Network-Wide Adaptive Tolling for Connected and Automated Vehicles

Guni Sharon[1], Michael W. Levin[2], Josiah P. Hanna[1], Tarun Rambha[3], Stephen D. Boyles[4] and Peter Stone[1]

[1]Computer Science, The University of Texas at Austin, Austin, Texas

[2]Civil, Environmental, and Geo- Engineering, University of Minnesota, Minneapolis, Minnesota

[3]Civil and Environmental Engineering, Cornell University, Ithaca, New-York

[4]Civil, Architectural and Environmental Engineering, The University of Texas at Austin, Austin, Texas

Abstract

This article proposes Δ-*tolling*, a simple adaptive pricing scheme which only requires travel time observations and two tuning parameters. These tolls are applied throughout a road network, and can be updated as frequently as travel time observations are made. Notably, Δ-*tolling* does not require any details of the traffic flow or travel demand models other than travel time observations, rendering it easy to apply in real-time. The flexibility of this tolling scheme is demonstrated in three specific traffic modeling contexts with varying traffic flow and user behavior assumptions: a day-to-day pricing model using static network equilibrium with link delay functions; a within-day adaptive pricing model using the cell transmission model and dynamic routing of vehicles; and a microsimulation of reservation-based intersection control for connected and autonomous vehicles with myopic routing. In all cases, Δ-*tolling* produces significant benefits over the no-toll case, measured in terms of average travel time and social welfare, while only requiring two parameters to be tuned. Some optimality results are also given for the special case of the static network equilibrium model with BPR-style delay functions.

1 Introduction

Road pricing as a tool for congestion management has a long history in the transportation world, being first suggested nearly a century ago (Pigou, 1920) as a method for internalizing congestion costs. Recent advances in connected and automated vehicle (CAV) technology offer unprecedented flexibility and scope for implementing these tolls. In principle, tolls can be charged on many or all network links, and changed frequently in response to real-time observations of traffic conditions. Toll values and traffic conditions can then be communicated to vehicles which instantly change routes in response, with minimal to no intervention needed on behalf of drivers, who might only indicate some measure of the trip urgency or other proxy for value of time before departing. Even before CAV technology reaches full penetration, communication capabilities and automated route selection software (as can be found on modern cell phones) may be sufficient to implement such a scheme.

Developing an optimal strategy for calculating tolls in such an environment appears difficult. There is a plethora of traffic flow models in the literature, and each of them suggests a different optimization problem. A similar issue exists for modeling route choice behavior. For these reasons, this article investigates a model-free adaptive[1] tolling procedure, which only requires travel time observations on links and which calculates

[1]In this article, we use *adaptive* to mean both time-varying and reactive to observed conditions. By this definition, time-of-day pricing according to a predetermined schedule is not adaptive.

tolls extremely quickly. We term this strategy Δ-*tolling*. In particular, Δ-*tolling* does not require knowing travel demand or the traffic flow model which maps route choices to congestion. In this procedure, we sacrifice provable optimality for generality, tractability, and decentralizability.

The past literature on adaptive tolling generally takes one of the following forms: optimization-based approaches which generally assume demand and the traffic flow model are known; or corridor pricing approaches which only toll a small set of links and are only interested in flow on a small set of links (as with managed lanes). This paper suggests a different approach that allows computing network-wide tolls in real-time, with the intent of reducing average travel time, and without requiring perfect knowledge of travel demand or how congestion forms.

The Δ-*tolling* scheme is exceptionally simple, and involves charging a toll on each link proportional to its delay (the difference between observed and free-flow travel times), updating this toll periodically. There are only two parameters to tune: the proportionality constant, and a smoothing parameter used to damp transient spikes in toll values. Furthermore, being model-free, it is readily adaptable to application and testing in multiple environments.

Despite the simplicity of the Δ-*tolling* scheme itself, practical implementation requires answering many interesting and non-trivial questions regarding its effectiveness in various settings, and how the tuning parameters should be chosen. In this article, we describe numerical tests of Δ-*tolling* with three widely different traffic models: Model A uses link performance functions which are commonly used in static traffic assignment, is deterministic, and travelers are fully rational; Model B uses the cell transmission model, commonly used in regional dynamic traffic assignment, and travelers update their routes by recalculating shortest paths at each node; and Model C uses microsimulation, commonly used for detailed corridor studies, with the same behavioral assumption. These models focus on how travelers' route choice is affected by Δ-*tolling*. Implementing Δ-*tolling* within these models achieved significant reductions in travel times. Furthermore, for the case of link performance functions (Model A), we are able to show that the Δ-*tolling* scheme is in fact optimal if the model parameters are correctly chosen. The effectiveness of Δ-*tolling* across three distinct traffic models is encouraging, given the complexity of traffic flow which renders all such models approximate.

