

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

# A cell transmission model for dynamic lane reversal with autonomous vehicles

Michael W. Levin  
Graduate Research Assistant

Stephen D. Boyles  
Assistant Professor

February 21, 2016

# Abstract

Autonomous vehicles admit consideration of novel traffic behaviors such as reservation-based intersection controls and dynamic lane reversal. We present a cell transmission model formulation for dynamic lane reversal. For deterministic demand, we formulate the dynamic lane reversal control problem for a single link as an integer program and derive theoretical results. In reality, demand is not known perfectly at arbitrary times in the future. To address stochastic demand, we present a Markov decision process formulation. Due to the large state size, the Markov decision process is intractable. However, based on theoretical results from the integer program, we derive an effective heuristic. We demonstrate significant improvements over a fixed lane configuration both on a single bottleneck link with varying demands, and on the downtown Austin network.

**Keywords:** dynamic lane reversal, autonomous vehicles, cell transmission model, dynamic traffic assignment

## 1 Introduction

The computer precision of autonomous vehicles (AVs) admits consideration of novel traffic behaviors such as reservation-based intersection control [1, 2], which uses wireless communications and highly specified vehicle trajectories to increase intersection utilization and reduce delays [3, 4]. Similar communication and behavior protocols may be used to implement lane reversal, in which lane direction changes in response to traffic conditions. Use of lane reversal and other capacity improvements may be particularly important for AVs because induced demand and empty repositioning trips may increase traffic congestion [5].

Lane reversal has already been explored through contraflow lanes. Most literature pertains to evacuation (see, for instance, [6–8]), because of the costs associated with reversing lanes for human drivers, but several papers study contraflow for daily operations. Zhou et al. [9] use machine learning on queue length and total delay for scheduling the lane reversal. Xue & Dong [10] similarly applied neural networks on fuzzy pattern clustering to contraflow for a bottleneck tunnel. Meng et al. [11] use a bi-level optimization to address the driver response to contraflow lanes through user equilibrium (UE) behavior. As demonstrated by the Braess [12] and Daganzo [13] paradoxes, consideration of UE routing behavior is important as it can adversely affect potential network improvements. Therefore, our results include solving dynamic traffic assignment (DTA) on a city network.

The primary constraint on existing work on contraflow lanes for daily operations is communication with and ensuring safety of human drivers. Reversing a lane with human drivers therefore often requires significant time and cannot be performed frequently. Furthermore, it is impractical to perform on every road segment (link), and, where it is used, the lane is reversed on the entire link. Partial lane reversal could increase flow by adding temporary turning bays. Consequently, a more frequent dynamic lane reversal (DLR) for AVs, controlled by a lane manager agent per link in communication with AVs on the link, could result in significant improvements over contraflow lanes.

1  
2  
3  
4 Our work is primarily motivated by the greater communications available for AVs due  
5 to the frequency of lane reversals we propose for DLR. We assume that lane direction can  
6 be changed at very small intervals of space-time, such as a few hundred feet of space and  
7 6 second time steps. Such frequent reversals of lane direction can be used to optimize lane  
8 direction for small variations in demand over time. Contraflow lanes are typically reversed  
9 for the duration of a peak period, whereas DLR could change lane direction many times  
10 within a peak period to reduce queueing and spillback. However, such small space-time  
11 intervals for DLR cannot be safely implemented with human vehicles. The more precise and  
12 greater bandwidth AV communications are necessary.

13  
14 In this paper, we assume that lane manager agents exist that can communicate the  
15 direction of each lane at space and time intervals to all vehicles on the link. Hausknecht et  
16 al. [14] suggest using AV intersection controllers as a lane manager to specify the direction  
17 of lanes for the entire link at different times. With some changes the intersection controllers  
18 could communicate lane direction at space intervals as well, and we also assume that AVs  
19 could be forced to obey these policies. Therefore, rather than study an enabling protocol,  
20 we focus on the potential benefits.

21  
22 Hausknecht et al. [14] found that DLR improved capacity on a micro-simulation of a  
23 small network and used optimization techniques on the lane reversal problem for static traffic  
24 assignment (STA). A natural extension is how to model DLR and construct optimal lane  
25 direction policies for city networks with dynamic demand and more realistic flow models.  
26 Computational tractability becomes a major concern. As noted by Hausknecht et al. [14],  
27 even for a static flow model, STA becomes a subproblem to finding the DLR policy, forming  
28 a bi-level optimization problem. As the number of lanes is integer, the top level involves  
29 integer programming (IP), a potentially NP-hard problem. Dynamic demand also introduces  
30 stochasticity from the perspective of the lane manager because future conditions may not  
31 be known perfectly. Therefore, finding the optimal DLR policy could require impractical  
32 computational resources. However, a heuristic that yields consistent improvements over  
33 current fixed lane configurations would be valuable.

34  
35 This paper incorporates DLR into the cell transmission model (CTM) and studies optimal  
36 policies for DLR. We consider two types of information availability for finding the optimal  
37 DLR policy. First, when future demand is known, we study DLR in the context of IPs and  
38 present theoretical results and motivating examples. When future demand is stochastic, we  
39 formulate DLR as a Markov decision process (MDP) and present a saturation-based heuristic  
40 for computational tractability that appears to perform well on a variety of demands for a  
41 single bottleneck link. We then solve DTA on a city network using this heuristic, and  
42 demonstrate significant improvements in system efficiency.

43  
44 The remainder of this paper is organized as follows: Section 2 discusses previous work on  
45 autonomous vehicles. Section 3 presents a CTM for DLR. Section 4 studies DLR under the  
46 perfect information scenario, and Section 5 considers stochastic demand. Section 6 discusses  
47 results on a city network, and we present our conclusions in Section 7.

## 2 Literature review

Because lane reversal at small time intervals cannot be safely implemented with human drivers, DLR has only been previously studied by Hausknecht et al. [14] in the context of AVs. Furthermore, reversing lanes at multiple spatial intervals within links has never been studied. Therefore, we begin by reviewing proposed intersection technologies for AVs that could be used to enable DLR. Then, we discuss flow models for AVs to justify the use of CTM.

### 2.1 Technology for dynamic lane reversal

The precision and communications potential of AVs have been used to propose several new traffic behaviors such as DLR. A primary topic of study is improving intersection efficiency, and the communications required for the proposed intersection controller can be adapted to the requirements of DLR.

Dresner & Stone [1, 2] introduced reservation-based intersection control, in which AVs communicate with an *intersection manager* to request intersection passage. The intersection manager simulates requests on a grid of space-time tiles, which are accepted only if they do not conflict with other requests. Fajardo et al. [3] and Li et al. [4] demonstrated that reservations can reduce delays beyond optimized signals. Therefore, when AVs are a sufficiently high proportion of vehicular demand, reservations are likely to be used in place of signals [15].

The seminal DLR paper of Hausknecht et al. [14] observed that the intersection manager could be used to control lane usage by restricting AVs from entering certain lanes. This could enforce DLR by ensuring that AVs do not enter a lane in the wrong direction. Therefore, the reservation protocol is sufficient for implementing lane reversal where lanes have the same direction for each link.

In this paper, we consider lane reversal at multiple spatial intervals within a link. This can also be handled by a modification to the intersection manager. In the reservation protocol, AVs communicate with the intersection manager well before reaching the intersection to request a reservation. These longer-range communications can be used to establish lane direction at small space-time intervals and require AVs to switch lanes to comply with lane reversals.

### 2.2 Flow models for autonomous vehicles

DTA [16] admits a variety of analytical and simulation-based flow models. Many current simulation-based DTA models approximately solve the hydrodynamic theory of traffic flow [17, 18], which is tractable for large-scale modeling. This paper builds from the CTM [19, 20], which is a finite differences approximation to the hydrodynamic theory. The hydrodynamic theory can also be solved through link-based methods such as the link transmission model [21] or through an analytical Lax-Hopf method [22]. However, CTM is used because the space-time intervals of the approximation can be aligned with the intervals for DLR.

1  
2  
3  
4 One significant question is whether the hydrodynamic theory is appropriate for modeling  
5 AV flow. AVs admit new technologies such as adaptive cruise control [23], which may behave  
6 differently from humans due to computer precision and reaction times. Microsimulation  
7 studies have shown that AVs are likely to increase capacity [24, 25] and stability [26] due to  
8 reduced following headways. Nevertheless, AV behavior is still bound by the constraints of  
9 non-compressible fluid flow, such as conservation of flow. Therefore, modeling AV traffic flow  
10 by the hydrodynamic theory (possibly with a different fundamental diagram) is a reasonable  
11 assumption. Levin & Boyles [27] developed a multiclass CTM for modeling shared human and  
12 AV links, and also developed a car-following model that predicted a triangular fundamental  
13 diagram based on reduced reaction times. Due to the Courant-Friedrich-Lewy [28] condition,  
14 if the backwards wave speed is less than free flow speed, CTM can model AV traffic flow.  
15 The modifications we make to model DLR are orthogonal to the capacity and backwards  
16 wave speed parameters of CTM.  
17  
18  
19  
20  
21  
22

### 23 **3 Cell transmission model**

24  
25 The original CTM was developed by Daganzo [19, 20] as a discrete approximation to the  
26 hydrodynamic theory of traffic flow [17, 18]. In this section, we modify the CTM to include a  
27 varying number of lanes in space and time. In this paper, we make the following assumptions  
28 to ensure safe and realistic behavior:  
29  
30  
31

