow from link ℓ is allowed as determined by the state of the traffic signal at intersection $\eta(\ell)$. At each intersection v,

$$\mathcal{S}_v \subseteq 2^{\mathcal{L}_v^{\text{in}}} \tag{4}$$

denotes the set of available signal *phases*, that is, each $s_v \in \mathcal{S}_v$, $s_v \subseteq \mathcal{L}_v^{\text{in}}$ denotes a set of incoming links at intersection v that may be actuated simultaneously. We define

$$\mathcal{S} = \{\cup_{v \in \mathcal{V}} s_v \mid s_v \in \mathcal{S}_v \ \forall v \in \mathcal{V}\} \subseteq 2^{\mathcal{L}} \tag{5}$$

so that each $s \in \mathcal{S}$, $s \subseteq \mathcal{L}$ denotes a set of links in the network that may be actuated simultaneously. We identify $s \in \mathcal{S}$ with its constituent phases so that $s = \{s_v\}_{v \in \mathcal{V}}$, and we interpret $\mathcal{S}$ as the set of allowed inputs to the traffic network.

When a link is actuated, a maximum of c_ℓ vehicles are allowed to flow from link ℓ to links $\mathcal{L}_{\eta(\ell)}^{\text{out}}$ per time step where

c_ℓ is the known *saturation flow* for link ℓ, [4]. The *turn ratio* $\beta_{\ell k}$ denotes the fraction of vehicles exiting link ℓ that are routed to link k, [2]. Then $\beta_{\ell k} \neq 0$ only if $\eta(\ell) = \tau(k)$, and

$$\sum_{k \in \mathcal{L}_{\eta(\ell)}^{\text{out}}} \beta_{\ell k} \leq 1. \tag{6}$$

Strict inequality in (6) implies that a fraction of vehicles on link ℓ are routed off the network via unmodeled roads that exit the network. Traffic flow can occur only if there is available capacity downstream. To this end, the supply ratio $\alpha_{\ell k}^{s_v}$ denotes the fraction of link k's capacity available to link ℓ during phase $s_v \in \mathcal{S}_{\tau(k)}$. That is, link ℓ may only send $\alpha_{\ell k}^{s_v}(x_k^{\text{cap}} - x_k[t])$ vehicles to link k in time period t under input s_v. As the supply is only divided among actuated incoming links, it follows that for each $k \in \mathcal{L}$

$$\sum_{\ell \in s_v} \alpha_{\ell k}^{s_v} = 1 \quad \forall s_v \in \mathcal{S}_{\tau(k)}, s_v \neq \emptyset. \tag{7}$$

Constant turn and supply ratios are a common modeling assumption justified by empirical observations; see [25] for further discussion.

We are now in a position to define the dynamics of the link queues. As we will see subsequently, the flow of vehicles out of link ℓ is only a function of the state of links in $\mathcal{L}_\ell^{\text{down}}$, and the update of link ℓ's state is only a function of links in $\mathcal{L}_\ell^{\text{loc}}$.

Let $\boldsymbol{x}[t] = \{x_\ell[t]\}_{\ell \in \mathcal{L}}$, $\boldsymbol{x}_\ell^{\text{down}}[t] = \{x_k[t]\}_{k \in \mathcal{L}_\ell^{\text{down}}}$, and $\boldsymbol{x}_\ell^{\text{loc}}[t] = \{x_k[t]\}_{k \in \mathcal{L}_\ell^{\text{loc}}}$. The outflow of link $\ell \in \mathcal{L}$ is as follows:

$$f_\ell^{\text{out}}(\boldsymbol{x}_\ell^{\text{down}}, s_{\eta(\ell)}) =$$
$$\begin{cases} \min\left\{x_\ell[t], c_\ell, \min_{\substack{k \text{ s.t.} \\ \beta_{\ell k} \neq 0}} \left\{\frac{\alpha_{\ell k}^{s_{\eta(\ell)}}}{\beta_{\ell k}}(x_k^{\text{cap}} - x_k[t])\right\}\right\} \\ \qquad\qquad\qquad\qquad\qquad \text{if } \ell \in s_{\eta(\ell)} \\ 0 \qquad\qquad\qquad\qquad\qquad\quad \text{else.} \end{cases} \tag{8}$$

The interpretation of (8) is that the flow of vehicles exiting a link ℓ when actuated is the minimum of the link's queue length, its saturation flow, and the downstream *supply* of capacity, weighted appropriately by turn and supply ratios. This modeling approach is based on the *cell transmission model* of traffic flow [26] which restricts flow if there is inadequate capacity downstream. A consequence of (8) is that inadequate capacity on one downstream link at an intersection causes congestion that blocks incoming flow to other downstream links. This phenomenon, sometimes called the *first-in-first-out* property, has been widely studied in the transportation literature and occurs even in multilane settings [27][1]. The number of vehicles in each link's queue then evolves according to the mass conservation equation

$$x_\ell[t+1] = F_\ell(\boldsymbol{x}_\ell^{\text{loc}}[t], \boldsymbol{s}_\ell^{\text{loc}}[t], d_\ell[t]) \tag{9}$$
$$\triangleq \min\left\{x_\ell^{\text{cap}}, x_\ell[t] - f_\ell^{\text{out}}(\boldsymbol{x}_\ell^{\text{down}}[t], s_{\eta(\ell)})\right.$$
$$\left. + \sum_{j \in \mathcal{L}_\ell^{\text{up}}} \beta_{j\ell} f_j^{\text{out}}(\boldsymbol{x}_j^{\text{down}}[t], s_{\eta(j)}) + d_\ell[t]\right\} \tag{10}$$

[1]Even if a turn pocket exists at an intersection, it is often too short to fully mitigate this blocking property. Nonetheless, if the road geometry is such that a sufficient number of dedicated lanes exist for a turning movement, these lanes may be modeled with a separate link.

where $d_\ell[t]$ is the number of vehicles that exogenously enters the queue on link ℓ in time step t, $\boldsymbol{d} = \{d_\ell[t]\}_{\ell \in \mathcal{L}}$, and $\boldsymbol{s}_\ell^{\text{loc}} = \{s_{\eta(\ell)}, s_{\tau(\ell)}\}$ if $\tau(\ell) \neq \emptyset$, $\boldsymbol{s}_\ell^{\text{loc}} = \{s_{\eta(\ell)}\}$ otherwise; that is, $\boldsymbol{s}_\ell^{\text{loc}}$ is the state of the signals that are "local" to link ℓ. The minimization in (10) is only needed in case the exogenous input $d_\ell[t]$ would cause the state of link ℓ to exceed x_ℓ^{cap} and ensures that the network dynamics maps

$$\mathcal{X} = \prod_{\ell \in \mathcal{L}}[0, x_\ell^{\text{cap}}] \tag{11}$$

to itself. We interpret this as refusal of vehicles attempting to exogenously enter the network when the link is full. Note in particular that the supply/demand formulation prevents upstream inflow from exceeding supply and thus for links with no exogenous input, x_ℓ^{cap} is never the unique minimizer in (10).

Remark 1. An alternative to the above approach is to define an auxiliary sink state Out in the transition systems $\mathcal{T}$ defined in Section VI which captures any trajectories that exit the domain $\mathcal{X}$. The temporal logic specification can then incorporate the requirement that the system never enters this Out state.

Assumption 1. *We assume there exists $\mathcal{D} \subset \mathbb{R}^\mathcal{L}$ such that*

$$\boldsymbol{d}[t] \in \mathcal{D} \quad \forall t \tag{12}$$

and $\mathcal{D}$ satisfies $\mathcal{D} \subset \cup_{i=1}^{n_\mathcal{D}} \mathcal{D}_i$ where each $\mathcal{D}_i$ is given by

$$\mathcal{D}_i = \{\boldsymbol{d} \mid \underline{\boldsymbol{d}}^i \leq \boldsymbol{d} \leq \overline{\boldsymbol{d}}^i\} \tag{13}$$

for some $\underline{\boldsymbol{d}}^i = \{\underline{d}_\ell^i\}_{\ell \in \mathcal{L}}$, $\overline{\boldsymbol{d}}^i = \{\overline{d}_\ell^i\}_{\ell \in \mathcal{L}}$.

In other words, we assume the disturbance is contained within a union of boxes given by (13). This assumption is not particularly restrictive, as any compact subset of $\mathbb{R}^\mathcal{L}$ can be approximated with boxes to arbitrary precision [28], however the number of boxes $n_\mathcal{D}$ affects the computation time as detailed in Section VII-B.

We let $F(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d}) = \{F_\ell(\boldsymbol{x}_\ell^{\text{loc}}, \boldsymbol{s}_\ell^{\text{loc}}, d_\ell)\}_{\ell \in \mathcal{L}} : \mathcal{X} \times \mathcal{S} \times \mathcal{D} \to \mathcal{X}$ so that

$$\boldsymbol{x}[t+1] = F(\boldsymbol{x}[t], \boldsymbol{s}[t], \boldsymbol{d}[t]). \tag{14}$$

The set of states of system (14) that are reachable from a set $Y \subset \mathcal{X}$ under the control signal $\boldsymbol{s} \in \mathcal{S}$ in one timestep is denoted by the Post operator and given by

$$\text{Post}(Y, \boldsymbol{s}) = \{\boldsymbol{x}' = F(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d}) \mid \boldsymbol{x} \in Y, \boldsymbol{d} \in \mathcal{D}\}. \tag{15}$$

We call $\text{Post}(Y, \boldsymbol{s})$ the *one step reachable set from Y under $\boldsymbol{s}$*. The main features of the queue-based modeling approach proposed above such as finite saturation rates, finite queue capacity, a set of available signaling phases, and fixed turn ratios are standard in many modeling and simulation approaches such as [2], [3], see also [4], [29] and references therein for discussions of queue-based modeling of traffic networks.