The remainder of the article is organized as follows. Section 2 provides a partial literature review on previous tolling approaches. Section 3 then specifies the Δ-*tolling* framework. Next, Sections 4–6 demonstrate Δ-*tolling* on the three different traffic models, describing the demand and traffic flow assumptions in each scenario, and providing numerical results testing its effectiveness. Finally, Section 7 concludes the article and discuss future research directions.

2 Related Work

Road pricing has received considerable attention due to its potential to reduce congestion, and the economic fairness of charging users for the delays they cause to other travelers. It has long been established that in a static equilibrium setting, marginal tolls can eliminate the inefficiency associated with selfish routing (Pigou, 1920; Beckmann et al., 1956). A detailed history along with practical aspects of congestion pricing can be found in (de Palma and Lindsey, 2011). However, such steady state conditions rarely exist in practice. Changes in supply, demand and other driver characteristics such as bounded rationality and value of time result in traffic that is dynamic both day-to-day and within-day. To control congestion in the presence of these factors, researchers have proposed a wide range of tolling models, based on different representations of traffic flow and different assumptions on the source of variability. In this section, we review these studies and highlight some of the gaps that will be addressed in this article.

2.1 Pricing models using link delay functions

The simplest way to model congestion is by using link delay functions that predict the travel time on links as a function of its traffic volume. In this subsection, we review adaptive pricing articles built on this assumption. These studies can be broadly grouped into the following three categories: (1) pricing models with route switching behavior, (2) congestion pricing under supply and demand side uncertainty, and (3) trial-and-error methods for congestion pricing.

(Friesz et al., 2004) proposed an optimal control formulation for finding tolls that maximize social welfare, while also achieving a minimum revenue target. Traveler choices were represented by an ordinary differential equation that corresponds to a *tâtonnement* route switching process. In a similar vein, (Yang et al., 2007) suggested an adaptive tolling framework that reaches the system optimum state, assuming that travelers follow the proportional switch adjustment process (Smith, 1984). (Tan et al., 2015) incorporated user heterogeneity in such day-to-day pricing models by proposing a multi-class flow evolution dynamic in which users with different values of time respond differently to current congestion levels. Adaptive tolls that minimize a weighted sum of system cost and time were sought. (Guo et al., 2016) and (Rambha and Boyles, 2016) studied similar problems with slightly different objectives in discrete time settings where travelers' choices depended on the previous day's flow. While the former focused on the asymptotic behavior of the system under the rational behavior adjustment process (Guo et al., 2015), the latter supposed that users select routes according to the logit choice model and the tolls were used to minimize the average system travel time over an infinite horizon.

Marginal tolls are usually computed assuming fixed trip tables and fixed network parameters such as capacity and free flow travel time. However, when these supply and demand inputs are uncertain, the marginal prices can be non-optimal and in some cases may worsen the network performance (Gardner et al., 2008, 2010; Boyles et al., 2010). To address this issue, (Gardner et al., 2011) defined six scenarios that take into account the information states of the system manager and travelers and suggested different optimization models in which responsive tolls are designed based on the realizations of the supply and demand. Recently, (Rambha et al., 2017) extended the problem of finding the optimal tolls under supply side uncertainty to cases in which travelers respond to online information by changing their decisions *en route*.

A third class of adaptive pricing models are what are called trial-and-error methods. Tolls in these models vary across different days but are not set to address the variability in network congestion. Instead, the tolls are adaptive because the system manager, may in reality, not know the demand and delay functions. By levying certain "trial-and-error tolls" and updating them over different days using observed link volumes, provable convergence to system optimal tolls can be guaranteed (Yang et al., 2004; Han and Yang, 2009; Yang et al., 2010). While travelers in these models are assumed to be aware of current day's tolls and react rationally, extensions in which travelers respond using day-to-day route dynamics also have been proposed (Ye et al., 2015). The Δ-*tolling* framework with link delay functions that will be described in Section 4 can be seen as a variant of the trial-and-error method.

2.2 Pricing models using macroscopic traffic simulators

The earliest work studying dynamic congestion used Vickrey's (1969) bottleneck model. Road pricing also affects the number of trips and their departure times, this elasticity was added by (Arnott et al., 1993). Demand elasticity could be affected by alternative modes, so (Danielis and Marcucci, 2002) combined the bottleneck model with a railroad mode. Similarly, (Huang, 2002) studied a bottleneck model with a parallel mass transit alternative mode, and compared the effects of several pricing schemes on congestion and overall system efficiency. (Verhoef, 2003) proposed a heuristic for adaptive tolling for dynamic traffic congestion in continuous time. (van den Berg and Verhoef, 2011) extended the results to continuous distributions of values of time, and found that congestion pricing could improve the social welfare of the majority of travelers even without returning toll revenues.