- 32 1. Vehicles can change lanes at most once per time step. For a typical time step of 6  
33 seconds with free flow speed of 30 miles per hour, the corresponding cell length is  
34 264 feet. That interval in space and time should be sufficient for one lane change.  
35 Lane changing may cause disruptions to the traffic stream because increases in den-  
36 sity from forcing vehicles to merge may reduce flow. This is modeled by scaling the  
37 fundamental diagram with the change in the numbers of lanes. When the number of  
38 lanes is reduced, the relative congestion increases, resulting in reductions in capacity  
39 and possibly maximum flow as per the new fundamental diagram.  
40  
41
- 42 2. The lane manager can specify the direction of each lane per cell and time step. Changes  
43 in lane direction are subject to constraints on jam density and lane changing.  
44  
45
- 46 3. All vehicles are autonomous and obey lane direction specified by the lane manager.  
47 (We do not admit human drivers because lanes cannot be safely reversed at small time  
48 and spatial intervals with human drivers.)  
49  
50
- 51 4. All lanes traveling in the same direction are contiguous. This simplifies lane changing  
52 and turning movement behavior.  
53  
54
- 55 5. DLR can be used for arterials and highway links that have a parallel, opposite direction  
56 link of the same length and free flow speed.  
57  
58  
59  
60  
61  
62  
63  
64  
65




Figure 1: Illustration of paired CTM links  $[a, b]$  and  $[b, a]$

### 3.1 Flow model

Consider a pair of links  $[a, b]$  and  $[b, a]$  from  $a$  to  $b$  and from  $a$  to  $b$ , respectively with contiguous lanes and identical free flow speed  $v$  and backwards wave speed  $w$ , so that DLR is possible and cells align. (For links without a parallel, opposite direction link, the number of lanes may be assumed to be fixed and follow the original CTM). Let  $C$  be the set of cells in  $[a, b]$ . We assume that because  $[a, b]$  and  $[b, a]$  have contiguous lanes, every cell  $i \in C$  has a parallel cell  $\overleftarrow{i}$  of the same length in the opposite direction. Link  $[a, b]$  has cells 1 through  $|C|$ .  $\overleftarrow{|C|}$  refers to the first cell of link  $[b, a]$ , and  $\overleftarrow{1}$  refers to the last cell. Figure 1 illustrates this notation. The cell length is  $v\Delta t$ , the distance a vehicle can travel in a time step of  $\Delta t$  at free flow speed  $v$ .

Assumptions 2 and 4 simplify defining the direction of each lane at each time step to specifying the number of lanes in each direction in space and time. This also opens the possibility for preventing use of a lane in any direction over some interval in spacetime. This could be used to clear a lane to reduce the congestion caused by a later lane reversal. We define a *lane policy*  $L$  to be a specification of the number of lanes for each space-time interval, denoted  $L_i(t)$  for cell  $i$  and time  $t$  or  $L(x, t)$  at position  $x \in \mathbb{R}$ .  $L(x, t)$  is used in Section 3 to verify that CTM with lane reversals is consistent with conservation of flow. In Section 3.2 we describe constraints on lane policies to follow the above assumptions. We use a trapezoidal fundamental diagram for link flow:

$$\Psi(k, x, t) = \min\{vk, QL(x, t), w(KL(x, t) - k)\} \quad (1)$$

where  $Q$  is the cell capacity per lane per time step,  $K$  is jam density per lane, and  $w$  is the backwards wave speed. As with Daganzo [19] we specify the cell transition flows, then demonstrate that they satisfy conservation of flow:  $\frac{\partial q}{\partial x} = -\frac{\partial k}{\partial t}$ . Cell occupancy  $n_i(t)$  is determined by transition flows  $y_i(t, L)$ , which depend on the lane policy  $L$ :

$$n_i(t+1) = n_i(t) - y_i(t, L) + y_{i-1}(t, L) \quad (2)$$

with

$$y_i(t, L) = \min\{S_i(t, L), R_{i+1}(t, L)\} \quad (3)$$

where

$$S_i(t, L) = \min\{n_i(t), QL_i(t)\} \quad (4)$$

is the sending flow and

$$R_i(t, L) = \min \left\{ QL_i(t), \frac{w}{v} (NL_i(t) - n_i(t)) \right\} \quad (5)$$

is the receiving flow, where  $N$  is the maximum number of vehicles that can fit in 1 lane of cell  $i$ . Since the links are interchangeable, equations (2) and (3) define cell evolution for cells 1 through  $|C|$  as well as cells  $\overleftarrow{|C|}$  through  $\overleftarrow{1}$ .

Let  $\Delta t$  be the time step and choose cell length such that  $v\Delta t = 1$ . If units are chosen so that  $\Delta t = 1$ , cell length is 1,  $v = 1$ ,  $x = i$ ,  $K = N$ , and  $k(x, t) = n_i(t)$ . This cell length satisfies the Courant-Friedrich-Lewy condition [28] for stability of these difference equations when  $w \leq v$ .

Then, as with Daganzo [19],

$$q(x, t) = \min \left\{ n_i(t), QL_i(t), QL_{i+1}(t), \frac{w}{v} (NL_{i+1}(t) - n_{i+1}(t)) \right\} = y_i(t, L) \quad (6)$$

which results in  $\frac{\partial q(x, t)}{\partial x} = y_{i+1}(t) - y_i(t)$ . Since  $\frac{\partial k(x, t)}{\partial t} = n_i(t+1) - n_i(t)$  is the rate of change in cell occupancy with respect to time, flow conservation  $\frac{\partial q}{\partial x} = -\frac{\partial k}{\partial t}$  is  $y_{i+1}(t) - y_i(t) = n_i(t) - n_i(t+1)$ , which is the cell propagation function of equation (2).

## 3.2 Constraints

In line with the aforementioned assumptions, the number of lanes per cell and time step must satisfy constraints for safety. These constraints admit lane reversals at each cell-time, and we later consider lane reversals at time intervals of 6 seconds. Therefore, we require that all vehicles are AVs because human drivers cannot safely navigate through such frequent changes in lane direction.

First, for all  $i \in C$  and for all  $t$  the total number of lanes across a cell and its opposite is limited by the maximum number of lanes available,  $\ell$ :

$$L_i(t) + L_{\overleftarrow{i}}(t) \leq \ell \quad (7)$$

We do not require equality because it may be desirable to empty a lane before reversing its direction.

Assumption 1 requires that if  $n_{\overleftarrow{i}}(t) > 0$  then

$$|L_i(t+1) - L_i(t)| \leq 1 \quad (8)$$

$$|L_{\overleftarrow{i}}(t+1) - L_{\overleftarrow{i}}(t)| \leq 1 \quad (9)$$

so that vehicles in cell  $i$  at time  $t$  that remain in  $i$  at  $t+1$  cannot be forced to change lanes more than once. Also,

$$|L_{i+1}(t+1) - L_i(t)| \leq 1 \quad (10)$$

$$|L_{\overleftarrow{i+1}}(t+1) - L_{\overleftarrow{i}}(t)| \leq 1 \quad (11)$$

so vehicles moving from cell  $i$  at time  $t$  to cell  $i + 1$  at time  $t + 1$  do not have to change lanes more than once.

When the lane direction changes, the number of vehicles in a cell could potentially exceed the jam density, which results in the following requirement:

$$NL_i(t) \geq n_i(t) \tag{12}$$

so that the available physical space in the cell (which changes based on its number of lanes) is sufficient to hold all vehicles in the cell.

### 3.3 Feasibility

The additional constraints require an analysis of feasibility. Because the initial conditions could potentially force a violation of constraint (12), a sufficient condition for feasibility is that constraints (7) through (12) are initially satisfied. This is easily achievable for DTA models that start with empty links at  $t = 0$  and load flow onto links in subsequent time steps. Proposition 1 shows that if the initial cell occupancies are feasible, then there exists a solution to DLR feasible for all time steps.

**Proposition 1.** *Let  $\mathcal{L}_T$  be the set of policies satisfying constraints (7) through (12) for  $0 \leq t \leq T$ . If for all cells  $i$ ,  $L_i(0) + L_{\leftarrow i}(0) \leq \ell$ ,  $|L_{i+1}(0) - L_i(0)| \leq 1$ ,  $|L_{\leftarrow i+1}(0) - L_{\leftarrow i}(0)| \leq 1$ ,  $NL_i(0) \geq n_i(0)$ , and  $NL_{\leftarrow i}(0) \geq n_{\leftarrow i}(0)$ , then  $\mathcal{L}_T \neq \emptyset$ .*

*Proof.* A fixed lane policy is a policy  $L$  such that for all  $i \in C$  and for all  $t$ ,  $L_i(t) = L_i(0)$ . Any fixed lane policy satisfies constraints (8) and inductively satisfies constraints (7), (10), and (12) if  $L_i(0) + L_{\leftarrow i}(0) \leq \ell$ ,  $|L_{i+1}(0) - L_i(0)| \leq 1$ , and  $NL_i(0) \geq n_i(0)$ , respectively.  $\square$

A fixed lane policy can be used to provide a lower bound on the value of the optimal DLR policy. However, naïve policies can perform worse than fixed lane policies.

## 4 Method for deterministic demand

In this section we study the optimal DLR policy for a single link when sending and receiving flows at the upstream and downstream ends are known. Although this level of knowledge may seem unrealistic, this is useful for developing theory about the DLR problem. Furthermore, the lane manager may be able to communicate with other lane and intersection managers across the city to acquire sending and receiving flows for a limited time horizon. Section 5 studies DLR policies for a single link with stochastic demand and downstream supply. Overall this paper focuses on policies from a single link perspective for computational tractability, which is nevertheless demonstrated to improve total system travel time on a city network in Section 6.






Figure 2: Example of bottleneck lane configuration

## 4.1 Motivation

We first motivate the discussion with a demonstration of the challenges in finding an optimal DLR policy. A naïve approach might consider the objective of maximizing flow on a per time step basis, i.e. at  $t$ , choose lanes to  $\max \sum_{i \in C} (y_i(t) + y_{\overleftarrow{i}}(t))$ . This objective is favorable because it exhibits the optimal substructure characteristic for constructing a dynamic programming algorithm. However, if we consider two parallel but opposite directional links with capacity 1200vph per lane, with 4 lanes between them, 4 cells, and 4800vph demand in each direction, the lane configuration shown in Figure 2, maximizes flow initially but results in a bottleneck in the middle of the link. Therefore, an optimal policy must consider future evolution of flows.