IV. PROBLEM FORMULATION AND APPROACH

We now define and motivate the need for control objectives expressible in LTL for traffic networks, and we outline a control synthesis approach which relies on a finite state representation of the traffic dynamics to meet these objectives.

LTL *formulae* are generated inductively using the Boolean operators $\vee$ (disjunction), $\wedge$ (conjunction), $\neg$ (negation), and the temporal operators $\bigcirc$ (next) and U (until). From these, we obtain a suite of derived logical and temporal operators such as $\to$ (implication), $\square$ (always), $\Diamond$ (eventually), $\square\Diamond$ (infinitely often), finite deadlines with repeated $\bigcirc$, and many others, see [5], [6].

Formally, such formulae are expressed over a set of atomic propositions, which we restrict to be indicator expressions over subsets of $\mathcal{X}$ or predicates over the signaling state. For example, the atomic proposition $x_\ell \leq 10$ is true for all $\boldsymbol{x} \in \mathcal{X}$ that satisfies the condition $x_\ell \leq 10$ (which constitutes a box subset of $\mathcal{X}$), and the atomic proposition $\ell \in \boldsymbol{s}$ is true for all signals that actuate link ℓ. We will see in Section V and Section VI that restricting to atomic propositions corresponding to box subsets of $\mathcal{X}$ offers significant computational advantages.

Semantically, LTL formulae are interpreted over a trajectory $\boldsymbol{x}[t]$ and the corresponding input sequence $\boldsymbol{s}[t]$ for $t = 0, 1, \ldots$. For example, the state/input sequence $(\boldsymbol{x}[t], \boldsymbol{s}[t])$ *satisfies* the LTL formula $\varphi = \square(x_\ell \leq 10) \wedge \square\Diamond(\ell \in \boldsymbol{s})$ if and only if $x_\ell[t] \leq 10$ for all t and $\ell \in \boldsymbol{s}[t]$ infinitely often (*i.e.*, for infinitely many t). Thus a trajectory satisfies a LTL formula if and only if the formula holds for the corresponding trace of atomic propositions that are valid at each time step. A formal definition of the semantics of LTL over traces is readily available in the literature, *e.g.*, [5], [6], and is a natural interpretation of the above Boolean and temporal operators. For example, a trace satisfies $\Diamond\varphi$ if and only if there exists a suffix of the trace satisfying φ.

Examples of LTL formulae representing desired control objectives relevant to traffic networks include those from the Introduction, as well as:

- $\varphi_1 = \Diamond\square(x_\ell \leq C)$ for some C
  "Eventually, link ℓ will have less than C vehicles and this will remain true for all time"
- $\varphi_2 = \square\Diamond(\ell \in \boldsymbol{s})$
  "Infinitely often, link ℓ is actuated"
- $\varphi_3 = \square((\ell \in s_{v_1}) \to \bigcirc(k \in s_{v_2}))$
  "Whenever signal v_1 actuates link ℓ, signal v_2 must actuate link k in the next time step"
- $\varphi_4 = \square(x_\ell \geq C_1 \to \Diamond(x_\ell \leq C_2))$
  "Whenever the number of vehicles on link ℓ exceeds C_1, it is eventually the case that the number of vehicles on link ℓ decreases below C_2."

The main problem considered in this paper is as follows:

Control Synthesis Problem. *Given a traffic network and an LTL formula φ over a set of atomic propositions as described above, find a control strategy that, at each time step, chooses a signaling input such that all trajectories of the traffic network satisfy φ from any initial condition.*

To solve the control synthesis problem, we propose computing a finite state abstraction that simulates (in a manner to be formalized below) the traffic network dynamics. As we discuss in Section VII-A, the result is a full-state feedback controller which requires finite memory. We rely on dynamical properties of the traffic network to compute the abstraction, and then apply tools from automata theory and formal methods

to synthesize a finite-memory, state feedback control strategy solving the control synthesis problem.

V. COMPONENTWISE MONOTONICITY OF TRAFFIC NETWORKS

To generate control strategies for the traffic network that guarantee satisfaction of a LTL formula, we first construct a finite state representation, or *abstraction*, of the model defined in Section III. We now define a *componentwise monotonicity* property that simplifies this task and next show that the model in Section III possesses this property.

Definition 1. *Consider the dynamical system*

$$z[t+1] = f(z[t], w[t]) \qquad (16)$$

for $z \in \mathcal{Z} \subseteq \mathbb{R}^n$, $w \in \mathcal{W} \subseteq \mathbb{R}^m$ with $f : \mathcal{Z} \times \mathcal{W} \to \mathcal{Z}$ continuous. System (16) *is* componentwise monotone *if there exists a signature matrix $\Delta = [\delta_{ij}]_{i,j=1}^n$ with each $\delta_{ij} \in \{-1, 1\}$ such that for all i,*

$$\delta_{ij}\underline{\xi}_j \leq \delta_{ij}\overline{\xi}_j \text{ and } \underline{w}_j \leq \overline{w}_j \ \forall j \in \mathcal{L} \qquad (17)$$

$$\text{implies} \qquad f(\underline{\xi}, \underline{w}) \leq f(\overline{\xi}, \overline{w})) \qquad (18)$$

for any $\underline{\xi}, \overline{\xi} \in \mathcal{Z}$, $\underline{w}, \overline{w} \in \mathcal{W}$. That is, (16) *is componentwise monotone if f_i is monotonic in each z variable and monotonically increasing in each w variable.*

A definition similar to Definition 1 appears in [22], but omits dependence on a disturbance input.

We now give a characterization of componentwise monotone systems which stipulates that $\partial f / \partial z$ be *sign-stable*, that is, the sign structure of the Jacobian does not change as z, w range over their domain.

Lemma 1. *Consider dynamical system* (16) *and further suppose that $f(z, w) = \begin{bmatrix} f_1(z,w) & \dots f_n(z,w) \end{bmatrix}^T$ is Lipschitz continuous so that partial derivatives exist almost everywhere. If for all $i \in \{1, \dots, n\}$:*

$$\forall j \in \{1, \dots, n\} \ \exists \delta_{ij} \in \{-1, 1\} : \delta_{ij}\frac{\partial f_i}{\partial z_j}(z, w) \geq 0 \text{ a.e.} \quad (19)$$

$$\text{and} \qquad \forall j \in \{1, \dots, m\} : \frac{\partial f_i}{\partial w_j}(z, w) \geq 0 \text{ a.e.} \quad (20)$$

where a.e.(almost everywhere) implies the condition must hold wherever the derivative exists, then (16) *is componentwise monotone.*

Proof. Let $[\delta_{ij}]_{i,j=1}^n$ be as in the hypothesis of the Lemma. By the Fundamental Theorem of Calculus, for all $\underline{w}$ and for almost all[2] $\underline{\xi}, \overline{\xi}$ satisfying $\delta_{ij}\underline{\xi}_j \leq \delta_{ij}\overline{\xi}_j$ for all j,

$$f_i(\overline{\xi}, \underline{w}) - f_i(\underline{\xi}, \underline{w}) = \qquad (21)$$

$$\left(\int_0^1 \sum_{j=1}^n \frac{\partial f_i}{\partial \xi_j}(\underline{\xi} + r(\overline{\xi} - \underline{\xi}), \underline{w})(\overline{\xi}_j - \underline{\xi}_j)dr \right) \geq 0 \quad (22)$$

where nonnegativity follows because $\delta_{ij}(\overline{\xi}_j - \underline{\xi}_j) \geq 0$, $\delta_{ij}\partial f_i/\partial \xi_j \geq 0$, and $\delta_{ij}^2 = 1$ for all i, j. Similarly, for almost

---

[2]Eq. (21) requires existence of $\partial g/\partial z$ almost everywhere along the line segment connecting $\underline{z}$ and $\overline{z}$, which holds for almost all $\overline{z}$ for fixed $\underline{z}$ [30, Ch. 2]. Similarly, $f_i(\overline{\xi}, \overline{w}) - f_i(\overline{\xi}, \underline{w})$ holds for almost all $\overline{w}$ for fixed $\underline{w}$.

all $\overline{w} \geq \underline{w}$, $f_i(\overline{\xi}, \overline{w}) - f_i(\overline{\xi}, \underline{w}) \geq 0$. It follows by continuity of f_i that $f_i(\overline{\xi}, \overline{w}) - f_i(\underline{\xi}, \underline{w}) \geq 0$ for all $\underline{\xi}, \overline{\xi}$ that satisfy $\delta_{ij}\underline{\xi}_j \leq \delta_{ij}\overline{\xi}_j$ and all $\underline{w} \leq \overline{w}$, for all i, completing the proof. $\square$

The critical feature of componentwise monotone systems we wish to exploit is that over approximating the one-step reachable set from a box of initial conditions is computationally efficient. In particular, the reach set is contained within a box defined by the value of f_i at two particular points for each i, regardless of the dimension of the spaces $\mathcal{Z}$ and $\mathcal{W}$:

Lemma 2. *Let* (16) *be componentwise monotone with signature matrix $\Delta = [\delta_{ij}]_{i,j=1}^n$ and assume $\mathcal{Z}$ is a closed box. Given $\underline{z}, \overline{z} \in \mathcal{Z}$ and $\underline{w}, \overline{w} \in \mathcal{W}$ with $\underline{z} \leq \overline{z}$ and $\underline{w} \leq \overline{w}$. Let $\underline{\xi}^i \in \mathcal{Z}$ and $\overline{\xi}^i \in \mathcal{Z}$ be defined elementwise as follows for each i:*

$$\underline{\xi}_j^i = \begin{cases} \underline{z}_j & \text{if } \delta_{ij} = 1 \\ \overline{z}_j & \text{if } \delta_{ij} = -1 \end{cases}, \quad \overline{\xi}_j^i = \begin{cases} \overline{z}_j & \text{if } \delta_{ij} = 1 \\ \underline{z}_j & \text{if } \delta_{ij} = -1. \end{cases} \quad (23)$$

Then

$$f_i(\underline{\xi}^i, \underline{w}) \leq f_i(z, w) \leq f_i(\overline{\xi}^i, \overline{w}) \qquad \forall i \quad (24)$$

for all z, w such that $\underline{z} \leq z \leq \overline{z}$ and $\underline{w} \leq w \leq \overline{w}$.