Adaptive tolling has also been widely studied in the context of managed lanes. (Yin and Lou, 2009) suggested a feedback control approach in which tolls are raised or lowered proportionally to the difference between the current and desired occupancy. They also proposed a self-learning approach in which the willingness to pay is estimated in an online manner and the lane choice is captured using a logit choice model, which was later extended to a multi-lane hybrid traffic flow model (Lou et al., 2011). (Gardner et al., 2013) analyzed managed lane pricing using an additive logit model and an all-or-nothing assignment. Extensions that incorporate demand uncertainty (Gardner et al., 2015) and departure time choices (Boyles et al., 2015) were also studied.

For pricing at a network level, (Carey and Srinivasan, 1993) define dynamic externalities and tolls using exit flow functions as defined by (Merchant and Nemhauser, 1978). (Wie and Tobin, 1998) formulated optimal control programs assuming point-queue models. (Wie, 2007) suggested a bi-level model in which the lower level involves a simplified dynamic traffic loading mechanism. It was assumed that a subset of arcs can be tolled and the objective was to maximize net consumer surplus. However, the traffic flow models in these approaches do not capture queue spill-backs. (Tsekeris and Voß, 2009) reviews several studies that used bottleneck or point queue models, and therefore did not capture the effects of queue spillback. (Waller et al., 2006) and (Lo and Szeto, 2005) showed that these traffic flow dynamics were important to the effectiveness of congestion pricing, and that ignoring them could result in tolls that *increased* congestion. Therefore, it is important to study tolling schemes such as Δ-*tolling* on mesoscopic or microscopic models that properly capture spatial propagation of congestion.

For more realistic flow models applied to large networks, such as the hydrodynamic model (Lighthill and Whitham, 1955; Richards, 1956), marginal costs are much more difficult to compute because of discontinuities in the flow model and congestion effects that transcend link boundaries. For such models it is not known how to reduce the problem beyond its fundamental bi-level form. The upper-level problem chooses the optimal tolls subject to route choice constraints, which form the lower-level problem. These route choice constraints are often in the form of dynamic traffic assignment (Chiu et al., 2011), which itself is a difficult optimization problem. Such bi-level problems have been studied extensively for both static and dynamic flow models as *network design problems* (Farahani et al., 2013), and are known to be NP-hard even when both the upper-level and lower-level problems are convex. Consequently, they are typically solved using heuristics or meta-heuristics. (Lin et al., 2011) formulated such a bi-level program in which route and departure time choices for a single destination network were captured with the cell transmission model. A dual variable based heuristic was used to solve the proposed MPEC (mathematical program with equilibrium constraints). (Joksimovic, Bliemer, Bovy and Verwater-Lukszo, 2005) and (Joksimovic, Bliemer and Bovy, 2005) included both departure time and route choice in the lower level through a discrete choice model (stochastic user equilibrium). Although results used a small test network, tolls were observed to encourage travel on less congested routes or departure times. (Ekstrm et al., 2016) devise a surrogate-based optimization method in which a small number of dynamic cordon tolling schemes are tested on the Stockholm network using VisumDUE, a dynamic traffic assignment tool.

When applied to real-world traffic networks, Δ-*tolling* can be classifed as a within-day pricing scheme, as described in Sections 5 and 6. *Within-day pricing* varies the tolls at different times of day in response to (expected) congestion. Within-day strategies can affect both route choice and departure times because travelers may respond to congestion pricing by delaying their trip until a less congested time. While many within-day strategies set a predictable schedule of tolls that human drivers can react to, Δ-*tolling* assumes the use of route guidance software that can more quickly react to rapidly changing tolls. Furthermore, Δ-*tolling* is responsive to fluctuations that may be caused due to uncertain demand.

2.3 Pricing models using microscopic traffic simulators

Very few researchers have used microscopic traffic simulators to study congestion pricing. The outputs of microscopic simulators cannot be analytically expressed, a challenge in optimizing control strategies. Thus, existing studies have used feedback mechanisms for computational tractability. The Δ-*tolling* approach proposed in this article can also be seen as a feedback mechanism much like ramp metering strategies such as ALINEA (Papageorgiou et al., 1991). (Zhang et al., 2008) developed a feedback control theory-based tolling for high-occupancy/toll lanes in VISSIM to avoid the potential hysteresis problem, and was later extended by (Cheng et al., 2014) to include the effects of travel time reliability and income levels of users. However, their pricing model is fairly complicated and may be difficult to apply to other traffic flow models, and testing was limited to several connected freeway segments. (Zheng et al., 2012) and (Simoni et al., 2015) used a hybrid approach involving a microscopic simulator MATSim and a macroscopic fundamental diagram flow model to set cordon tolls in the city of Zurich. This model was used to calculate the aggregate density, from which the cordon tolls were inflated or deflated based on a linear feedback control strategy. (Grether et al., 2008) also used MATSim along with an activity based model that simulated users plans, modes of travel, and values of time, but only evaluated fixed time-of-day dynamic tolls.