## 4.2 Integer program

Because naïve methods for DLR policies may reduce flow, we formulate the DLR problem for a single link as an IP. We then analyze this IP to derive some theoretical results about the solution that inspire our heuristic in Section 5. To be consistent with the cell notation, let cells 0 and  $\overleftarrow{|C|+1}$  be source cells connected to cells 1 and  $\overleftarrow{|C|}$ , respectively. Because we assume demand and supply for the pair of links under consideration are deterministic in this section, we model the upstream and downstream links as point queues on source and sink cells. (This assumption is relaxed in Section 5.) Let the number of vehicles entering the queues on 0 and  $\overleftarrow{|C|+1}$  at time  $t$  be given by  $d_0(t)$  and  $d_{\overleftarrow{|C|+1}}(t)$ . Then the queues of vehicles waiting to enter the link at time  $T$  are  $\sum_{t=0}^T (d_0(t) - y_{|C|}(t, L))$  and  $\sum_{t=0}^T (d_{\overleftarrow{|C|+1}}(t) - y_{\overleftarrow{0}}(t, L))$ , the differences between upstream demand and vehicles that entered the pair of links.

For the downstream ends, let cells  $|C|+1$  and  $\overleftarrow{0}$  be sink cells connected to cells  $|C|$  and  $\overleftarrow{1}$  with receiving flows are  $R_{|C|+1}(t)$  and  $R_{\overleftarrow{0}}(t)$ . Denote by  $y_0(t)$  flow entering cell 1 and by  $y_{\overleftarrow{|C|+1}}(t)$  flow entering cell  $\overleftarrow{|C|}$ . In the following IP, we consider the objective of maximizing the link throughput. Let  $L^*$  be an optimal solution to the following:

$$\begin{aligned}
& \max && Z(L) = \sum_{t=0}^T \alpha^t (y_{|C|}(t, L) + y_{\overleftarrow{\Gamma}}(t, L)) && (13) \\
& \text{s.t.} && y_i(t, L) = \min \{S_i(t, L), R_{i+1}(t, L)\} && \forall i, \overleftarrow{i} \in C, \forall t \in [0, T] \\
& && S_i(t, L) = \min \{n_i(t), QL_i(t)\} && \forall i, \overleftarrow{i} \in C, \forall t \in [0, T] \\
& && R_i(t, L) = \min \left\{ QL_i(t), \frac{w}{v} (NL_i(t) - n_i(t)) \right\} && \forall i, \overleftarrow{i} \in C, \forall t \in [0, T] \\
& && L_i(t) + L_{\overleftarrow{i}}(t) \leq \ell && \forall i \in C, \forall t \in [0, T] \\
& && |L_i(t) - L_i(t+1)| \leq 1 && \forall i \in C, \forall t \in [0, T] \\
& && |L_i(t) - L_{i+1}(t+1)| \leq 1 && \forall i \in C, \forall t \in [0, T] \\
& && L_i(t) \in \mathbb{Z}_+ && \forall i \in C, \forall t \in [0, T] \\
& && y_i(t, L) \geq 0 && \forall i, \overleftarrow{i} \in C, \forall t \in [0, T]
\end{aligned}$$

where  $\alpha \in (0, 1]$  is a discount factor to discourage delayed throughput.  $\alpha < 1$  penalizes delaying throughput to later time steps.  $\alpha < 1$  is necessary for analyses for which  $T \rightarrow \infty$ , as  $\alpha = 1$  would result in  $Z(L) \rightarrow \infty$  as  $T \rightarrow \infty$ .

Cell transition flows and vehicle movement may be specified for the single link under consideration because it is assumed that vehicles will move forward if possible. However, if multiple links were to be considered, the IP would have to include vehicle route choice.

$Z(L)$  as defined in the IP (13) maximizes discounted flow through the single link under consideration. This IP does not directly apply to traffic networks because of queue spillback. However, the DLR policy problem for a single link is sufficiently complex to require heuristics when used with DTA. Solving the IP for a network would introduce additional complexity in the form of route choice and intersection conflicts. Therefore, we restrict our attention to flow on a single link. In Section 6.3, we show that the single link heuristic yields significant improvements for a city network.

**Proposition 2.** *The IP (13) has at least one feasible solution if for all cells  $i$ ,  $L_i(0) + L_{\overleftarrow{i}}(0) \leq \ell$ ,  $|L_{i+1}(0) - L_i(0)| \leq 1$ ,  $|L_{\overleftarrow{i+1}}(0) - L_{\overleftarrow{i}}(0)| \leq 1$ ,  $NL_i(0) \geq n_i(0)$ , and  $NL_{\overleftarrow{i}}(0) \geq n_{\overleftarrow{i}}(0)$ .*

*Proof.* From Proposition 1, there exists a solution satisfying the DLR constraints on the number of lanes at each cell-time for  $t \geq 0$ . The feasibility of flow propagation constraints follows from CTM.  $\square$

Although solving the IP (13) yields the optimal DLR policy for a single link, it is not sufficient for network analyses. The single-link model does not consider queue spillback effects because demand waiting to enter the link is modeled as point queues. Furthermore, solving this IP for every link in a city network is not tractable, especially when UE route choice is taken into consideration. Therefore, the remainder of this section develops structure and intuition about the single-link IP. This structure is used to construct an effective heuristic in Section 5.




Figure 3: Flow through a single cell




Figure 4: Flow between a pair of cells

### 4.3 Bottlenecks

This section further explores the creation of bottlenecks on the links by allowing  $y_0(t) + y_{|C|+1}^{\leftarrow}(t) > Q\ell$  demand to enter in one time step. As seen in Section 4.1, bottlenecks can adversely affect the objective of maximizing total discounted flow through the link. In Proposition 3, we prove that creating a bottleneck is not necessary for optimality.

The total flow between any pair of parallel opposing cells is limited to  $Q\ell$  because only  $\ell$  lanes are available at a single cell (see Figures 3 and 4). Specifically, for all  $L \in \mathcal{L}_T$  and  $i \in C$ ,  $y_i(t, L) + y_{\bar{i}}^{\leftarrow}(t, L) \leq Q\ell$  and  $y_i(t, L) + y_{\bar{i+1}}^{\leftarrow}(t, L) \leq Q\ell$ . This yields a general result about bottlenecks:

**Proposition 3.** *Suppose that there exists an  $i \in C$ ,  $j \geq i$ ,  $t$  such that  $y_i(t, L^*) + y_{\bar{j}}^{\leftarrow}(t, L^*) \geq Q\ell$ . Then there exists an  $L' \in \mathcal{L}_T$  with  $Z(L') \geq Z(L^*)$  and  $y_i(t, L') + y_{\bar{j}}^{\leftarrow}(t, L') \leq Q\ell$ .*

*Proof.* By induction on  $j - i$ . The proof is split into two cases: whether the difference between  $i$  and  $j$  is even or odd.

*Basis:*  $j = i$ : because  $y_i(t, L) + y_{\bar{i}}^{\leftarrow}(t, L) \leq Q\ell$ , flow through cell  $i$  is limited to  $Q\ell$ , thereby limiting the future reward.  $j = i + 1$ : because  $y_i(t, L) + y_{\bar{i+1}}^{\leftarrow}(t, L) \leq Q\ell$ , flow through cell  $i$  is limited to  $Q\ell$ , thereby limiting the future reward.

*Inductive step:* Suppose  $j - i = n + 1$  (with  $n + 1 \geq 2$ ). If  $y_i(t, L^*) + y_{\bar{j}}^{\leftarrow}(t, L^*) > Q\ell$ , then such an  $L'$  exists by the inductive hypothesis applied at  $i + 1$ ,  $\bar{j} - 1$ ,  $t + 1$  (with  $(j - 1) - (i + 1) = n - 1$ ).  $\square$

Proposition 3 notes that if in  $L^*$ , two cells  $i$  and  $j$  at some time  $t$  have  $y_i(t, L^*) + y_{\bar{j}}^{\leftarrow}(t, L^*) > Q\ell$ , then some alternate policy  $L'$  with  $y_i(t, L') + y_{\bar{j}}^{\leftarrow}(t, L') \leq Q\ell$  is also optimal.

Note that this applies for flow entering in opposite directions at the time step, or for flow entering and flow already on the link. This allows restrictions to be placed on the solution. For instance, in a pair of links with 4 lanes total, if it is optimal to assign 3 lanes to one direction in one time step and entering flow exceeds  $2Q$ , then it is optimal to assign 3 lanes to succeeding cells in succeeding time steps to allow that flow to reach the end of the link.

#### 4.4 Partial lane reversal

A major modeling decision in the above formulation is deciding lane direction at the cell level, as opposed to the entire link. From Proposition 3,  $L_1(t) + L_{\overleftarrow{|C|}}(t) > \ell$  is not necessary for optimality. However, the opposite, where  $L_1(t) + L_{\overleftarrow{|C|}}(t) < \ell$ , could be beneficial to add additional turning lanes for exiting vehicles. To prevent queue spillback for one turning movement from interfering with another until vehicles exit, additional turning lanes longer than one cell could also improve flow. Although lane reversals to improve short-term flow at the end of the link may not be optimal in the long term, a discount factor of  $\alpha < 1$  encourages giving preference to exiting flow due to the total discount of at least  $\alpha^{|C|}$  from the minimum time required to traverse the link. Proposition 4 demonstrates that under certain conditions, a partial lane reversal on cell  $|C|$  will improve the total discounted flow through the link. Symmetric conditions apply to cell  $\overleftarrow{|C|}$ . These conditions are likely to occur at some time step for many networks.