Proof. Observe that $\delta_{ij}\underline{\xi}_j^i \leq \delta_{ij}z_j$ for all i, j for all $\underline{z} \leq z \leq \overline{z}$, and symmetrically, $\delta_{ij}z_j \leq \delta_{ij}\overline{\xi}_j^i$ for all i, j for all $\underline{z} \leq z \leq \overline{z}$. The Lemma then follows immediately from Definition 1. $\square$

This remarkable feature of componentwise monotone systems is analogous to well-known results for monotone systems [31]–[33], but componentwise monotonicity allows consideration of a much broader class of systems, including traffic networks, which are generally not monotone.

Remark 2. Observe that the lower and upper bounds in (24) are achieved for appropriate choice of z and w, thus the approximation of the one-step reachable set is tight.

To prove that the traffic network dynamics developed in Section III are componentwise monotone, we first require a technical assumption:

Assumption 2. *For all $\ell \in \mathcal{L}$,*

$$c_\ell \leq x_\ell^{cap} - \frac{\beta_{k\ell}}{\alpha_{k\ell}}c_k \quad \forall k \in \mathcal{L}_\ell^{up}. \quad (25)$$

Assumption 2 is a sufficient condition for ensuring that if a link has inadequate capacity and blocks upstream flow, then this link's queue will not empty in one time step. This effectively is an assumption that the time step is sufficiently small to appropriately capture the queuing phenomenon. Specifically, the saturation flow rate c_ℓ of link ℓ is in units of vehicles per time step and, thus, is implicitly a function of the chosen time step. Physically, c_ℓ is required to decrease with decreased time step and thus Assumption 2 is satisfied when a sufficiently small time step is used for the model.

Theorem 1. *The traffic network model is componentwise monotone for any signaling input $s \in \mathcal{S}$. In particular, F_ℓ is increasing in x_k for k downstream or upstream of link ℓ or equal to ℓ, and decreasing in x_k for k adjacent to link ℓ.*

Proof. For fixed $\boldsymbol{s} \in \mathcal{S}$, we show that $F(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d})$ satisfies conditions (19) and (20) of Lemma 1 with $\boldsymbol{x}, \boldsymbol{d}$ replacing z, w. Observe that F is continuous and piecewise differentiable by (8)–(10) and (14), thus it is Lipschitz continuous [34]. The minimum function in (8) implies that F is differentiable almost everywhere. We first have $\frac{\partial F_\ell}{\partial d_\ell} \in \{0, 1\}$ a.e. by (10), satisfying (20). Now consider $\partial F_\ell / \partial x_k$. For (20), we consider four exhaustive cases:

- Case 1, $k \in (\mathcal{L}_\ell^{\text{down}} \cup \mathcal{L}_\ell^{\text{up}}) \backslash \{\ell\}$. From (8)–(10), link k may block the outflow of link ℓ when $k \in \mathcal{L}_\ell^{\text{down}}$, or link k may contribute to the inflow to link ℓ if $k \in \mathcal{L}_\ell^{\text{up}}$, thus we have $\frac{\partial F_\ell}{\partial x_k} \in \{0, -\frac{\partial f_k^{\text{out}}}{\partial x_k}, \beta_{k\ell} \frac{\partial f_k^{\text{out}}}{\partial x_k}, -\frac{\partial f_\ell^{\text{out}}}{\partial x_k} + \beta_{k\ell} \frac{\partial f_k^{\text{out}}}{\partial x_k}\}$ a.e. where the fourth possibility occurs only if $k \in \mathcal{L}_\ell^{\text{down}} \cap \mathcal{L}_\ell^{\text{up}}$. But $\frac{\partial f_\ell^{\text{out}}}{\partial x_k} \in \{0, -\alpha_{\ell k}^{s_{\eta(\ell)}}/\beta_{\ell k}\}$ a.e. and $\frac{\partial f_k^{\text{out}}}{\partial x_k} \in \{0, 1\}$ a.e., thus $\frac{\partial F_\ell}{\partial x_k} \geq 0$ a.e., satisfying (19).

- Case 2, $k = \ell$. We have $\frac{\partial f_\ell^{\text{out}}}{\partial x_\ell} \in \{0, 1\}$ a.e. and, for $j \in \mathcal{L}_\ell^{\text{up}}$, $\frac{\partial f_j^{\text{out}}}{\partial x_\ell} \in \{0, \alpha_{j\ell}^{s_{\eta(j)}}/\beta_{j\ell}\}$ a.e., however, Assumption 2 ensures that, a.e., either $\frac{\partial f_\ell^{\text{out}}}{\partial x_\ell} = 0$ or $\frac{\partial f_\ell^{\text{in}}}{\partial x_\ell} = 0$, *i.e.*, $\frac{\partial f_\ell^{\text{out}}}{\partial x_\ell} = 0$ for all $j \in \mathcal{L}_\ell^{\text{up}}$. Thus $\frac{\partial F_\ell}{\partial x_\ell} \in \{0, 1, 1 + \sum_{j \in \mathcal{L}_\ell^{\text{up}}} \beta_{j\ell} \frac{\partial f_j^{\text{out}}}{\partial x_\ell}\}$ a.e. But $\sum_{j \in \mathcal{L}_\ell^{\text{up}}} \beta_{j\ell} \frac{\partial f_j^{\text{out}}}{\partial x_\ell} \geq -\sum_{j \in \mathcal{L}_\ell^{\text{up}}} \alpha_{j\ell}^{s_{\eta(j)}} = -1$ by (7) (recall that $\eta(j) = \tau(\ell)$ for all $j \in \mathcal{L}_\ell^{\text{up}}$), that is, $\partial f_\ell^{\text{in}} / \partial x_\ell \geq -1$, thus $\frac{\partial F_\ell}{\partial x_\ell} \geq 0$ a.e., satisfying (19).

- Case 3, $k \in \mathcal{L}_\ell^{\text{adj}}$. In this case, inadequate capacity of link k may block flow to link ℓ, as discussed above. We have $\frac{\partial F_\ell}{\partial x_k} = \sum_{j \in \mathcal{L}_\ell^{\text{up}}} \beta_{j\ell} \frac{\partial f_j^{\text{out}}}{\partial x_k}$. Since $\frac{\partial f_j^{\text{out}}}{\partial x_k} \in \{0, -\alpha_{jk}^{s_{\eta(j)}}/\beta_{jk}\}$ a.e., we have $\frac{\partial F_\ell}{\partial x_k} \leq 0$ a.e., satisfying (19).

- Case 4, $k \notin \mathcal{L}_\ell^{\text{loc}}$. Then $\frac{\partial F_\ell}{\partial x_k} = 0$, trivially satisfying (19). □

The following corollary implies that the one-step reachable set of the traffic dynamics from a (closed) box $\mathcal{I}$ for any given signaling input $\boldsymbol{s}$ is over-approximated by the union of boxes, one box for each $i = 1, \ldots n_{\mathcal{D}}$, where each of these boxes is efficiently computed by evaluating F_ℓ at two particular points for each $\ell \in \mathcal{L}$. The obtained over-approximation is denoted with the $\overline{\texttt{Post}}$ operator. This critical result allows efficient computation of a finite state representation of the traffic dynamics, as detailed in Section VI.

Corollary 1. *Consider the set $\mathcal{I} = \{\boldsymbol{x} \mid \underline{\boldsymbol{x}} \leq \boldsymbol{x} \leq \bar{\boldsymbol{x}}\}$ for $\underline{\boldsymbol{x}}, \bar{\boldsymbol{x}} \in \mathcal{X}$, and for each $\ell \in \mathcal{L}$, define $\underline{\boldsymbol{\xi}}^\ell(\underline{\boldsymbol{x}}, \bar{\boldsymbol{x}}) = \{\underline{\xi}_k^\ell(\underline{x}_k, \bar{x}_k)\}_{k \in \mathcal{L}_\ell^{\text{loc}}}$, $\bar{\boldsymbol{\xi}}^\ell(\underline{\boldsymbol{x}}, \bar{\boldsymbol{x}}) = \{\bar{\xi}_k^\ell(\underline{x}_k, \bar{x}_k)\}_{k \in \mathcal{L}_\ell^{\text{loc}}}$ where*

$$\underline{\xi}_k^\ell(\underline{x}_k, \bar{x}_k) = \begin{cases} \underline{x}_k & \text{if } k \in \mathcal{L}_\ell^{\text{down}} \cup \mathcal{L}_\ell^{\text{up}} \\ \bar{x}_k & \text{if } k \in \mathcal{L}_\ell^{\text{adj}} \end{cases} \quad (26)$$

$$\bar{\xi}_k^\ell(\underline{x}_k, \bar{x}_k) = \begin{cases} \bar{x}_k & \text{if } k \in \mathcal{L}_\ell^{\text{down}} \cup \mathcal{L}_\ell^{\text{up}} \\ \underline{x}_k & \text{if } k \in \mathcal{L}_\ell^{\text{adj}}. \end{cases} \quad (27)$$

Then for all $\boldsymbol{s} \in \mathcal{S}$, $\texttt{Post}(\mathcal{I}, \boldsymbol{s}) \subseteq \overline{\texttt{Post}}(\mathcal{I}, \boldsymbol{s})$ where

$$\overline{\texttt{Post}}(\mathcal{I}, \boldsymbol{s}) := \\ \bigcup_{i=1}^{n_{\mathcal{D}}} \{\boldsymbol{x}' \mid F_\ell(\underline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \underline{d}_\ell^i) \leq x_\ell' \leq F_\ell(\bar{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \bar{d}_\ell^i) \quad \forall \ell \in \mathcal{L}\}. \quad (28)$$
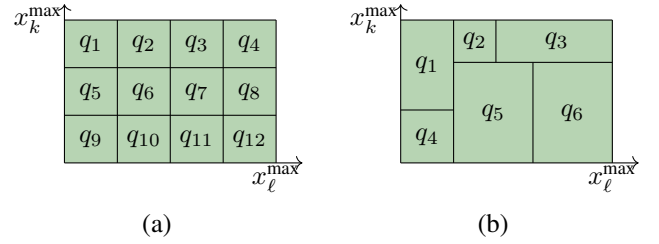

Fig. 2. Stylized depictions of two box partitions. (a) A gridded box partition with regularly sized intervals. (b) A nongridded box partition.