3 Framework

The Δ-*tolling* framework is designed to be widely applicable across a broad range of traffic and user behavior assumptions. The modeling framework has three major components:

- The traffic model.
- The travel time calculation model.
- The tolling model.

These three models make use of four variables:

- τ - the vector of tolls applied to each link.
- d - the travel demand, expressed in amount of vehicles departing each origin towards each destination.
- X - the system state, a tuple of sets or vectors reflecting current traffic conditions corresponding to a particular traffic model.
- T - the vector of measured link travel times.

If the underlying network is represented by $G = (N, A, Z)$ where N and A are the sets of nodes and links, and $Z \subseteq N$ is the set of origins and destinations where trips start and end, then we have $\tau \in \mathbb{R}^{|A|}$, $d \in \mathbb{R}^{|Z|\times|Z|}$, and $T \in \mathbb{R}^{|A|}$. We use $T_{a\in A}$ to represent the travel time on link a, same goes for X_a and τ_a (a link (a) might also be expressed as a pair i, j representing the link connecting node i with j). Similarly, $d_{r,s\in Z}$ is the demand originating at node r towards node s.

Each of these variables evolves over time according to the traffic flow, travel time, and tolling models, which are described next. The reader may find it useful to refer to Figure 1 during this discussion. Arrows in this figure reflect direct dependencies between the variables, as described below.

The **traffic flow** model $\mathcal{M}$ encompasses the routing decisions made by drivers, as well as the congestion effects caused by interactions amongst drivers. We express this relationship as

$$X^t = \mathcal{M}(X^{t-1}, T^{t-1}, d^t, \tau^t)\,. \tag{1}$$

This equation represents the following potential dependencies: the system state at time t may depend on the system state at the previous time interval (X^{t-1}); the measured link travel times at the previous time interval (T^{t-1}); the vehicles departing during time t (d^t), and the current tolls (τ^t).[2] Particular traffic flow

[2]The time interval between time steps (t, $t + 1$) may differ between models and between instances of the same model. Examples are given in the following sections.

models may not make use of all of these dependencies (for instance, the traffic state for Model A in Section 4 does not explicitly depend on previous time intervals) but they are included for generality. All of the models in this article assume that the last measured travel times T^{t-1} and current tolls τ^t are communicated to all vehicles (based on the assumption of CAV technology), but in principle the framework could allow for route choice decisions made without perfect knowledge of these. Specific examples of traffic models $\mathcal{M}$ and the corresponding system states X are given in Sections 4–6.

The **travel time calculation** model $\mathcal{T}$ maps the system state to link travel times used for tolling:

$$T^t = \mathcal{T}(X^t)\,. \tag{2}$$

Although the travel times T^{t-1} may be part of the system state X^{t-1}, we include a separate dependence on T^{t-1} to allow for measurement errors, as might occur in practice if travel times are measured from sensors or probe vehicles in the field, or to allow for approximate travel time calculations, as are often obtained from simulation-based traffic models (such as Models B and C below). This separation also emphasizes that the toll calculation only relies on travel times, and does not require additional information about the system state or demand. Examples of travel time models are also found in the sections that follow.

The **toll calculation** model for Δ-*tolling* is the equation

$$\tau^t = (1 - R)\tau^{t-1} + R\beta(T^{t-1} - T^0) \tag{3}$$

where T^0 is the vector of link free-flow times,[3] and R and β are tuning parameters. Parameter β is the proportionality constant relating link delay and the toll value, while R is a weighting parameter that results in an exponential decay effect for tolls assigned in previous time steps. For further explanation regarding the motivation behind the R parameter please refer to Appendix A. Both the R and β parameters must be tuned for a given network.

To summarize, the dependencies between the variables are intended to fit the following story: at time step t all users are informed with the tolls (τ^t) imposed during time interval t to $t+1$, as well as the travel times (T^{t-1}) measured at the end of the previous time interval. Given this data, users choose and follow a route leading from their current location to their destination that optimizes their utility. Based on the routes they choose, the system evolves to state X^t, and at the end of this time interval, the updated travel times (T^t) are measured. These updated travel times, along with the updated tolls, are fed-back to all users which, once again, re-optimize their route. Figure 1 presents a schematic illustrating these dependencies.

Note that allowing users to predict and react to future congestion and tolls can have a positive effect on the system as the convergence towards a user equilibrium would be faster or even instant. On the one hand, assuming such capabilities is reasonable from a practical standpoint (recurring congestion can be predicted). On the other hand, considering such capabilities significantly complicates our theoretical and empirical models. As a result they are not assumed in this study. Nonetheless, due to its relevance to the application of Δ-*tolling*, we intend to explore this topic in future work.