**Proposition 4.** *If*

$$(i) \quad L_{|C|}^*(t) \leq L_{|C|}^*(t-1)$$

$$(ii) \quad \sum_{t'=0}^t d_{\overleftarrow{|C|+1}}(t) - \sum_{t'=0}^{t-1} y_{\overleftarrow{|C|+1}}(t, L^*) \leq Q \left( L_{\overleftarrow{|C|}}^*(t) - 1 \right)$$

$$(iii) \quad n_{\overleftarrow{|C|}}(t) \leq Q \left( L_{\overleftarrow{|C|}}^*(t) - 1 \right)$$

$$(iv) \quad n_{|C|}(t, L^*) \geq R_{|C|+1}(t) > QL_{|C|}^*(t)$$

$$(v) \quad \alpha < 1$$

*then there exists an  $L' \in \mathcal{L}_T$  with  $Z(L') > Z(L^*)$ .*

*Proof.* Construct  $L'$  as follows:  $L'_i(t) = L_i^*(t)$  and  $L'_{\overleftarrow{i}}(t) = L_{\overleftarrow{i}}^*(t)$  for all  $i \in C$  and  $t' \in [0, t]$ , except that  $L'_{\overleftarrow{|C|}}(t) = L_{\overleftarrow{|C|}}^*(t) - 1$  and  $L'_{|C|}(t) = L_{|C|}^*(t) + 1$ .  $L'$  is feasible because of condition

(i). Because of conditions (ii) and (iii) removing a lane from  $\overleftarrow{|C|}$  does not restrict flow. Then for all  $i \in C$  and  $t' \in [0, t]$ ,  $y_i(t', L') \geq y_i(t', L^*)$  and  $y_{\overleftarrow{i}}(t', L') \geq y_{\overleftarrow{i}}(t', L^*)$ . Furthermore, because of condition (iv),  $y_{|C|}(t, L') > y_{|C|}(t, L^*)$ . From condition (v),  $Z(L') > Z(L^*)$ .  $\square$

Condition (i) in Proposition 4 ensures feasibility of adding a lane to  $|C|$ . Conditions (ii) and (iii) state that the numbers of vehicles in cells  $\overleftarrow{|C|+1}$  and  $\overleftarrow{|C|}$  are sufficiently small that removing a lane from  $\overleftarrow{|C|}$  will not obstruct flow. Finally, condition (iv) states that the number of vehicles in  $|C|$  and the receiving flow on  $|C|+1$  are greater than the capacity allocation from  $L^*$ , and thus adding a lane to  $|C|$  will result in  $y_{|C|}(t, L') > y_{|C|}(t, L^*)$  when moving flow later is discounted in accordance with condition (v). Condition (v) is necessary because it rewards reducing delays, and adding a temporary extra turning lane is designed to reduce delays. With  $\alpha = 1$ , there might be no difference in objective from providing an extra turning lane instead of forcing some vehicles to wait until the next time step.

## 4.5 Stability

Because the objective function in the IP (13) maximizes discounted flow through the link, the optimal solution without a discount has a superstable property: if any policy prevents queues from growing to infinity as  $T \rightarrow \infty$ , then  $L^*$  is such a policy. First, we bound the queue lengths when  $L^*$  is used. Let  $\mathcal{d}(L)$  be the sum of the queue lengths at the end of the time horizon,  $T$ , for policy  $L \in \mathcal{L}_T$ . Then

$$\mathcal{d}(L) = \sum_{t=0}^T \left( (d_0(t) - y_{|C|}(t, L)) + \left( d_{\overleftarrow{|C|+1}}(t) - y_{\overleftarrow{1}}(t, L) \right) \right) \quad (14)$$

**Proposition 5.** *Let  $T \geq 0$ ,  $L \in \mathcal{L}_T$ , and  $\alpha = 1$ . Then  $\mathcal{d}(L) \geq \mathcal{d}(L^*)$ .*

*Proof.* From the objective function with  $\alpha = 1$ ,

$$\sum_{t=0}^T \left( (y_{|C|}(t, L^*)) + (y_{\overleftarrow{1}}(t, L^*)) \right) \geq \sum_{t=0}^T \left( (y_{|C|}(t, L)) + (y_{\overleftarrow{1}}(t, L)) \right) \quad (15)$$

Therefore

$$\begin{aligned} \mathcal{d}(L^*) &= \sum_{t=0}^T \left( (d_0(t) - y_{|C|}(t, L^*)) + \left( d_{\overleftarrow{|C|+1}}(t) - y_{\overleftarrow{1}}(t, L^*) \right) \right) \\ &\leq \sum_{t=0}^T \left( (d_0(t) - y_{|C|}(t, L)) + \left( d_{\overleftarrow{|C|+1}}(t) - y_{\overleftarrow{1}}(t, L) \right) \right) \\ &= \mathcal{d}(L) \end{aligned} \quad (16)$$

□

Denote by  $(L_T) = (L_T : L_T \in \mathcal{L}_T, T \in \mathbb{Z}_+)$  a sequence of feasible policies where every  $T \in \mathbb{Z}_+$  is mapped to a policy  $L_T \in \mathcal{L}_T$ . Similarly, denote by  $(L_T^*)$  a sequence of optimal policies to the IP (13). For any sequence of policies  $(L_T)$ , the resulting remaining queue lengths also form a sequence  $(\mathcal{d}(L_T))$ . Obviously,  $(\mathcal{d}(L_T))$  is bounded below as  $\mathcal{d}(L_T) \geq 0$

for any  $L_T \in \mathcal{L}_T$ . However  $(\mathcal{d}(L_T))$  may not be bounded above (and if it is bounded, the sequence may not converge). Nevertheless, we can use such sequences to establish the superstability of  $(L_T^*)$ .

**Proposition 6.** *Let  $\alpha = 1$ , and suppose that there exists a sequence of feasible policies  $(L_T)$  and a  $\zeta \in \mathbb{R}_+$  such that  $(\mathcal{d}(L_T))$  is bounded by  $\zeta$  (i.e. for all  $T \in \mathbb{Z}_+$ ,  $\mathcal{d}(L_T) \leq \zeta$ ). Then  $(\mathcal{d}(L_T^*))$  is also bounded by  $\zeta$ .*

*Proof.* For any  $T \in \mathbb{Z}_+$ ,  $\mathcal{d}(L_T^*) \leq \mathcal{d}(L_T)$  by Proposition 5. Since  $\mathcal{d}(L_T) \leq \zeta$ ,  $\mathcal{d}(L_T^*) \leq \zeta$ .  $\square$

Proposition 6 states the superstable property: if some sequence of feasible policies  $(L_T)$  results in bounded queue lengths, then  $(L_T^*)$  also has bounded queue lengths. However, these stability results require that  $\alpha = 1$ , i.e. that delaying exiting flow has no effect on the objective, as long as flow exits before  $T$ . This is due to the relationship between the queue length and the objective function. When a discount is used, equation (15) becomes

$$\sum_{t=0}^T ((\alpha^t y_{|C|}(t, L^*)) + (\alpha^t y_{\overleftarrow{1}}(t, L^*))) \geq \sum_{t=0}^T ((\alpha^t y_{|C|}(t, L)) + (\alpha^t y_{\overleftarrow{1}}(t, L))) \quad (17)$$

which only yields

$$\begin{aligned} & \sum_{t=0}^T \left( (d_{-1}(t) - \alpha^t y_{|C|}(t, L^*)) + \left( d_{\overleftarrow{|C|+1}}(t) - \alpha^t y_{\overleftarrow{1}}(t, L^*) \right) \right) \\ & \leq \sum_{t=0}^T \left( (d_{-1}(t) - \alpha^t y_{|C|}(t, L)) + \left( d_{\overleftarrow{|C|+1}}(t) - \alpha^t y_{\overleftarrow{1}}(t, L) \right) \right) \end{aligned} \quad (18)$$

Equation (18) shows how a policy  $L^*$  for  $\alpha < 1$  may not be optimal for bounding queue length. As a counterexample, consider a scenario in which the policy may shift lanes to cells 0 through  $|C|$  to allow more vehicles on the link to exit, or shift lanes to cells  $\overleftarrow{|C|}$  through  $\overleftarrow{0}$  to allow more queued vehicles to enter (and exit  $|C|$  time steps later). For sufficiently small  $\alpha$ , the optimal policy will prioritize vehicles already on the link because they can exit sooner, although this may result in a longer queue for entering cell  $\overleftarrow{|C|}$  at the end of the time horizon.

Although  $\alpha = 1$  is necessary for superstability to hold,  $\alpha < 1$  does not necessarily prevent the optimal policy from bounding queues for some demand scenarios. However,  $L^*$  cannot be guaranteed to bound queues if  $\alpha < 1$ . The choice of discount factor is similar to the capacity-delay tradeoff for traffic signals, where longer cycle lengths increase both capacity and delay. As  $\alpha$  increases, the optimal policy prioritizes capacity more than delay.  $\alpha = 1$  maximizes capacity but also removes any penalty for delaying vehicles. On the other hand, as discussed in Section 4.4,  $\alpha < 1$  is a necessary condition for adding extra turning lanes to increase the objective function in some scenarios.

The stability discussion also demonstrates some weaknesses of the IP approach. Besides requiring perfect information about demand, the IP also is solved for a fixed, finite time

horizon. The solution results in a policy optimized for a specific demand scenario. Because it is an IP, minor changes to the demand could result in major changes to the optimal policy. In the next section, we study DLR with stochastic demand as a Markov decision process. The resulting heuristic policy is more robust and tractable than the solution to this IP.