Proof. By substituting $\underline{\boldsymbol{x}}, \bar{\boldsymbol{x}}$ for $z, \bar{z}$ and $\underline{d}_\ell^i, \bar{d}_\ell^i$ for $\underline{w}, \overline{w}$ in Lemma 2 and defining $f(\boldsymbol{x}, \boldsymbol{d}) \triangleq F(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d})$, we obtain $\{\boldsymbol{x}' = F(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d}) \mid \boldsymbol{x} \in \mathcal{I}, \boldsymbol{d} \in \mathcal{D}_i\} \subseteq \{\boldsymbol{x}' \mid F_\ell(\underline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \underline{d}_\ell^i) \leq x_\ell' \leq F_\ell(\bar{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \bar{d}_\ell^i) \ \forall \ell \in \mathcal{L}\}$ for all $i = 1, \ldots, n_{\mathcal{D}}$. The corollary follows from the trivial fact that $\texttt{Post}(\mathcal{I}, \boldsymbol{s}) = \cup_{i=1}^{n_{\mathcal{D}}} \{\boldsymbol{x}' = F(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d}) \mid \boldsymbol{x} \in \mathcal{I}, d \in \mathcal{D}_i\}$. □

VI. FINITE STATE REPRESENTATION

To apply the powerful tools of LTL synthesis, we require a finite state representation of the traffic network model. In general, obtaining finite state abstractions is a difficult problem and existing techniques do not scale well. In this section, we exploit the componentwise monotonicity properties developed above and propose an efficient method for determining a finite state representation of the traffic network dynamics.

A. Finite State Abstraction

Definition 2 (Box partition)**.** *For finite index set $\mathcal{Q}$, the set $\{\mathcal{I}_q\}_{q \in \mathcal{Q}}$ is a* box partition *of $\mathcal{X}$ (or simply a* box partition*), if each $\mathcal{I}_q \subseteq \mathcal{X}$ is a box, $\cup_{q \in \mathcal{Q}} \mathcal{I}_q = \mathcal{X}$, and $\mathcal{I}_q \cap \mathcal{I}_{q'} = \emptyset$ for all $q, q' \in \mathcal{Q}$. For $q \in \mathcal{Q}$, let $\underline{\boldsymbol{x}}_q = \{\underline{x}_{q,\ell}\}_{\ell \in \mathcal{L}}$, $\bar{\boldsymbol{x}}_q = \{\bar{x}_{q,\ell}\}_{\ell \in \mathcal{L}}$ denote the lower and upper corners, respectively, of $\mathcal{I}_q$, that is, $\mathcal{I}_q = \{\boldsymbol{x} \mid \underline{\boldsymbol{x}}_q \prec_q^1 \boldsymbol{x} \prec_q^2 \bar{\boldsymbol{x}}_q\}$ where $\prec_q^1 = \{\prec_{q,\ell}^1\}_{\ell \in \mathcal{L}}$, $\prec_q^2 = \{\prec_{q,\ell}^2\}_{\ell \in \mathcal{L}}$, and $\prec_{q,\ell}^1, \prec_{q,\ell}^2 \in \{<, \leq\}$.*

For a box partition $\{\mathcal{I}_q\}_{q \in \mathcal{Q}}$ of $\mathcal{X}$, let $\pi : \mathcal{X} \to \mathcal{Q}$ be uniquely defined by the condition $x \in \mathcal{I}_{\pi(x)}$, that is, $\pi(\cdot)$ is the natural projection from the domain $\mathcal{X}$ to the (index set of) boxes. A special case of a box partition of a rectangular domain is the following:

Definition 3 (Gridded box partition)**.** *For $\mathcal{X} = \{\boldsymbol{x} = \{x_\ell\}_{\ell \in \mathcal{L}} \mid \underline{x}_\ell \leq x_\ell \leq \bar{x}_\ell\}$, a box partition $\{\mathcal{I}_q\}_{q \in \mathcal{Q}}$ of $\mathcal{X}$ is a* gridded *box partition if for each $\ell \in \mathcal{L}$, there exists $N_\ell \in \{1, 2, \ldots\}$ and a set of intervals $\{I_1^\ell, \ldots, I_{N_\ell}^\ell\}$ such that $\cup_{i=1}^{N_\ell} I_i^\ell = [\underline{x}_\ell, \bar{x}_\ell]$ and for each $q \in \mathcal{Q}$, there exists indices $q_\ell \in \{1, \ldots, N_\ell\}$ such that $\mathcal{I}_q = \prod_{\ell \in \mathcal{L}} I_{q_\ell}^\ell$. For gridded box partitions, we make the identification $\mathcal{Q} \cong \prod_{\ell \in \mathcal{L}} \{1, \ldots, N_\ell\}$ for all $\ell \in \mathcal{L}$.*

When a box partition is not a gridded box partition, we say it is *nongridded*. Fig. 2 shows two examples of box partitions, one of which is a gridded box partition. From a box partition of the traffic network domain $\mathcal{X}$, we obtain a finite state representation, or *abstraction*, of the traffic network model as follows. Each element of the box partition corresponds to a single state in the resulting finite state transition system, and

to obtain a computationally tractable approach, we propose a method for efficiently obtaining a finite state abstraction using the componentwise monotonicity properties developed above:

Definition 4 (CM-induced finite state abstraction). *Given a box partition $\{\mathcal{I}_q\}_{q\in\mathcal{Q}}$ of $\mathcal{X}$, the nondeterministic componentwise monotonicity-induced (CM-induced) finite state abstraction, or simply the* finite state abstraction*, of the traffic model is the transition system $\mathcal{T} = (\mathcal{Q}, \mathcal{S}, \rightarrow)$ where $\mathcal{Q}$ is the index set of the box partition, $\mathcal{S}$ is the available signaling inputs, and $\rightarrow$ is defined by:*

$$(q, \boldsymbol{s}, q') \in\rightarrow \quad \text{if and only if} \quad \mathcal{I}_{q'} \cap \overline{\text{Post}}(\textbf{cl}(\mathcal{I}_q), \boldsymbol{s}) \neq \emptyset. \tag{29}$$

Remark 3. We must take the closure of $\mathcal{I}_q$ in (29) as the $\overline{\text{Post}}$ operator and relevant properties (*e.g.*, (28)) assume a closed box. This allows efficient algorithms for constructing $\rightarrow$ via (29) as detailed below.

Note that the CM-induced finite state abstraction is nondeterministic. Nondeterminism arises from the disturbance input $\boldsymbol{d}$ and from the fact that a collection of continuous states is abstracted to one discrete state.

By the definition of the finite state abstraction above, for any trajectory $\boldsymbol{x}[t]$, $t \in \mathbb{N}$ generated by the traffic model under input sequence $\boldsymbol{s}[t]$, $t \in \mathbb{N}$, there exists a unique sequence $q[t]$, $t \in \mathbb{N}$ with each $q[t] \in \mathcal{Q}$ such that $x[t] \in \mathcal{I}_{q[t]}$ and $q[t] \overset{\boldsymbol{s}[t]}{\rightarrow} q[t+1]$. A transition system satisfying this property is said to be a *discrete abstraction* of the dynamical system (14). A controller synthesized from the abstraction to satisfy an LTL formula as described in Section V can be applied to the original traffic network with the same guarantees because the abstraction *simulates* the original traffic network [6]. However, abstractions generally result in unavoidable conservatism, that is, nonexistence of an appropriate control strategy from the abstraction does not imply nonexistence of a control strategy for the original traffic network.

The following corollary to Remark 2 implies that the finite state abstraction suggested in Definition 4 does not introduce excessive conservatism; specifically, Corollary 2 tells us that if $(q, \boldsymbol{s}, q') \in\rightarrow$, then for each link ℓ, it is possible for the state of link ℓ to transition from a state in box $\mathcal{I}_q$ to a state in $\mathcal{I}_{q'}$.