The following three sections show specific instantiations of this framework for varied traffic models: one inspired by day-to-day pricing in static traffic assignment, and two meant to represent within-day pricing in dynamic models (the cell transmission model and a microsimulator). The intent of these demonstrations is to study the performance of Δ-*tolling* across different modeling contexts, in contrast to many prior studies which demonstrate effectiveness only in a single traffic model (often the same model used to derive the tolling scheme). The focus of the presented experiments is on robustness of performance across widely-varying models, rather than claiming that any of the specific models is the "right" one for any particular application.

[3]Free flow travel time T^0 may change over time e.g., due to weather conditions, and should be appropriately updated.

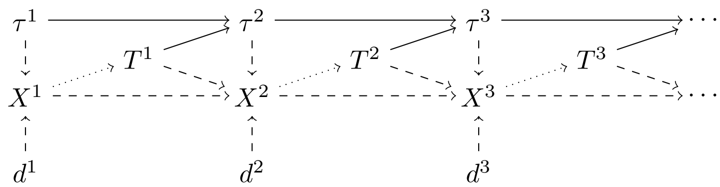

Figure 1: Schematic for Δ-*tolling* framework. Dashed lines reflect the traffic model, dotted lines the travel time calculation model, and solid lines the tolling model. Note that the toll calculation only requires travel times, and no other details of the demand or traffic model.

4 Model A: Link Performance Functions

In the first model, the traffic model is based on link performance functions which map the flow on each link to its travel time. The presumption is that the intervals between successive time steps (and toll updates) are large enough that most trips can be completed during a single time interval, and that delays and flows can be well-approximated by deterministic, steady-state conditions. In this model, we assume that drivers choose routes to minimize their travel cost (time+tolls). Because delays are deterministic and the interval between time steps is assumed long, we further assume that CAV technology can choose routes rationally, leading to a Nash equilibrium in which each vehicle chooses a route with minimum cost (cf. Wardrop, 1952).

4.1 Model specification

For an origin-destination (OD) pair $(r, s) \in Z^2$, let Π_{rs} denote the set of simple paths in G connecting origin r to destination s, $\Pi_s = \cup_{r \in Z} \Pi_{rs}$ the set of paths ending at s, and $\Pi = \cup_{s \in Z} \Pi_s$ the set of all network paths. Let d^t be the demand during time period t. The vector of path flows $h^t \in \mathbb{R}^{|\Pi|}$ is feasible if each element is nonnegative, and if $\sum_{\pi \in \Pi_{rs}} h^t_\pi = d^t_{rs}$ (each vehicle is assigned a path for its OD pair). Each vector of path flows generates a vector of link flows $x^t \in \mathbb{R}^{|A|}$ given by $x^t_{ij} = \sum_{\pi \in \Pi : (i,j) \in \pi} h^t_\pi$. Furthermore, we assume that each link $(i, j) \in A$ is equipped with an increasing and differentiable link performance function $\Psi_{ij}(x_{ij})$, giving the travel time on link (i, j) as a function of its flow x_{ij} alone. That is, $T_{i,j} = \Psi_{i,j}(x_{i,j})$. Let $\mathcal{X}(d^t)$ denote the set of feasible link flows when the demand is d^t, that is, the set of link flow vectors corresponding to a feasible path flow vector.

For this model, the state vector is simply the vector of link flows:

$$X^t = x^t \tag{4}$$

and the traffic model $\mathcal{M}$ in equation (1) is specified with the following formula:

$$x^t = \arg \min_{x \in \mathcal{X}(d^t)} \sum_{(i,j) \in A} \int_0^{x_{ij}} \left(\Psi_{ij}(x) + \tau^t_{ij} \right) \, dx \, . \tag{5}$$

In this formula we assume that the value of time is homogeneous among all travelers, allowing us to choose units so that the tolls τ and travel time t can be directly added. The minimizer of the function on the right-hand side is known to satisfy the Nash equilibrium principle, and to be unique under the assumption that the Ψ_{ij} are increasing. Note that there is no explicit dependence on X^{t-1} in Model A.

One special case of Model A occurs when the demand is stationary with time, $d^t \equiv d$. In this case, Model A can be thought of as a day-to-day tolling model, where the tolls are updated on sequential days as drivers make the same trips.

4.2 Scenario specification

Every scenario simulated through Model A follows the following principles:

- Demand model - demand is modeled as a fixed amount of flow that needs to be routed between any two given nodes in the network at any time step.
- Vehicle model - in model A there is no notion of atomic vehicles, traffic is viewed as a set of infinitely divisible flows.
- Path assignment model - instead of assigning paths to vehicles, Model A assigns flows to paths. The vehicles comprising each flow are assumed to be self interested and are assigned the minimal generalized cost path (travel time + tolls). Such a policy leads to the Nash equilibrium that is defined by Equation 5.