## 5 Stochastic demand

Although perfect information about demand yields ideal scenarios and corresponding theoretical results, in reality acquiring perfect information for arbitrary time horizons (such as the entire AM peak) requires knowledge of both vehicle route choice and departure times. Changes in either would potentially require solving the entire model again for some subinterval of time. Therefore, developing a DLR policy for stochastic supply and demand is also valuable. From the perspective of the link manager at time  $t$ , we assume that the change in demand  $d_0(t)$ ,  $d_{|C|+1}^{\leftarrow}(t)$  and supply  $R_{|C|+1}(t)$ ,  $R_{\overline{0}}^{\leftarrow}(t)$  for the next time step are known, but future demand and supply are given by stochastic processes. In general, upstream sending flow at  $t + 1$  is not independent of upstream sending flow at  $t$  because vehicles that do not enter at  $t$  will wait for the next time step. Similarly, if downstream receiving flows are limited by congestion at time  $t$ , there is a higher probability they will be limited by congestion at time  $t + 1$ . Since all vehicles are in communication with the link manager, we assume that for all  $i \in C$ ,  $n_i(t)$  and  $n_i^{\leftarrow}(t)$  are deterministic. Therefore, we consider the following infinite-horizon Markov decision process (MDP) with state space  $\mathcal{S}$ , control space  $\mathcal{U}$ , and one-step rewards  $\mathbf{g}(t)$ :

- The state at time  $t$  is the cell occupancies and number of lanes. Therefore, the state space is

$$\mathcal{S} = [0, N\ell]^{2|C|} \times \mathbb{Z}_+^2 \times [0, \ell]^{2|C|} \quad (19)$$

The integer  $2|C|$ -vectors of  $[0, N\ell]^{2|C|}$  are the possible combinations of cell occupancies because  $N\ell$  is the maximum occupancy of any single cell, and there are  $|C|$  cells in each direction.  $\mathbb{Z}_+^2$  is the possible lengths of the queues of vehicles waiting to enter the links. The integer  $2|C|$ -vectors of  $[0, \ell]^{2|C|}$  are the possible lane configurations.

- The control is how many lanes are assigned to each cell. Therefore,

$$\mathcal{U} \subset [0, \ell]^{2|C|} \quad (20)$$

The control space is limited by constraints (7) through (12) to ensure that vehicles do not change lanes more than once per time step, and that each cell has enough lanes that vehicles in the cell have sufficient physical space.

- The one-step rewards are given by the objective function to the IP:

$$\mathbf{g}(t) = y_{|C|}(t, L) + y_{\overline{1}}(t, L) \quad (21)$$

where transition flows  $y_{|C|}(t, L)$  and  $y_{\overline{1}}(t, L)$  are determined by equation (3).

- The state transitions are determined by entering demand and transition flows. Entering demand is  $d_0(t)$  and  $d_{|C|+1}^{\leftarrow}(t)$ . Transition flows are described by equation (3) for CTM. The transition flows are affected by the number of lanes assigned to each cell.
- The objective is to find a policy  $L^*$  of lane assignments that maximizes the long-run expected reward.

With a countable state space and finite action space, the MDP has an optimal stationary policy. Unfortunately, solving this MDP is fairly difficult. Due to the simulation-based CTM state, solving it analytically encounters similar issues to solving CTM DTA analytically. Computational methods for solving MDPs, based on dynamic programming, are polynomial in the state space. However, the state space is computationally intractable due to the curse of dimensionality. A typical 0.5 mile, 4 lane pair of links with free flow speed 30 miles per hour and jam density 120 vehicles per mile has 10 cells in each direction, and each cell could contain up to 24 vehicles. This alone results in a state space of  $4.01 \times 10^{27}$  elements. Choosing out of 5 possible lane configurations (0 through 4) per cell in one direction results in a further orthogonal  $9.8 \times 10^6$  possibilities.

Based on this complexity, it would be ideal to derive theoretical results for the MDP similar to the analyses in Section 4. Propositions 2 and 5 can be extended to the MDP with similar proofs. However, Proposition 3 does not have a direct counterpart in the stochastic case. Consider a pair of links with 2 lanes in each direction, 900 vph capacity per lane, expected 900 vph demand in each direction, time step of  $\Delta t$ , and  $\alpha = 1$ . If, due to randomness,  $3600\Delta t$  demand in direction 1 appears at time  $t$ , based on expected future demand assigning 3 lanes to direction 1 at time  $t$  is a maximum throughput policy. If at times  $t+1$  through  $t+5$ ,  $3600\Delta t$  demand also appears in direction 2, to maximize throughput a bottleneck on the link should be created in direction 1 as the potential reward in direction 2 is greater.

## 5.1 Heuristic algorithm

Therefore, instead of attempting to solve this MDP computationally, we use the structure developed in Section 4 to inspire a saturation-based heuristic. Hausknecht et al. [14] briefly discuss a theorem on DLR with respect to saturation, but it assumes stationary, constant flow and does not include downstream receiving flow limitations. We extend it to the stochastic demand, CTM model to determine expected saturation levels for two links  $[a, b]$  and  $[b, a]$ . At time  $t$ , we first determine the number of lanes per cell, then propagate flow.

To simplify the possible actions, we choose two modes of control. First, all but the last cell is assigned the same number of lanes, formally cells 1 through  $|C| - 1$  and cells  $|C|$  through  $\overleftarrow{2}$ . Although Proposition 3 may not hold in its most general sense, allowing more than  $Q\ell$  of flow to enter in one time step still cannot increase the reward. Furthermore, we add the restriction that each direction must always have at least one lane, even if no flow is presently using it. This prevents flow in one direction from being completely obstructed



due to high demand in the other direction. In most practical scenarios, it is unlikely for one direction to have completely zero demand.

### 5.1.1 Overall lane direction

Inspired by Theorem 1 of Hausknecht et al. [14], this heuristic estimates the difference between demand and capacity for each direction. If demand exceeds capacity in one direction, and the other direction has unused capacity, then it may be beneficial to reverse one lane. Since the number of lanes is integer, we choose to reverse a lane only if shifting  $Q$  capacity from one direction to the other is expected to improve flow. Formally, define  $\delta_\gamma(t)$  as the saturation estimation for direction  $\gamma \in \{1, 2\}$ , where the direction index is assigned arbitrarily.  $\delta_\gamma(t) > 0$  and  $\delta_\gamma(t) < 0$  indicate over- and under-saturation, respectively. To avoid confusion with  $L_i(t)$ , let  $l_\gamma(t)$  represent the number of lanes in direction  $\gamma$ . The initial condition is  $l_1(t-1) + l_2(t-1) = \ell$ . Set

$$\delta_1(t) = \min \left\{ \begin{array}{l} \sum_{1 \leq i \leq |C|} n_i(t) + \sum_{0 \leq t' \leq T} \mathbb{E}[S_{-1}(t+t')], \\ \sum_{0 \leq t' \leq T} \mathbb{E}[R_{|C|+1}(t+t')] \end{array} \right\} - Ql_1(t-1)T \quad (22)$$

and

$$\delta_2(t) = \min \left\{ \begin{array}{l} \sum_{1 \leq i \leq |C|} n_i^{\leftarrow}(t) + \sum_{0 \leq t' \leq T} \mathbb{E}[S_{|C|+1}^{\leftarrow}(t+t')], \\ \sum_{0 \leq t' \leq T} \mathbb{E}[R_{-1}^{\leftarrow}(t+t')] \end{array} \right\} - Ql_2(t-1)T \quad (23)$$

$T$  defines how far ahead into the future the heuristic considers when estimating saturation. A low value of  $T$  will not allow all vehicles to exit, and will result in the heuristic being highly reactive to specific realizations of supply and demand. Therefore we recommend  $T$  be at least the number of cells in the link. On the other hand, a high value of  $T$  might prevent the heuristic from reacting optimally to dynamic congestion.

The minimum term in  $\delta_\gamma(t)$  is the possible expected throughput of the link, accounting for expected upstream sending flow and constraints of expected downstream receiving flow. The subtracted term is the maximum throughput possible over  $T$  using lanes  $l_\gamma(t-1)$ .

If  $\delta_1(t) > (Q + \delta_2(t))^+$  and lane constraints (7) through (12) allow it, set  $l_1(t) = l_1(t-1) + 1$ , where  $(\cdot)^+ = \max\{0, \cdot\}$ . Similarly, if  $\delta_2(t) > (Q + \delta_1(t))^+$  and lane constraints (7) through (12) allow it, set  $l_2(t) = l_2(t-1) + 1$ . These two conditions cannot both be true because if  $\delta_1(t) > (Q + \delta_2(t))^+$  then  $\delta_2(t) < \delta_1(t)$ , and vice versa.

### 5.1.2 Additional turning bays

In addition, the last cell can be assigned extra turning lanes to allow more flow to exit, based on Proposition 4. We refer to the number of lanes at the start and end cells of direction 1 as  $l_1^1(t)$  and  $l_1^{|C|}(t)$ , respectively. For direction 2, the number of lanes at the

start and end cells are  $l_2^{\overleftarrow{C|}}(t)$  and  $l_2^{\overleftarrow{1}}(t)$ , respectively. Initially, set  $l_1^{\overleftarrow{C|}}(t) = l_1^{\overleftarrow{C|}}(t) = l_1(t)$  and  $l_2^{\overleftarrow{C|}}(t) = l_2^{\overleftarrow{1}}(t) = l_2(t)$ . For direction 1, if  $l_1(t) \leq l_1(t-1)$  (to satisfy at most 1 additional lane per time step), set

$$\delta'_{11} = \min \{n_{|C|}(t), Q(l_1(t) + 1), R_{|C|+1}(t)\} - \min \{n_{|C|}(t), Ql_1(t), R_{|C|+1}(t)\} \quad (24)$$

and

$$\delta'_{21} = \min \{S_{|C|+1}^{\overleftarrow{C|}}(t), Ql_2(t)\} - \min \{S_{|C|+1}^{\overleftarrow{C|}}(t), Q(l_2(t) - 1)\} \quad (25)$$

$\delta'_{11}$  is the difference in flow for the cases of  $l_1^{\overleftarrow{C|}}(t) = l_1(t) + 1$  and  $l_1^{\overleftarrow{C|}}(t) = l_1(t)$ , and  $\delta'_{21}$  is similarly the difference in flow for the cases of  $l_2^{\overleftarrow{C|}}(t) = l_2(t) - 1$  and  $l_2^{\overleftarrow{C|}}(t) = l_2(t)$ . If the improvement is sufficient, i.e. if  $\delta'_{11} > 0$  and  $\delta'_{11} > \delta'_{21}$ , then set  $l_1^{\overleftarrow{C|}}(t) = l_1(t) + 1$  and  $l_2^{\overleftarrow{C|}}(t) = l_2(t) - 1$ . An analogous operation is performed for direction 2.