Corollary 2. *For the CM-induced finite state abstraction defined above, $(q, \boldsymbol{s}, q') \in\rightarrow$ if and only if*

$$\exists \boldsymbol{d} = \{d_\ell\}_{\ell\in\mathcal{L}} \in \mathcal{D}, \exists \boldsymbol{x}' = \{x'_\ell\}_{\ell\in\mathcal{L}} \in \mathcal{I}_{q'} \text{ such that} \tag{30}$$
$$\forall \ell \in \mathcal{L}, \exists \boldsymbol{x} \in \textbf{cl}(\mathcal{I}_q) \text{ s.t. } x'_\ell = F_\ell(\boldsymbol{x}_\ell^{loc}[t], \boldsymbol{s}^{loc}[t], d_\ell[t]). \tag{31}$$

Proof. (if). Suppose (30)–(31) holds for some $q, q' \in \mathcal{Q}$ and $\boldsymbol{s} \in \mathcal{S}$, and let $\boldsymbol{d} \in \mathcal{D}$ and $\boldsymbol{x}' \in \mathcal{I}_{q'}$ be a particular solution such that (31) holds for all ℓ. We will show that $\boldsymbol{x}' \in \overline{\text{Post}}(\textbf{cl}(\mathcal{I}_q), \boldsymbol{s})$. Let i^* be such that $\boldsymbol{d} \in \mathcal{D}_{i^*}$, and let $\underline{\boldsymbol{\xi}}^\ell$, $\overline{\boldsymbol{\xi}}^\ell$ be as in Corollary 1. We must have

$$F_\ell(\underline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \underline{d}_\ell^{i^*}) \leq x'_\ell \leq F_\ell(\overline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \overline{d}_\ell^{i^*}) \tag{32}$$

by Lemma 2 where we make the same substitutions as in the proof of Corollary 1 because (24) holds for $\boldsymbol{x}$ satisfying (31)

```
1: function ABSTRACTION(network model, D, {I_q}_{q∈Q})
       returns T
2:    inputs: network model, a traffic network model with
               update functions {F_ℓ}_{ℓ∈L} with domain X
               and signal input set S
3:            D, the disturbance set D = ∪_{i=1}^{n_D} D^i
4:            {I_q}_{q∈Q}, a box partition X
5:    →:= ∅
6:    for each s ∈ S do
7:       for each q ∈ Q do
8:          for i := 1 to n_D do
9:             ξ^ℓ := as in (26)
10:            ξ̄^ℓ := as in (27)
11:            y := F_ℓ(ξ^ℓ, s_ℓ^loc, d_ℓ^i)
12:            ȳ := F_ℓ(ξ̄^ℓ, s_ℓ^loc, d̄_ℓ^i)
13:            Q' := SUCCESSORS(y, ȳ, {I_q}_{q∈Q})
14:            →:=→ ∪(q × s × Q')
15:         end for
16:      end for
17:   end for
18:   return T := (Q, S, →)
19: end function
```

Fig. 3. Algorithm for computing a finite state abstraction of the traffic dynamics. The algorithm requires function SUCCESSORS, which can be implemented using different algorithms, depending on the structure of the box partition.

for each $\ell \in \mathcal{L}$. By (28), it follows that $\boldsymbol{x}' \in \overline{\text{Post}}(\textbf{cl}(\mathcal{I}_q), \boldsymbol{s})$, and thus $(q, \boldsymbol{s}, q') \in\rightarrow$.

(only if). Suppose $(q, \boldsymbol{s}, q') \in\rightarrow$, it follows that $\mathcal{I}_{q'} \cap \overline{\text{Post}}(\textbf{cl}(\mathcal{I}_q), \boldsymbol{s}) \neq \emptyset$, let $\boldsymbol{x}' \in \mathcal{I}_{q'} \cap \overline{\text{Post}}(\textbf{cl}(\mathcal{I}_q), \boldsymbol{s})$ and let $i^* \in \{1, \ldots, n_\mathcal{D}\}$ be such that $F_\ell(\underline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \underline{d}_\ell^{i^*}) \leq x'_\ell \leq F_\ell(\overline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \overline{d}_\ell^{i^*})$ for all $\ell \in \mathcal{L}$. Remark 2 implies that for each ℓ, there exists $\boldsymbol{x} \in \textbf{cl}(\mathcal{I}_q)$ and $d_\ell^\dagger \in [\underline{d}_\ell^{i^*}, \overline{d}_\ell^{i^*}]$ such that $x'_\ell = F_\ell(\boldsymbol{x}_\ell^{\text{loc}}, \boldsymbol{s}_\ell^{\text{loc}}, d_\ell^\dagger)$. Indeed, suppose not, then

$$\tilde{x}_\ell \triangleq \sup_{\boldsymbol{x}\in\textbf{cl}(\mathcal{I}_q), d_\ell\in[\underline{d}_\ell^{i^*}, \overline{d}_\ell^{i^*}]} F_\ell(\boldsymbol{x}_\ell^{\text{loc}}, \boldsymbol{s}_\ell^{\text{loc}}, d_\ell) < x'_\ell, \quad \text{or} \tag{33}$$

$$\underset{\sim}{x}_\ell \triangleq \inf_{\boldsymbol{x}\in\textbf{cl}(\mathcal{I}_q), d_\ell\in[\underline{d}_\ell^{i^*}, \overline{d}_\ell^{i^*}]} F_\ell(\boldsymbol{x}_\ell^{\text{loc}}, \boldsymbol{s}_\ell^{\text{loc}}, d_\ell) > x'_\ell. \tag{34}$$

If (33) holds, then $F_\ell(\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{d}) \leq \tilde{x}_\ell < x'_\ell \leq F_\ell(\overline{\boldsymbol{\xi}}^\ell, \boldsymbol{s}_\ell^{\text{loc}}, \overline{d}_\ell^{i^*})$ for all $\underline{\boldsymbol{x}}_q \leq \boldsymbol{x} \leq \overline{\boldsymbol{x}}_q$ and all $\underline{\boldsymbol{d}}^{i^*} \leq \boldsymbol{d} \leq \overline{\boldsymbol{d}}^{i^*}$, which implies the upper bound in (24) is not achieved, contradicting the first statement of Remark 2. A symmetric argument shows that if (34) holds, then Remark 2 is again contradicted. Defining $\boldsymbol{d} = \{d_\ell^\dagger\}_{\ell\in\mathcal{L}}$ for the particular collection $\{d_\ell^\dagger\}_{\ell\in\mathcal{L}}$ above implies that (30)–(31) holds, completing the proof. $\square$

We remark that, in (31), the same choice of $x \in \textbf{cl}(\mathcal{I}_q)$ will generally not work for all $\ell \in \mathcal{L}$ due to the over-approximation of the reachable set; see [24] for further discussion.

B. Constructing The Transition System $\mathcal{T}$

We begin with the primary algorithm for calculating $\mathcal{T}$ shown in Fig. 3, which relies on Corollary 1 to compute $\overline{\text{Post}}$ and to construct the finite state abstraction as defined

in Definition 4. This algorithm requires a function called SUCCESSORS that takes the lower and upper corners of a box Y as input, as well as a box partition of $\mathcal{X}$, and returns the indices of the box partitions which intersects Y. We first present a generic algorithm for SUCCESSORS applicable to any box partition. To this end, consider the nonempty box $\mathcal{I}_q = \{\boldsymbol{x} \mid \underline{\boldsymbol{x}} \prec_q^1 \boldsymbol{x} \prec_q^2 \overline{\boldsymbol{x}}\}$ and let $Y \triangleq \{\boldsymbol{x} \mid \underline{\boldsymbol{y}} \leq \boldsymbol{x} \leq \overline{\boldsymbol{y}}\}$. It is straightforward to show that $\mathcal{I}_q \cap Y \neq \emptyset$ if and only if $\underline{\boldsymbol{x}} \prec_q^1 \overline{\boldsymbol{y}}$ and $\underline{\boldsymbol{y}} \prec_q^2 \overline{\boldsymbol{x}}$.

The algorithm in Fig. 4 utilizes this fact to compute $\mathcal{Q}'$, the indices of the partitions that intersect a box defined by the corners $\underline{\boldsymbol{y}}$ and $\overline{\boldsymbol{y}}$. The algorithm is convenient because it works for any box partition of $\mathcal{X}$, however it requires comparing the corners $\underline{\boldsymbol{y}}, \overline{\boldsymbol{y}}$ to the corners of each box $\mathcal{I}_q$, $q \in \mathcal{Q}$. Thus, computing $\mathcal{T}$ scales quadratically with $|\mathcal{Q}|$ since we must determine if $\overline{\text{Post}}(s, \mathcal{I}_q)$ intersects each box $\mathcal{I}_{q'}$, $q' \in \mathcal{Q}$ for each $q \in \mathcal{Q}$.