4.3 Theoretical results

One advantage of Model A is that the analytical form of the traffic model (5) is amenable to mathematical analysis. In particular, we are able to show several optimality results if demand is stationary with time and the link performance functions are of the form specified by the Bureau of Public Roads (BPR),

$$\Psi_{ij}(x_{ij}) = T_{ij}^0 \left(1 + A\left(\frac{x_{ij}}{u_{ij}}\right)^B\right) \tag{6}$$

where T_{ij}^0 is the free-flow travel time, u_{ij} the practical capacity, and A and B are shape parameters assumed uniform throughout the network. In this section, we mean "optimality" in the sense of minimizing the average travel time, which is proportional to $\sum_{(i,j)\in A} x_{ij}\Psi_{ij}(x_{ij})$. This is a convex function of the link flows, so optimal link flows x_{ij}^* exist and are unique.

The next results make use of the following well-known facts. (Both can easily be shown by writing the optimality conditions of the associated convex minimization programs.)

Fact 1. *Let $x^t \in \mathcal{X}(d^t)$ and let h^t be any feasible path flow vector generating x^t. The flows x^t satisfy* (5) *if and only if every positive component of h^t corresponds to a path whose generalized cost (the sum of $T_{ij}^t + \tau_{ij}^t$ along its links) is minimal for its OD pair.*

Fact 2. *Let $x^t \in \mathcal{X}(d^t)$ and let h^t be any feasible path flow vector generating x^t. The flows x^t are optimal if and only if every positive component of h^t corresponds to a path whose marginal cost (the sum of $T_{ij}^t + x_{ij}^t\Psi_{ij}'(x_{ij}^t)$ along its links) is minimal for its OD pair.*

The first result shows that the Δ-*tolling* rule is equivalent to marginal cost pricing if the parameter β is chosen correctly:

Proposition 1. *If the link performance functions are of the BPR form and $\beta = B$, then $\beta(T_{ij} - T_{ij}^0) = x_{ij}\Psi_{ij}'(x_{ij}) = BAT^0(x_{ij}^t/u_{ij})^B$.*

Proof. Routine. $\square$

Next, we show that if the tolls do not change from one time step to the next, then the resulting link flows must be optimal. Furthermore, the system has reached a stable state, and the optimal state will persist for future time iterations. In other words, if the tolls are stable, the flows are optimal and stable.

Proposition 2. *Let the link performance functions have the BPR form, and let demand be stationary with time. If $\beta = B$ and $\tau^t = \tau^{t+1}$ for any time interval t, then x^t is optimal, and furthermore x^T is optimal for any $T > t$.*

Proof. Stationary demand implies that the set of feasible x and h are stationary. By the Δ-toll update rule (3), if $\tau^t = \tau^{t+1}$ then we must have $\tau^t = \tau^{t+1} = \beta(T^t - T^0) = BAT^0(x^t/u)^B$, where this vector equation holds componentwise. Let h^t be a feasible path flow vector generating x^t. By Fact 1, every positive component of h^t corresponds to a path whose sum of $T^t_{ij} + BAT^0_{ij}(x^t_{ij}/u_{ij})^B$ along its links is minimal. But by Proposition 1, this implies that every component of h^t corresponds to a path whose sum of $T^t_{ij} + x^t_{ij}T'_{ij}(x^t_{ij})$ is minimal, and thus x^t is optimal by Fact 2. Furthermore, since the minimizer of (5) is unique, $\tau^t = \tau^{t+1}$ implies $x^t = x^{t+1}$, and thus $\tau^T = \tau^t$ and $x^T = x^t$ whenever $T > t$. □

The converse of Proposition 2 is not true, because the vector of tolls which produce a particular flow x under the mapping (5) is not unique, and only the marginal-cost tolls corresponding to Proposition 1 are fixed points of the toll update rule (3). If x^t is optimal but generated by a different toll vector, we will not have $\tau^t = \tau^{t+1}$. Nevertheless, we can show that if x^t is optimal, the flows are stable in subsequent iterations, even if the tolls still change.