### 5.1.3 Simulation algorithm

This heuristic is part of the `Simulate` procedure in Algorithm 1. Every time step, we use the above heuristic to determine the number of lanes in each direction for each pair of parallel cells using equations (22) through (25). Then, we calculate transition flows using equations (3) through (5), and propagate flow according to equation (2). We repeat this each time step until all vehicles have exited. The simulation is illustrated in Figure 5. DLR adds the step of deciding lane directions before propagating flow. The remainder of the simulation is the same in conventional CTM.

## 5.2 Demonstration

To demonstrate the effectiveness of the above heuristic, we performed a suite of tests on a single pair of links with varying combinations of stationary demand. Each link was 0.4 miles long, had 2 lanes, 1200 vph capacity, 30 mph free flow speed, 15 mph backwards wave speed, and arrivals were Poisson each time step based on demand. A time step of 6 seconds was used for CTM (as used by [29]), and the lookahead parameter  $T$  was set to 40 time steps. Due to randomness in the demand, each scenario was simulated 100 times for 1 hour, and average results are presented. Figure 6 graphs the difference in throughput between DLR and a fixed lane configuration of 2 lanes in each direction.

Figure 6 demonstrates that in asymmetric demand scenarios where the total demand is less than the total link capacity including lanes in both directions, the DLR heuristic tends to improve over the fixed base lane configuration. Although this is not surprising, these results are important for several reasons.

First, although contraflow lanes would achieve similar results in some of the demand scenarios considered, they are difficult to implement due to human drivers. When AV intersection controllers are in use, DLR may be implemented on *every* link, and this demonstrates some of the benefits of doing so.

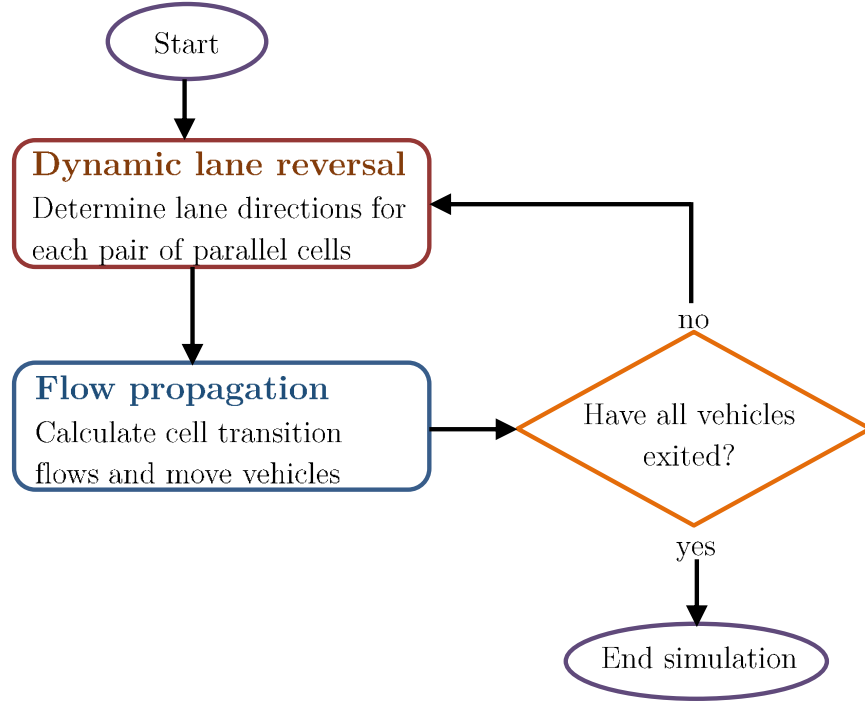


Figure 5: CTM simulation with DLR

Second, this heuristic responds particularly well to scenarios in which one direction is slightly oversaturated and the other is slightly undersaturated, but reversing a lane would not improve the total flow. For example, consider a link with 4 lanes, with 1200vph capacity per lane, and with demand of 2700vph in one direction and 1500vph in the other. With 2 lanes in each direction, 300vph of demand will not be served, but this is also true for a 3–1 lane configuration. DLR allows frequent changing between 2–2 and 3–1 configurations, allowing that additional 300vph to use the link. The proposed heuristic switches automatically based on the queues of vehicles waiting to enter.

Finally, this DLR heuristic was not observed to perform significantly worse than a fixed lane configuration. In several demand scenarios the average throughput of DLR was slightly worse than that of fixed lanes. However, the decrease was two orders of magnitude less than the potential improvement. Overall, these results suggest that while this heuristic may not be the optimal policy for DLR, in many cases it improves over a fixed lane configuration, and it will probably not be much worse. Therefore, this heuristic is worth consideration on larger networks.

## 6 Dynamic lane reversal on networks

Although the heuristic developed in Section 5 proved effective on single link bottlenecks with stationary demand, the ultimate goal is to apply DLR to larger networks with the additional

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65




Figure 6: Change in total throughput from DLR heuristic

1  
2  
3  
4 variables of intersection constraints and UE routing. Therefore, we incorporate the heuristic  
5 into DTA, presented in Figure 1.  
6  
7

## 8 **6.1 Determining expected sending and receiving flows** 9

10 The saturation definitions in equations (22) and (23) use expected demand, which depends  
11 on traveler route choice. To determine this endogenously, each link stores expected sending  
12 and receiving flows per assignment interval (AST). In DTA, ASTs are used to reduce the  
13 computational complexity of routing demand. Typically, each iteration a single shortest  
14 path is found for every origin-destination-AST (ODT) tuple. For DLR, we also use ASTs as  
15 the aggregation level for expected sending and receiving flows because it corresponds to the  
16 path assignment aggregation. Because changes in route choice affect expected sending and  
17 receiving flows, each iteration, the expected values per link are updated based on average  
18 observations from the simulation. Average upstream sending flows for link  $[a, b]$  are calculated  
19 as the average number of vehicles wanting to enter  $[a, b]$ . (For general networks, this requires  
20 disaggregation of sending flows of upstream links by destination link). Receiving flows are  
21 more difficult to calculate because of intersection constraints on crossing flow. Instead, we  
22 used the average exiting flow as the expected receiving flow for the heuristic. For congested  
23 links this is an accurate measure because exiting flow is bounded by receiving flows. For  
24 uncongested links, DLR is not necessary anyways.  
25  
26  
27  
28  
29  
30

## 31 **6.2 Algorithm description** 32

33 The first step of the algorithm is to determine which links can be paired for DLR. We paired  
34 together any links  $[a, b]$  and  $[b, a]$  with the same length and free flow speed. In practice, some  
35 pairs of opposite and parallel links are separated by a median or divider. We assume that for  
36 AVs, such dividers are not necessary for safety purposes. We also did not have specific data  
37 on which pairs of links had dividers or not. However, if such dividers are used in practice,  
38 they would prevent DLR from being applied.  
39  
40

41 For the first iteration, expected sending and receiving flows are not known, so a partial  
42 demand initialization [30] is used to both improve convergence and provide initial inputs to  
43 DLR. DLR is embedded in the simulation step of DTA, as illustrated by Figure 5. Every  
44 time step, lane assignments are chosen using the heuristic in Section 5. After each simula-  
45 tion, expected sending and receiving flows are recorded. This definition of DLR uses values  
46 from only the last iteration. However, because the number of vehicles moved continuously  
47 decreases through MSA, the change in the DLR policy gradually decreases as well.  
48  
49  
50

## 51 **6.3 City network results** 52

53 To demonstrate the tractability and effectiveness of our proposed heuristic, we tested it on  
54 the downtown Austin network, shown in Figure 7, which has 62836 trips over 2 hours, 171  
55 zones, 546 intersections, and 1247 links in the AM peak. CTM was used with a time step of  
56 6 seconds and an AST duration of 15 minutes. DLR was implemented on all pairs of parallel  
57  
58  
59  
60  
61  
62  
63  
64  
65

---

**Algorithm 1** Dynamic lane reversal in dynamic traffic assignment

---

```
1: procedure INITIALIZATION
2:   for each link  $[a, b]$  do
3:     if there exists a link  $[b, a]$  with the same free flow speed and length then
4:       Pair  $[a, b]$  and  $[b, a]$  together for DLR
5:     end if
6:   end for
7:   for  $m = 1$  to  $M$  do
8:     Add  $\frac{1}{M}$  of unassigned vehicles to the network
9:     PATH-GENERATION(1)
10:    SIMULATE
11:  end for
12: end procedure
13:
14: procedure METHOD OF SUCCESSIVE AVERAGES
15:   for  $m = 1$  to  $M$  do
16:     PATH-GENERATION( $\frac{1}{m}$ )
17:     SIMULATE
18:   end for
19: end procedure
20:
21: procedure PATH-GENERATION( $\lambda$ )
22:   for each ODT  $(r, s, t)$ , find shortest path  $\pi_{rst}^*$  do
23:     for each vehicle  $v$  traveling from  $r$  to  $s$  departing within  $t$  do
24:       Assign  $v$  to  $\pi_{rst}^*$  with probability  $\lambda$ 
25:     end for
26:   end for
27: end procedure
28:
29: procedure SIMULATE
30:   for  $t = 0$  to  $\infty$  do
31:     for each link  $[a, b]$  do
32:       if  $[a, b]$  is paired with  $[b, a]$  then
33:         Determine lane assignment for  $[a, b]$  and  $[b, a]$ 
34:       end if
35:       Propagate flow through  $[a, b]$ 
36:       Update expected sending and receiving flows
37:     end for
38:   end for
39: end procedure
```

---

Table 1: Total system travel time

Scenario	TSTT (hr)
Fixed lanes	8420.966
DLR heuristic	6588.828

and opposite direction links with the same speed and length. As much of the network is a downtown grid, DLR was implemented on most links in the network. Because DLR is most applicable when all vehicles are AVs, the reservation control model of Levin & Boyles [31] for DTA with first-come-first-serve priority was used for intersections.