However, the general algorithm in Fig. 4 fails to take into account any structure in the partition itself. For example, for gridded box partitions, we can identify $\mathcal{Q}'$ by comparing the corners $\underline{\boldsymbol{y}}, \overline{\boldsymbol{y}}$ componentwise to the partition's constituent coordinate intervals. For simplicity of presentation, we consider gridded box partitions $\{\mathcal{I}_q\}_{q \in \mathcal{Q}}$ where, for each $\ell \in \mathcal{L}$, there exists a set of intervals $\{I_1^\ell, \ldots, I_{N_\ell}^\ell\}$ of the form

$$I_1^\ell = [\eta_0^\ell, \eta_1^\ell], \quad I_j^\ell = (\eta_{j-1}^\ell, \eta_j^\ell], \; j = 2, \ldots, N_\ell \quad (35)$$

for $0 = \eta_0^\ell \leq \eta_1^\ell < \eta_2^\ell < \ldots < \eta_{N_\ell - 1}^\ell < \eta_{N_\ell}^\ell = x_\ell^{\text{cap}}$ such that $\mathcal{I}_q = \prod_{\ell \in \mathcal{L}} I_{q_\ell}^\ell$ for $q = \{q_\ell\}_{\ell \in \mathcal{L}} \in \mathcal{Q} \cong \prod_{\ell \in \mathcal{L}} \{1, \ldots, N_\ell\}$. Define

$$\overline{j}_\ell = \begin{cases} 1 & \text{if } \overline{y}_\ell = 0 \\ \max_{j \in \{1, \ldots, N_\ell\}} j \text{ s.t. } \eta_{j-1}^\ell < \overline{y}_\ell & \text{else} \end{cases} \quad (36)$$

$$\underline{j}_\ell = \min_{j \in \{1, \ldots, N_\ell\}} j \text{ s.t. } \underline{y}_\ell \leq \eta_j^\ell \quad (37)$$

and let $\mathcal{Q}' = \{\{q_\ell\}_{\ell \in \mathcal{L}} \mid q_\ell \in \{\underline{j}_\ell, \underline{j}_\ell + 1, \ldots, \overline{j}_\ell\}\}$. Then $\mathcal{I}_q \cap Y \neq \emptyset$ if and only if $q' \in \mathcal{Q}'$. Thus, to determine the partitions $\mathcal{Q}'$ that intersect a given box Y, we simply identify the indices of the intervals that intersects Y along each dimension. Finding $\underline{j}_\ell$ and $\overline{j}_\ell$ can be done in $O(N_\ell)$ time for each ℓ, thus solving for $\mathcal{Q}'$ requires $O(|\mathcal{L}| \max_{\ell \in \mathcal{L}} \{N_\ell\})$ time. Thus, for gridded box partitions, we can instead use the implementation of SUCCESSORS found in Fig. 5.

The algorithm in Fig. 5 may be applied to nongridded box partitions with some modification. In particular, a nongridded box partition $\{\mathcal{I}_q\}_{q \in Q}$ can be refined to obtain the coarsest possible gridded box partition with the property that each box $\mathcal{I}_q$ is the union of boxes from the refinement. This refinement is used as an index set; to compute the possible transitions from $\mathcal{I}_q$ for $q \in \mathcal{Q}$ under signaling $\boldsymbol{s} \in \mathcal{S}$, we compute $\underline{\boldsymbol{y}}$ and $\overline{\boldsymbol{y}}$ as in lines 11 and 12 of the algorithm in Fig. 3, and then use the refinement along with the algorithm in Fig. 5 to determine $\mathcal{Q}'$, the set of intersected boxes. The refinement does not introduce additional states in the transition system or require addition reach computations; it is only used to efficiently determine $\mathcal{Q}'$. For example, the coarsest refinement of Fig. 2(b) partitions the box labeled q_5 into four boxes, which are all labeled q_5. This method will be faster if the total number of intervals in the

```
1: function SUCCESSORS(y, ȳ, {I_q}_{q∈Q}) returns Q'
2:     inputs: y and ȳ, points in domain X
3:            {I_q}_{q∈Q}, an interval partition of X
4:     initialize: Q' = ∅
5:     for each q' ∈ Q do
6:         if (x_q ≺¹_q ȳ)∧(y ≺²_q x̄_q) then
7:             Q' := Q' ∪ {q'}
8:         end if
9:     end for
10:    return Q'
11: end function
```

Fig. 4. A generic algorithm for overapproximating successor states applicable to any box partition. The algorithm returns $\mathcal{Q}'$, the set of indices of boxes that intersect the box defined by the corners $\underline{\boldsymbol{y}}, \overline{\boldsymbol{y}}$, that is, $q' \in \mathcal{Q}'$ if and only if $\mathcal{I}_{q'} \cap \{\boldsymbol{x} \in \mathcal{X} \mid \underline{\boldsymbol{y}} \leq \boldsymbol{x} \leq \overline{\boldsymbol{y}}\} \neq \emptyset$.

```
1: function SUCCESSORS(y, ȳ, {I_q}_{q∈Q}) returns Q'
2:     inputs: y = {y_ℓ}_{ℓ∈L} and ȳ = {ȳ_ℓ}_{ℓ∈L},
                   points in domain X
3:            Q, a grid interval partition of X
4:     for each ℓ ∈ L do
5:         j̄_ℓ := as in (36)
6:         j_ℓ := as in (37)
7:     end for
8:     return Q' := {(j_ℓ)_{ℓ∈L} | j_ℓ ∈ {j_ℓ,...,j̄_ℓ} ∀ℓ ∈ L}
9: end function
```

Fig. 5. An algorithm for identifying successor states when $\mathcal{Q}$ is a gridded box partition.

refinement is less than $|\mathcal{Q}|$.

C. Augmenting the State Space with Signaling

To capture control objectives that include the state of the signals themselves (which are modeled as inputs in the finite state abstraction $\mathcal{T}$), we augment the discrete state space. Examples of specifications that require this augmention include φ_2 and φ_3 above or the specifications "the state of an intersection cannot change more than once per n^{min} time steps" or "an input signal cannot remain unchanged for n^{max} time steps." In particular, we propose augmenting the finite state abstraction to encompass both the current state of the finite state abstraction and the current state of the traffic signals.

Definition 5 (Augmented finite state abstraction). *The augmented finite state abstraction of the traffic network is the transition system $\mathcal{T}_{aug} = (\mathbb{Q}, \mathcal{S}, \rightarrow_{aug})$ where*
- *$\mathbb{Q} = \mathcal{Q} \times \mathcal{S}$ is the set of discrete states consisting of the box partition index set and the set of allowed input signals,*
- *$\mathcal{S}$ is the set of allowed input signals,*
- *$\rightarrow_{aug} \subseteq \mathbb{Q} \times \mathcal{S} \times \mathbb{Q}$ is the set of transitions given by $((q, \boldsymbol{\sigma}), \boldsymbol{s}, (q', \boldsymbol{\sigma}')) \in \rightarrow_{aug}$ for $(q, \boldsymbol{\sigma}), (q', \boldsymbol{\sigma}') \in \mathbb{Q}$ if and only if $(q, \boldsymbol{s}, q') \in \rightarrow$ and $\boldsymbol{\sigma}' = \boldsymbol{s}$.*

VII. SYNTHESIZING CONTROLLERS FROM LTL SPECIFICATIONS

A. Synthesis Summary

We omit the details of how a control strategy is synthesized from the nondeterministic transition system $\mathcal{T}_{\text{aug}}$ for a given

LTL control objective, as this is well-documented in the literature, see *e.g.* [14], [35]. Instead, we summarize the main steps of this synthesis as follows: from the LTL control objective, we obtain a deterministic Rabin automaton that accepts all and only trajectories that satisfy the LTL specification using off-the-shelf software. We then construct the synchronous product of the Rabin automaton and $\mathcal{T}_{\text{aug}}$ in Definition 5, resulting in a nondeterministic product Rabin automaton from which a control strategy is found by solving a Rabin game [35]. The result is a control strategy for which trajectories of the traffic network are guaranteed to satisfy the LTL specification.

As the discrete state space is finite, the signaling control strategy takes the form of a collection of "lookup" tables over the discrete states of the system, $\mathbb{Q}$, and there is one such table for each state in the Rabin automaton. Thus, implementing the control strategy requires implementing the underlying deterministic transition system of the specification Rabin automaton, which is interpreted as a finite memory controller that "tracks" progress of the LTL specification and updates at each time step. Given the current state of the Rabin transition system, the controller chooses the signaling input dictated by the current state of the augmented system $\mathbb{Q}$. Thus, we obtain a state feedback, finite memory controller. Additionally, the controller update only requires knowledge of the currently occupied partition of $\mathcal{Q}$, and thus does not require precise knowledge of the state $\boldsymbol{x}$.

B. Computational Requirements

For each $q \in \mathcal{Q}$ and each $\boldsymbol{s} \in \mathcal{S}$, determining the set $\{q' \mid q \xrightarrow{\boldsymbol{s}} q'\}$ requires first computing $\overline{\text{Post}}(\mathcal{I}_q, \boldsymbol{s})$, which requires computing $F_\ell(\cdot)$ at $2n_\mathcal{D}$ points for each $\ell \in \mathcal{L}$. Since $F_\ell(\cdot, \boldsymbol{s}, \cdot)$ is only a function of the links in $\mathcal{L}^{\text{loc}}$, each computation of this function requires time $O(1)$ assuming the average number of links at an intersection does not change with network size. Thus $\overline{\text{Post}}(\mathcal{I}_q, \boldsymbol{s})$ is computed in time $O(|\mathcal{L}|n_\mathcal{D})$. Then, we identify the set $\mathcal{Q}'$ of boxes that intersect $\overline{\text{Post}}(\mathcal{I}_q, \boldsymbol{s})$. As described in Section VI-B, this requires $2|\mathcal{Q}|$ comparisons of vectors of length $|\mathcal{L}|$ and thus is done in time $O(|\mathcal{Q}||\mathcal{L}|)$ via the algorithm in Fig. 4. However, for gridded box partitions, $\mathcal{Q}'$ is computed in time $O(|\mathcal{L}| \max_{\ell \in \mathcal{L}} \{N_\ell\})$ by the algorithm in Fig. 5. Even for nongridded box partitions, $\mathcal{Q}'$ can be computed in time $O(|\mathcal{L}| \max_{\ell \in \mathcal{L}} \{N_\ell\})$ where N_ℓ is interpreted as the number of intervals of link ℓ resulting from the coarsest refinement of the box partition that results in a gridded box partition. For a gridded partition, $|\mathcal{Q}| = \prod_{\ell \in \mathcal{L}} N_\ell$ and thus the number of boxes grows exponentially with the number of links in the network. For a nongridded box partition, the number of partitions can be substantially lower. Since $\{q' \mid q \xrightarrow{\boldsymbol{s}} q'\}$ must be computed for each q and $\boldsymbol{s}$, constructing $\mathcal{T}$ requires time $O(|\mathcal{Q}|^2|\mathcal{S}||\mathcal{L}|^2 n_\mathcal{D})$ when using the algorithm in Fig. 4 or time $O(|\mathcal{Q}||\mathcal{S}| \max_{\ell \in \mathcal{L}} \{N_\ell\}|\mathcal{L}|^2 n_\mathcal{D})$ for the algorithm in Fig. 5.