Proposition 3. *Let the link performance functions have the BPR form, and let demand be stationary with time. If $\beta = B$ and x^t is optimal at any time interval t, then x^T is optimal whenever $T > t$.*

Proof. Let h^t be any feasible path flow vector generating x^t. By Fact 1, every positive component of h^t corresponds to a shortest path with link weights $T^t_{ij} + \tau^t_{ij}$. Since these flows are optimal, Fact 2 and Proposition 1 imply that these are also shortest paths with link weights $T^t_{ij} + \beta(T^t_{ij} - T^0_{ij})$. That is, for each node i and origin r, there exist node potentials π^r_i and ρ^r_i such that

$$\pi^r_i + T^t_{ij} + \tau^t_{ij} \geq \pi^r_j \tag{7}$$

$$\rho^r_i + T^t_{ij} + \beta(T^t_{ij} - T^0_{ij}) \geq \rho^r_j \tag{8}$$

for each link $(i,j) \in A$ and origin $r \in Z$, with equality holding along all of the paths with positive flow at h^t. Multiplying inequality (7) by $1 - R$, inequality (8) by R, and adding, we have

$$(1-R)\pi^r_i + R\rho^r_i + T^t_{ij} + \tau^{t+1} \geq (1-R)\pi^r_j + R\rho^r_j\,, \tag{9}$$

implying that $(1-R)\pi + R\rho$ form valid node potentials for the new toll vector τ^{t+1}, with equality holding for exactly the same links and origins as before. Thus, the shortest paths with respect to $T^t + \tau^{t+1}$ are the same as those with respect to $T^t + \tau^t$, and x^t remains optimal for x^{t+1}. The argument can be repeated for any $T > t$. □

4.4 Experiments and results

Model A was implemented in C using Algorithm B (Dial, 2006) to solve the equilibrium subproblem. It was tested on two city networks, representing the cities of Sioux Falls, SD and Austin, TX. The Sioux Falls network is a standard test instance in the transportation network literature (Bar-Gera, 2014), with 76 links, 24 nodes, and 360,600 trips spanning 24 hours. The Austin network represents the central business district of the city, and contains 1247 links, 546 nodes, and 62,836 trips over a two-hour morning peak period. Additional details on the Austin network can be found in (Levin et al., 2015). Figure 2 shows schematics of both networks.

In both networks, the link performance functions are BPR functions, using the standard values of the shape parameters: $A = 0.15$ and $B = 4$. For the experiments in this article, the time intervals t were interpreted

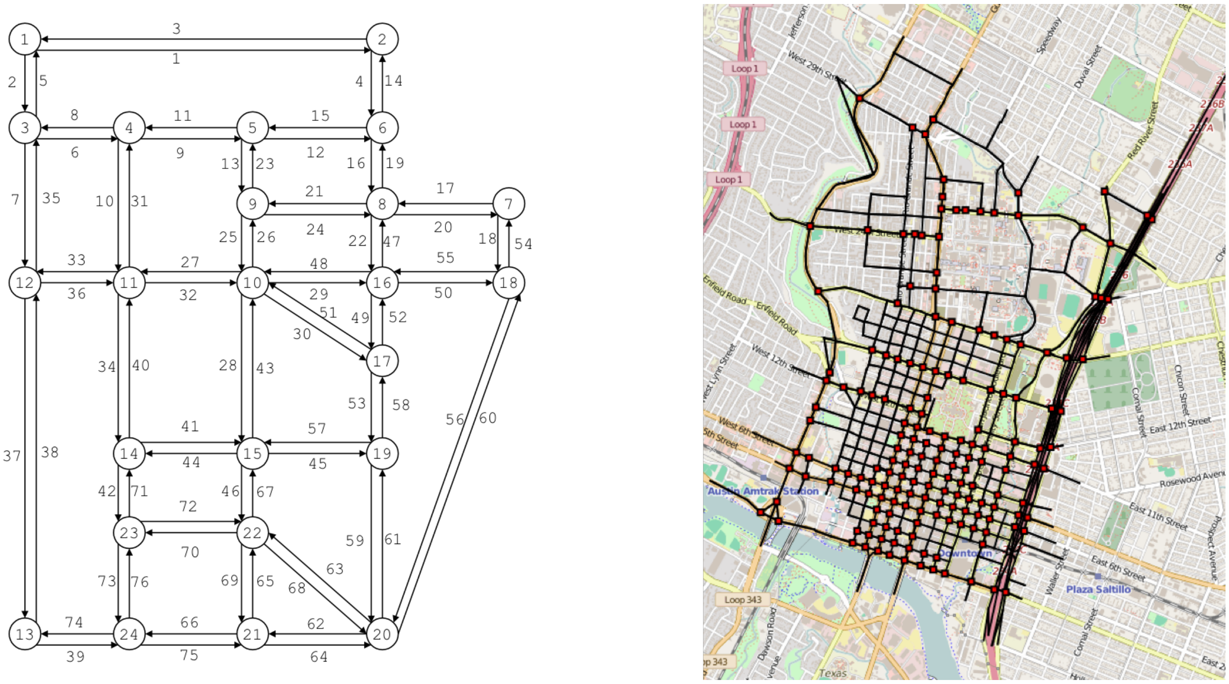

Figure 2: Sioux Falls (left) and Austin (right) networks.

as subsequent days, so demand was assumed stationary and the experiment represents a day-to-day pricing scenario.