The demand is completely deterministic. However, because route choice changes through the process of solving for dynamic user equilibrium, determining the demand for individual links in the network would require forward simulation. Due to the computational cost of simulating many possible lane direction scenarios, it is easier to model link-specific demand as a random variable.

To demonstrate our DLR heuristic, we solved DTA for two scenarios: current (fixed) lane configuration, and DLR. We then compared the travel times at UE for both scenarios. Both scenarios were solved to the same cost gap of 2% of total system travel time.

Convergence of DTA with DLR is demonstrated in Figure 8. The partial demand initialization resulted in a relatively small initial gap. Around 4%–5%, the cost gap percent oscillated, which was probably due in part to DLR. However, after iteration 9 the cost gap steadily decreased, suggesting it found a local equilibrium. With the addition of DLR, DTA required 8.16 minutes to solve on an Intel Xeon CPU at 3.07 GHz. This makes it tractable for study on large city networks.

Our heuristic was developed for a single link, and the results in Section 5.2 show its effectiveness. However, the network level introduces route choice and queue spillback, neither of which are considered in our analysis of DLR policy for single link flow. The results presented here could be further improved by including network effects.

Table 1 shows the total system travel time (TSTT) for both fixed lanes and when the DLR heuristic was used. The DLR heuristic resulted in an improvement of 21.8% over fixed lanes. This demonstrates the potential benefits of using DLR during peak hour demand. As this is the AM peak, most of the demand is headed towards the downtown region, shown in Figure 2. The extra capacity afforded by DLR helps alleviate the congestion caused by the asymmetric use of right-of-way. On average, distance traveled by the same vehicle was observed to decrease by 23.9% when DLR was used. This suggests that greater capacity on shorter distance routes increased their utility in UE routing.

Figure 9 shows the average improvement in travel time from DLR for vehicles at different departure times. Vehicles departing later receive the greatest benefit because those vehicles experience a more congested network, and DLR alleviates much of the congestion. Overall, these results demonstrate that the DLR heuristic is effective at improving efficiency in congested large city networks.








Figure 8: Convergence of dynamic lane reversal on downtown Austin




Figure 9: Average reduction in travel time at different assignment intervals

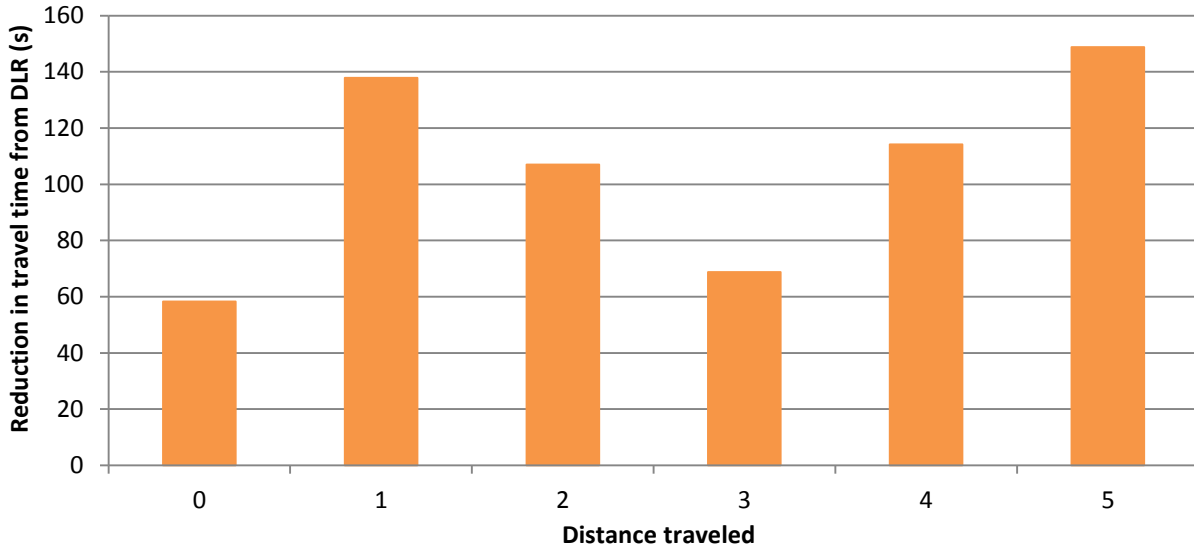


Figure 10: Average reduction in travel time from DLR with respect to vehicle miles traveled

This particular test network contains both freeways on the east and west boundaries and a detailed downtown region. (Some links in downtown are two-way, while others are one-way and do not have a counterpart for DLR). Vehicles traveling shorter distances are more likely to take arterials or the downtown grid, whereas vehicles traveling longer distances are more likely to take freeways and downtown roads due to the geometry of the network. Figure 10 demonstrates that vehicles traveling between 1-2 miles and 5+ miles in the fixed lane configuration experienced similar reductions in travel times. This suggests that this DLR heuristic is more effective for arterials than restricted access freeways because vehicles traveling longer distances have a greater potential for reductions in travel time. This is likely due to the limited number of lanes on exit ramps; DLR would not be able to add extra lane capacity to these ramps. As with the VISTA DTA simulator [29], the queueing model allows queues for these ramps to block entire cells. However, regardless of the distance, vehicles experienced average reductions in travel time, which suggests this heuristic consistently improves over the fixed lane configuration.

Figure 11 shows which DLR-enabled links actually have lane directions changed at 30 and 60 minutes into the demand. DLR was used on surprisingly few links considering the travel time reduction of 21.8%. The network was more congested after 60 minutes, so more links had reversed lanes. The links that DLR was used on were mainly arterials, and were mostly on the boundaries of the network. This allowed greater flow on the bridges in the south part of the city, and on other arterials, so vehicles could enter the downtown grid faster. The downtown grid did not have much DLR used because vehicles use the local streets there in both directions.

These results show that DLR can be significantly effective when used only on a few links. In terms of practical implementation, DLR could require significant infrastructure

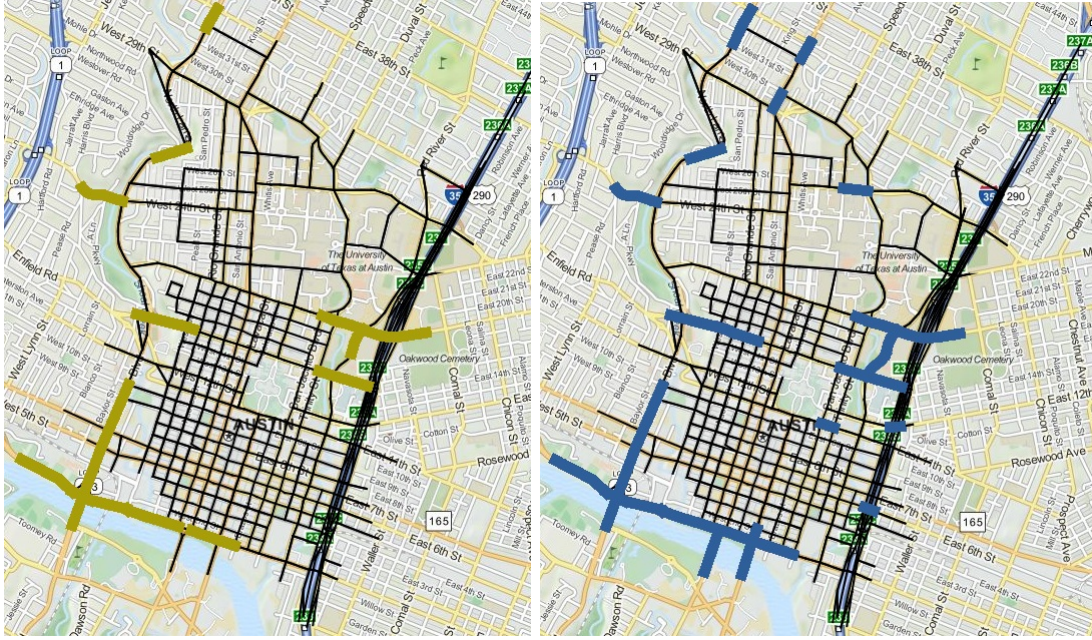


Figure 11: Links with DLR active around 30 minutes (left) and 60 minutes (right)

modifications in terms of installing lane manager hardware and modifying roadway geometry so that lane reversal is feasible. These modifications might only need to be made on a small set of carefully chosen links to achieve considerable benefits.

Of course, the fact that DLR was only used on such few links also indicates that greater improvement may be possible with a better DLR policy. The heuristic used here made lane direction decisions at an individual link level, and a policy considering network-level flows would likely result in further reductions in congestion. Therefore, these results warrant further study and optimization of DLR policies.