We briefly compare these computational requirements to that of polyhedral methods such as those in [14]. As the dynamics in (8)–(10) are piecewise affine, such methods can in principle be applied here. Computing $\text{Post}(\mathcal{I}_q, \boldsymbol{s})$ requires polyhedral affine transformations and polyhedral geometric
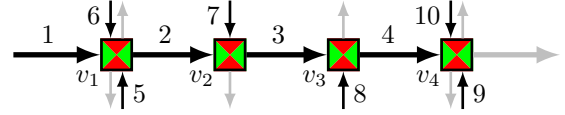

Fig. 6. Signalized network consisting of a major corridor road (links 1, 2, 3, and 4) which intersects minor cross streets (links 5, 6, 7, 8, 9, and 10). The gray links are not explicitly modeled.

sums, operations that scale exponentially in $|\mathcal{L}|$ [36], [37]. To determine if $\text{Post}(\mathcal{I}_q, \boldsymbol{s})$ intersects another polytope, geometric differences are required, which again scales exponentially with $|\mathcal{L}|$.

VIII. CASE STUDY

We consider the example network in Fig. 6 which consists of a main corridor (links 1, 2, 3, and 4) with intersecting cross streets (links 5, 6, 7, 8, 9, and 10) and four intersections, a commonly encountered network configuration. The gray links exit the network and are not explicitly modeled. The network parameters are $(x_1^{\text{cap}}, \ldots, x_{10}^{\text{cap}}) = (40, 50, 50, 50, 40, 40, 40, 40, 40, 40)$, $(c_1, \ldots, c_{10}) = (20, 20, 20, 20, 10, 10, 10, 10, 10, 10)$, $\beta_{12} = \beta_{23} = \beta_{34} = \beta_{62} = \beta_{52} = 0.5$, $\beta_{73} = \beta_{84} = 0.9$, $\alpha_{62}^{\{1\}} = \alpha_{52}^{\{1\}} = 0.5$, and all other supply ratios are one, where the time step is 15 seconds. We assume

$$\mathcal{D} = \{\boldsymbol{d} \mid \boldsymbol{0} \le \boldsymbol{d} \le [10\ 0\ 0\ 0\ 10\ 10\ 0\ 0\ 10\ 10]\}$$
$$\cup \{\boldsymbol{d} \mid \boldsymbol{0} \le \boldsymbol{d} \le [10\ 0\ 0\ 0\ 10\ 10\ 10\ 10\ 0\ 0]\}. \quad (38)$$

We further assume the available signals are $\mathcal{S}_{v_1} = \{\{1\}, \{5, 6\}\}$, $\mathcal{S}_{v_2} = \{\{2\}, \{7\}\}$, $\mathcal{S}_{v_3} = \{\{3\}, \{8\}\}$, and $\mathcal{S}_{v_4} = \{\{4\}, \{9, 10\}\}$. We wish to find a control policy for the four signalized intersections that satisfies the LTL property $\varphi = \varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \varphi_4$ where

$$\varphi_1 = \Box \Diamond (\boldsymbol{s}_{v_1} = \{5, 6\}) \wedge \Box \Diamond (\boldsymbol{s}_{v_2} = \{7\})$$
$$\wedge \Box \Diamond (\boldsymbol{s}_{v_3} = \{8\}) \wedge \Box \Diamond (\boldsymbol{s}_{v_4} = \{9, 10\}) \quad (39)$$

"Each signal actuates cross street traffic infinitely often"

$$\varphi_2 = \Diamond \Box \big((x_1 \le 30) \wedge (x_2 \le 30) \wedge (x_3 \le 30) \wedge (x_4 \le 30) \big) \quad (40)$$

"Eventually, links 1, 2, 3, and 4 have fewer than 30 vehicles on each link and this remains true for all time"

$$\varphi_3 = \Box \big(\neg(\boldsymbol{s}_{v_4} = \{4\}) \wedge \bigcirc(\boldsymbol{s}_{v_4} = \{4\}) \rightarrow \bigcirc\bigcirc(\boldsymbol{s}_{v_4} = \{4\}) \big) \quad (41)$$

$$\varphi_4 = \Box \big(\neg(\boldsymbol{s}_{v_4} = \{9, 10\}) \wedge \bigcirc(\boldsymbol{s}_{v_4} = \{9, 10\}) \rightarrow \bigcirc\bigcirc(\boldsymbol{s}_{v_4} = \{9, 10\}) \big) \quad (42)$$

For φ_3 (resp. φ_4), "The signal at intersection v_4 must actuate corridor traffic (resp. cross street traffic) for at least two sequential time-steps."

Thus, φ_2 reflects our preference for actuating corridor traffic and ensures that eventually, links 2, 3, and 4 have "adequate supply" because if the number of vehicles on these links is less than 30, then these links can always accept upstream demand, thus avoiding *congestion* (congestion occurs when demand is greater than supply). Condition φ_1 ensures that, despite the
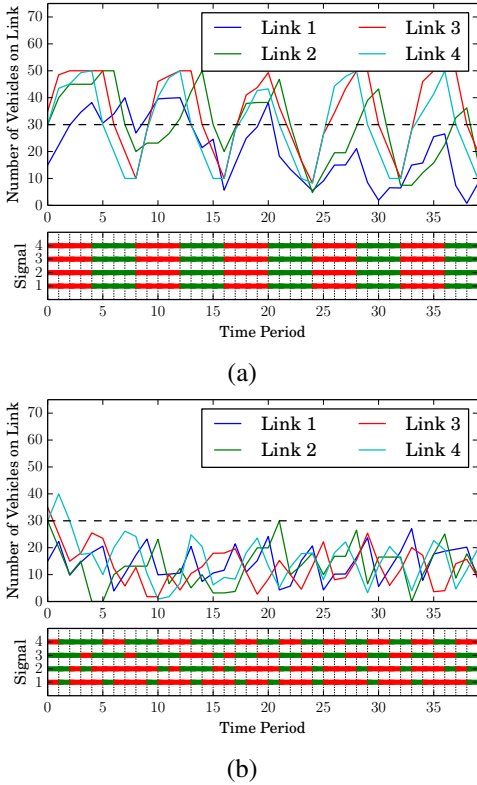
Fig. 7. (a) A sample trajectory of a naïve strategy that alternately actuates corridor traffic and then cross street traffic for four time steps each in a synchronized fashion. This policy does not satisfy the desired control objective, in particular, (40) is not satisfied. (b) A sample trajectory resulting from the synthesized control policy that is guaranteed to satisfy the LTL policy (39)–(42). In the lower plots of (a) and (b), green (resp., red) for the signal trace indicates corridor traffic (resp., cross street traffic) is actuated.

preference for facilitating traffic along the corridor, we must infinitely often actuate traffic at the cross streets. Conditions φ_3 and φ_4 are needed if, *e.g.* there exists crosswalks at intersection v_4 and a minimum amount of time is required to allow pedestrians to cross. Note that repeated application of the $\bigcirc$ ("next") operator allows us to consider finite time horizons as in (41) and (42).

We partition the state space into 408 boxes that favors larger boxes when there are fewer total vehicles in the network. There are 16 signaling inputs, and thus, the number of states in the transition system $\mathcal{T}_{\text{aug}}$ is $|\mathbb{Q}| = 6528$. The Rabin automaton generated from φ contains 62 states and one acceptance pair. Computing the finite state abstraction $\mathcal{T}$ took 22.4 seconds. In contrast, the computation would be intractable using polyhedral methods. Computing the product automaton took 30.9 minutes and computing the control strategy took 15.5 minutes on a Macbook Pro with a 2.3 GHz processor where we use the Rabin game solver in `conPAS2` [14], however `conPAS2` is written in MATLAB and the synthesis process is likely to be much more efficient if implemented in C or C++ and optimized. Furthermore, all computations can be performed offline and some are parallelizable, such as computing the product automaton. Finally, we note that the computed control strategy is implemented with minimal online costs.

Fig. 7(a) shows a sample trajectory of the network using a naïve coordinated signaling strategy whereby each intersection actuates corridor traffic for three time steps and then cross traffic for three time steps. The exogenous disturbance is generated uniformly randomly from $\mathcal{D}$. The trajectories are not guaranteed to satisfy the control objective, in particular, φ_2 is violated. Fig. 7(b) shows a sample trajectory of the system with a control strategy synthesized using the finite state abstraction augmented with signal history and the LTL requirement above. The control strategy is correct-by-construction and thus guaranteed to satisfy φ from any initial state.

We see that the synthesized controller reacts to increased vehicles on the corridor by actuating the corridor links, thereby preventing congestion (inadequate supply) along the corridor. At the same time, the controller actuates cross streets when doing so does not adversely affect conditions on the corridor (*i.e.*, cause congestion). In contrast, the fixed time controller in Fig 7(a) is not able to react to the current conditions of the network and fails to prevent congestion along the corridor; in fact, links 2, 3, and 4 periodically reach full capacity.

IX. CONCLUSIONS

We have proposed a framework for synthesizing a control strategy for a traffic network that ensures the resulting traffic dynamics satisfy a control objective expressed in linear temporal logic (LTL). In addition to offering a novel domain for applying formal methods tools in a control theory setting, we have identified and exploited key properties of traffic networks to allow efficient computation of a finite state abstraction.