## 7 Conclusions

This paper developed a cell transmission model with variable number of lanes in space and time consistent with the hydrodynamic theory of traffic flow to model dynamic lane reversal. Focusing on a single link, the scenarios of known and stochastic demand were considered. When demand is known, we demonstrated that a solution algorithm should consider future demand and receiving flows and formulated dynamic lane reversal as an integer program. We derived theoretical results about the optimal solution(s), noting that using lane reversals to create a bottleneck on a link is never necessary for optimality when demand is known, and proving that the optimal solution will stabilize queues if they can be stabilized.

Because demand is often not known perfectly at arbitrary times in the future, we formulated the dynamic lane reversal problem with stochastic demand as a Markov decision

1  
2  
3  
4 process. The Markov decision process was analytically difficult to solve because it is built  
5 on a simulation-based dynamic traffic assignment model, and the curse of dimensionality led  
6 to computational intractability. Nevertheless, we developed a heuristic based on saturation  
7 estimates that was demonstrated to work well on a single bottleneck link at various com-  
8 binations of stationary demand. We then presented an algorithm for using the heuristic in  
9 dynamic traffic assignment, and tested it on a city network. This converged to an equilibrium  
10 and resulted in a 21.8% reduction in total system travel time.

11  
12  
13 While our proposed heuristic improves significantly over the current fixed lane configura-  
14 tion, it may not be the optimal dynamic lane reversal policy. Therefore better solutions  
15 to the single link and network-wide dynamic lane reversal problem could further improve  
16 traffic. Also, dynamic lane reversal for networks should account for UE routing behavior  
17 and queue spillback. In short, the model presented here was an initial foray into dynamic  
18 lane reversal in dynamic traffic assignment. Many possibilities for greater improvements and  
19 further research exist.  
20  
21  
22

## 23 Acknowledgements

24  
25 The authors gratefully acknowledge the support of the Data-Supported Transportation Op-  
26 erations & Planning Center, the NSF CAREER Program, Grant No. 1254921, and the  
27 Dwight D. Eisenhower Fellowship program.  
28  
29  
30

## 31 References

- 32  
33  
34 [1] K. Dresner and P. Stone, “Multiagent traffic management: A reservation-based intersec-  
35 tion control mechanism,” in *Proceedings of the Third International Joint Conference on*  
36 *Autonomous Agents and Multiagent Systems-Volume 2*, pp. 530–537, IEEE Computer  
37 Society, 2004.  
38  
39 [2] K. Dresner and P. Stone, “Traffic intersections of the future,” in *Proceedings of the*  
40 *National Conference on Artificial Intelligence*, vol. 21, p. 1593, Menlo Park, CA; Cam-  
41 bridge, MA; London; AAAI Press; MIT Press; 1999, 2006.  
42  
43 [3] D. Fajardo, T.-C. Au, S. Waller, P. Stone, and D. Yang, “Automated intersection  
44 control: Performance of future innovation versus current traffic signal control,” *Trans-*  
45 *portation Research Record: Journal of the Transportation Research Board*, no. 2259,  
46 pp. 223–232, 2011.  
47  
48 [4] Z. Li, M. Chitturi, D. Zheng, A. Bill, and D. Noyce, “Modeling reservation-based au-  
49 tonomous intersection control in vissim,” *Transportation Research Record: Journal of*  
50 *the Transportation Research Board*, no. 2381, pp. 81–90, 2013.  
51  
52 [5] M. W. Levin and S. D. Boyles, “Effects of autonomous vehicle ownership on trip, mode,  
53 and route choice,” in *Transportation Research Board 94th Annual Meeting*, no. 15-2147,  
54 2015.  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

- 1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65
- [6] X. M. Zhang, S. An, and B. L. Xie, "A cell-based regional evacuation model with contraflow lane deployment," in *Advanced Engineering Forum*, vol. 5, pp. 20–25, Trans Tech Publ, 2012.
  - [7] J. Wang, H. Wang, W. Zhang, W. Ip, and K. Furuta, "Evacuation planning based on the contraflow technique with consideration of evacuation priorities and traffic setup time," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 14, no. 1, pp. 480–485, 2013.
  - [8] V. Dixit and B. Wolshon, "Evacuation traffic dynamics," *Transportation Research Part C: Emerging Technologies*, vol. 49, pp. 114–125, 2014.
  - [9] W. Zhou, P. Livolsi, E. Miska, H. Zhang, J. Wu, and D. Yang, "An intelligent traffic responsive contraflow lane control system," in *Vehicle Navigation and Information Systems Conference, 1993., Proceedings of the IEEE-IEE*, pp. 174–181, IEEE, 1993.
  - [10] D. Xue and Z. Dong, "An intelligent contraflow control method for real-time optimal traffic scheduling using artificial neural network, fuzzy pattern recognition, and optimization," *Control Systems Technology, IEEE Transactions on*, vol. 8, no. 1, pp. 183–191, 2000.
  - [11] Q. Meng, H. L. Khoo, and R. L. Cheu, "Microscopic traffic simulation model-based optimization approach for the contraflow lane configuration problem," *Journal of Transportation Engineering*, vol. 134, no. 1, pp. 41–49, 2008.
  - [12] P.-D. D. D. Braess, "Über ein paradoxon aus der verkehrsplanung," *Unternehmensforschung*, vol. 12, no. 1, pp. 258–268, 1968.
  - [13] C. F. Daganzo, "Queue spillovers in transportation networks with a route choice," *Transportation Science*, vol. 32, no. 1, pp. 3–11, 1998.
  - [14] M. Hausknecht, T.-C. Au, P. Stone, D. Fajardo, and T. Waller, "Dynamic lane reversal in traffic management," in *Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on*, pp. 1929–1934, IEEE, 2011.
  - [15] K. M. Dresner and P. Stone, "Sharing the road: Autonomous vehicles meet human drivers.," in *IJCAI*, vol. 7, pp. 1263–1268, 2007.
  - [16] Y.-C. Chiu, J. Bottom, M. Mahut, A. Paz, R. Balakrishna, T. Waller, and J. Hicks, "Dynamic traffic assignment: A primer," *Transportation Research E-Circular*, no. E-C153, 2011.
  - [17] M. J. Lighthill and G. B. Whitham, "On kinematic waves. ii. a theory of traffic flow on long crowded roads," in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 229, pp. 317–345, The Royal Society, 1955.

- 1  
2  
3  
4 [18] P. I. Richards, “Shock waves on the highway,” *Operations research*, vol. 4, no. 1, pp. 42–  
5 51, 1956.  
6  
7  
8 [19] C. F. Daganzo, “The cell transmission model: A dynamic representation of highway traf-  
9 fic consistent with the hydrodynamic theory,” *Transportation Research Part B: Method-*  
10 *ological*, vol. 28, no. 4, pp. 269–287, 1994.  
11  
12 [20] C. F. Daganzo, “The cell transmission model, part ii: network traffic,” *Transportation*  
13 *Research Part B: Methodological*, vol. 29, no. 2, pp. 79–93, 1995.  
14  
15 [21] I. Yperman, S. Logghe, and B. Immers, “The link transmission model: An efficient  
16 implementation of the kinematic wave theory in traffic networks,” in *Proceedings of the*  
17 *10th EWGT Meeting*, 2005.  
18  
19 [22] P.-E. Mazaré, A. H. Dehwah, C. G. Claudel, and A. M. Bayen, “Analytical and grid-free  
20 solutions to the lighthill–whitham–richards traffic flow model,” *Transportation Research*  
21 *Part B: Methodological*, vol. 45, no. 10, pp. 1727–1748, 2011.  
22  
23 [23] G. Marsden, M. McDonald, and M. Brackstone, “Towards an understanding of adaptive  
24 cruise control,” *Transportation Research Part C: Emerging Technologies*, vol. 9, no. 1,  
25 pp. 33–51, 2001.  
26  
27 [24] B. Van Arem, C. J. Van Driel, and R. Visser, “The impact of cooperative adaptive  
28 cruise control on traffic-flow characteristics,” *Intelligent Transportation Systems, IEEE*  
29 *Transactions on*, vol. 7, no. 4, pp. 429–436, 2006.  
30  
31 [25] A. Kesting, M. Treiber, and D. Helbing, “Enhanced intelligent driver model to access  
32 the impact of driving strategies on traffic capacity,” *arXiv preprint arXiv:0912.3613*,  
33 2009.  
34  
35 [26] W. J. Schakel, B. Van Arem, and B. D. Netten, “Effects of cooperative adaptive cruise  
36 control on traffic flow stability,” in *Intelligent Transportation Systems (ITSC), 2010*  
37 *13th International IEEE Conference on*, pp. 759–764, IEEE, 2010.  
38  
39 [27] M. W. Levin and S. D. Boyles, “A multiclass cell transmission model for shared human  
40 and autonomous vehicle roads,” *Transportation Research Part C: Emerging Technolo-*  
41 *gies*, vol. 62, pp. 103–116, 2016.  
42  
43 [28] R. Courant, K. Friedrichs, and H. Lewy, “Über die partiellen differenzgleichungen der  
44 mathematischen physik,” *Mathematische Annalen*, vol. 100, no. 1, pp. 32–74, 1928.  
45  
46 [29] A. K. Ziliaskopoulos and S. T. Waller, “An internet-based geographic information sys-  
47 tem that integrates data, models and users for transportation applications,” *Transporta-*  
48 *tion Research Part C: Emerging Technologies*, vol. 8, no. 1, pp. 427–444, 2000.  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

- [30] M. W. Levin, M. Pool, T. Owens, N. R. Juri, and S. T. Waller, “Improving the convergence of simulation-based dynamic traffic assignment methodologies,” *Networks and Spatial Economics*, pp. 1–22, 2014.
- [31] M. W. Levin and S. D. Boyles, “Intersection auctions and reservation-based control in dynamic traffic assignment,” in *Transportation Research Board 94th Annual Meeting*, no. 15-2149, 2015.