Future research will investigate systematic methods for determining an appropriate box partition to further reduce the number of states in the computed abstraction. Additionally, traffic networks are often composed of tightly coupled neighborhoods and towns connected by sparse longer roads, and such networks may be amenable to a compositional formal methods approach using an assume-guarantee framework [5].

REFERENCES

[1] T. Wongpiromsarn, T. Uthaicharoenpong, Y. Wang, E. Frazzoli, and D. Wang, "Distributed traffic signal control for maximum network throughput," in *Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on*, pp. 588–595, Sept 2012.

[2] P. Varaiya, "The max-pressure controller for arbitrary networks of signalized intersections," in *Advances in Dynamic Network Modeling in Complex Transportation Systems*, pp. 27–66, Springer, 2013.

[3] P. Varaiya, "Max pressure control of a network of signalized intersections," *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 177–195, 2013.

[4] M. Papageorgiou, C. Diakaki, V. Dinopoulou, A. Kotsialos, and Y. Wang, "Review of road traffic control strategies," *Proceedings of the IEEE*, vol. 91, no. 12, pp. 2043–2067, 2003.

[5] E. M. Clarke, O. Grumberg, and D. A. Peled, *Model checking*. MIT press, 1999.

[6] C. Baier and J. Katoen, *Principals of Model Checking*. MIT Press, 2008.

[7] P. Tabuada and G. Pappas, "Linear time logic control of discrete-time linear systems," *IEEE Transactions on Automatic Control*, vol. 51, no. 12, pp. 1862–1877, 2006.

[8] P. Tabuada, "Controller synthesis for bisimulation equivalence," *Systems & Control Letters*, vol. 57, no. 6, pp. 443–452, 2008.

[9] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for dynamic robots," *Automatica*, vol. 45, no. 2, pp. 343–352, 2009.

[10] H. Kress-Gazit, G. Fainekos, and G. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, pp. 1370–1381, Dec 2009.

[11] M. Kloetzer and C. Belta, "Automatic deployment of distributed teams of robots from temporal logic motion specifications," *IEEE Transactions on Robotics*, vol. 26, pp. 48–61, Feb 2010.

[12] A. Abate, A. D'Innocenzo, and M. Di Benedetto, "Approximate abstractions of stochastic hybrid systems," *IEEE Transactions on Automatic Control*, vol. 56, pp. 2688–2694, Nov 2011.

[13] T. Wongpiromsarn, U. Topcu, and R. Murray, "Receding horizon temporal logic planning," *IEEE Transactions on Automatic Control*, vol. 57, pp. 2817–2830, Nov 2012.

[14] B. Yordanov, J. Tůmová, I. Černá, J. Barnat, and C. Belta, "Temporal logic control of discrete-time piecewise affine systems," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1491–1504, 2012.

[15] E. A. Gol, M. Lazar, and C. Belta, "Language-guided controller synthesis for linear systems," *IEEE Transactions on Automatic Control*, vol. 59, pp. 1163–1176, May 2014.

[16] A. A. Julius and A. K. Winn, "Safety controller synthesis using human generated trajectories: Nonlinear dynamics with feedback linearization and differential flatness," in *Proceedings of the 2012 American Control Conference*, pp. 709–714, 2012.

[17] U. Topcu, N. Ozay, J. Liu, and R. M. Murray, "On synthesizing robust discrete controllers under modeling uncertainty," in *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control*, HSCC '12, (New York, NY, USA), pp. 85–94, ACM, 2012.

[18] J. Liu, N. Ozay, U. Topcu, and R. Murray, "Synthesis of reactive switching protocols from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 58, pp. 1771–1785, July 2013.

[19] E. Aydin Gol, M. Lazar, and C. Belta, "Temporal logic model predictive control for discrete-time systems," in *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, pp. 343–352, ACM, 2013.

[20] E. Plaku, L. E. Kavraki, and M. Y. Vardi, "Falsification of LTL safety properties in hybrid systems," *International Journal on Software Tools for Technology Transfer*, vol. 15, no. 4, pp. 305–320, 2013.

[21] S. Coogan and M. Arcak, "Freeway traffic control from linear temporal logic specifications," in *Proceedings of the 5th ACM/IEEE International Conference on Cyber-Physical Systems*, pp. 36–47, 2014.

[22] M. Kulenovic and O. Merino, "A global attractivity result for maps with invariant boxes," *Discrete and Continuous Dynamical Systems Series B*, vol. 6, no. 1, p. 97, 2006.

[23] T. Moor and J. Raisch, "Abstraction based supervisory controller synthesis for high order monotone continuous systems," in *Modelling, Analysis, and Design of Hybrid Systems*, pp. 247–265, Springer, 2002.

[24] S. Coogan, E. Aydin Gol, M. Arcak, and C. Belta, "Controlling a network of signalized intersections from temporal logical specifications," in *American Control Conference (ACC)*, 2015. To appear.

[25] J. Lebacque, "Intersection modeling, application to macroscopic network traffic flow models and traffic management," in *Traffic and Granular Flow'03*, pp. 261–278, Springer, 2005.

[26] C. F. Daganzo, "The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory," *Transportation Research Part B: Methodological*, vol. 28, no. 4, pp. 269–287, 1994.

[27] J. C. Munoz and C. F. Daganzo, "The bottleneck mechanism of a freeway diverge," *Transportation Research Part A: Policy and Practice*, vol. 36, no. 6, pp. 483–505, 2002.

[28] M. Kieffer, L. Jaulin, and E. Walter, "Guaranteed recursive non-linear state bounding using interval analysis," *International Journal of Adaptive Control and Signal Processing*, vol. 16, no. 3, pp. 193–218, 2002.

[29] M. Papageorgiou, "An integrated control approach for traffic corridors," *Transportation Research Part C: Emerging Technologies*, vol. 3, no. 1, pp. 19–30, 1995.

[30] F. H. Clarke, *Optimization and nonsmooth analysis*, vol. 5. Siam, 1990.

[31] M. W. Hirsch, "Systems of differential equations that are competitive or cooperative II: Convergence almost everywhere," *SIAM Journal on Mathematical Analysis*, vol. 16, no. 3, pp. 423–439, 1985.

[32] D. Angeli and E. Sontag, "Monotone control systems," *IEEE Transactions on Automatic Control*, vol. 48, no. 10, pp. 1684–1698, 2003.

[33] H. L. Smith, *Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems*. American Math. Soc., 1995.

[34] S. Scholtes, *Introduction to piecewise differentiable equations*. Springer, 2012.

[35] F. Horn, "Streett games on finite graphs," *Proc. 2nd Workshop Games in Design Verification (GDV)*, 2005.

[36] A. Kurzhanskiy and P. Varaiya, "Computation of reach sets for dynamical systems," in *The Control Systems Handbook*, ch. 29, CRC Press, second ed., 2010.

[37] M. Herceg, M. Kvasnica, C. Jones, and M. Morari, "Multi-Parametric Toolbox 3.0," in *Proceedings of the European Control Conference*, (Zürich, Switzerland), pp. 502–510, July 17–19 2013. http://control.ee.ethz.ch/~mpt.

Samuel Coogan is a Ph.D. candidate in Electrical Engineering and Computer Sciences at the University of California, Berkeley. He received his B.S. in Electrical Engineering from Georgia Tech in 2010 and his M.S. in Electrical Engineering from UC Berkeley in 2012. His research interests are in control theory, nonlinear and hybrid systems, and formal methods. He is particularly interested in applying techniques from these domains to the control and design of transportation systems. He received an NSF Graduate Research Fellowship in 2010 and the Leon O. Chua Award for outstanding achievement in nonlinear science from UC Berkeley in 2014.

Ebru Aydin Gol received her B.Sc. degree in computer engineering from Orta Dogu Teknik Universitesi, Ankara, Turkey, in 2008, M.Sc. degree in computer science from Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, in 2010 and Ph.D. degree in systems engineering from Boston University, Boston, MA, USA in 2014. She has been a Site Reliability Engineer at Google since 2014. Her research interests include verification and control of dynamical systems, optimal control, and synthetic biology.

Murat Arcak is a professor at U.C. Berkeley in the Electrical Engineering and Computer Sciences Department. He received the B.S. degree in Electrical Engineering from the Bogazici University, Istanbul, Turkey (1996) and the M.S. and Ph.D. degrees from the University of California, Santa Barbara (1997 and 2000). His research is in dynamical systems and control theory with applications to synthetic biology, multi-agent systems, and transportation. Prior to joining Berkeley in 2008, he was a faculty member at the Rensselaer Polytechnic Institute. He received a CAREER Award from the National Science Foundation in 2003, the Donald P. Eckman Award from the American Automatic Control Council in 2006, the Control and Systems Theory Prize from the Society for Industrial and Applied Mathematics (SIAM) in 2007, and the Antonio Ruberti Young Researcher Prize from the IEEE Control Systems Society in 2014. He is a member of SIAM and a fellow of IEEE.

Calin Belta is a Professor in the Department of Mechanical Engineering, Department of Electrical and Computer Engineering, and the Division of Systems Engineering at Boston University, where he is also affiliated with the Center for Information and Systems Engineering (CISE) and the Bioinformatics Program. His research focuses on dynamics and control theory, with particular emphasis on hybrid and cyber-physical systems, formal synthesis and verification, and applications in robotics and systems biology. Calin Belta is a Senior Member of the IEEE and an Associate Editor for the SIAM Journal on Control and Optimization (SICON) and the IEEE Transactions on Automatic Control. He received the Air Force Office of Scientific Research Young Investigator Award and the National Science Foundation CAREER Award.