Because Model A assumes fixed demand and homogeneous travelers, we can use the average travel time $ATT^t = (x^t \cdot T^t)/(d \cdot 1)$ as a performance metric. In this demonstration, the weights $R^t = 1/(t+1)$ were chosen, effectively setting the toll during time step t to the average of the daily "target" tolls $\beta(T^t - T^0)$. This choice of R^t value was inspired by the method of successive averages (Liu et al., 2009).

Table 1 shows the steady-state average travel time for both networks as the sensitivity parameter β varies. For both networks, with the above choice of R^t, fairly rapid convergence was obtained to a steady state.[4] For β values of 1, 2, 4, 8, convergence on Sioux Falls required 95, 27, 11, and 94 iterations, respectively. For Austin, these values required 24, 16, 27, and 42 iterations to converge, respectively.

Over the range of β values tested, the Δ-*tolling* strategy always reduced average travel time from the no-toll value. When β was set equal to the B exponent in the link performance functions, the travel times were the lowest observed, and in fact correspond to the system-optimal solution, as suggested by Proposition 1. Note that the performance of Δ-*tolling* seems to be insensitive to the chosen β value as *beta* values twice as big (8) or small (2) from the optimal (4) result in a system performance which is almost identical to the optimal one.

5 Model B: Cell Transmission Model

Model B implements the Δ-*tolling* framework in the cell transmission model (CTM) developed by (Daganzo, 1994, 1995*a*) as an explicit solution method for the hydrodynamic theory of traffic flow proposed by (Lighthill and Whitham, 1955) and (Richards, 1956). CTM is frequently used in dynamic traffic assignment. The time

[4]The system is said to converge to a steady state if the change in average travel time between successive time steps was less than a tenth of a millisecond

Scenario	No tolls	SO	$\beta = 1$	$\beta = 2$	$\beta = 4$	$\beta = 8$
Sioux Falls						
Model A	20.74	19.95	20.09	19.98	**19.95**	19.96
Model B	24.74	NA	20.28	20.08	**19.92**	20.26
Downtown Austin						
Model A	21.92	21.78	21.81	21.79	**21.78**	21.79
Model B	20.67	NA	16.06	**15.64**	15.82	17.39

Table 1: Average travel time (minutes) at UE for different β values using Models A ($R^t = 1/(t+1)$) and B ($R = 10^{-4}$). Note that for Model A with $\beta = 4$ the UE and SO allign (this is the provable SO for model A). The SO for Model B is not applicable .

step used in this model is typically short, on the order of a few seconds. When used with Δ-*tolling*, this allows for a truly adaptive toll which can be updated based on traffic conditions. Because the tolls are updated frequently, we believe that the equilibrium behavior assumed in Model A is not appropriate. Rather, we assume that CAVs dynamically receive updates of travel times and tolls, and may alter their route to the destination in response to receiving such information. This update is based on the currently reported travel times and tolls when they reach the diverge point.[5] Modeling the behavior of specific vehicles (each individual vehicle has a value-of-time, and chooses routes according to its own generalized cost function) requires that model B follows a variant of CTM that models discretized flow (Ziliaskopoulos and Waller, 2000; Chiu et al., 2011).

5.1 Model specification

The CTM divides each link into a set of cells, each of length equal to the distance a vehicle would travel in one time step at free-flow conditions. Cells result from a fixed division of a link into discrete segments. The length of each segment corresponds to the distance a vehicle would travel on that link at free-flow speed in one time step. This choice of cell length ensures stability of the cell transmission model (it satisfies the Cournout-Friedrich-Lewy conditions for the underlying system of partial differential equations). Let C be the set of cells in the entire network, and for a given cell $c \in C$, let $C^+(c)$ denote the set of cells immediately downstream of c, and $C^-(c)$ the set of sells immediately upstream of c. For cells in the interior of a link $|C^+(c)| = |C^-(c)| = 1$, whereas if c is at the upstream end of a link we may have $|C^-(c)| > 1$ (if there are multiple links incident from upstream) and if c is at the downstream end we may have $|C^+(c)| > 1$ (if there are multiple links adjacent downstream). Let $x^t_{c,\pi} \in \mathbb{N}$ denote the number of vehicles in cell c at the start of time interval t which are currently following path $\pi \in \Pi$, and $x^t_c = \sum_{\pi\in\Pi} x^t_{c,\pi}$ the total number of vehicles in the cell. Based on these values, the CTM respectively defines the *sending flow* S^t_c and *receiving flow* R^t_c as the maximum number of vehicles which could possibly exit cell c during time step t, and the maximum number of vehicles which could possibly enter cell c during time t. If we denote Q_c as the capacity of cell c, N_c the maximum number of vehicles which can physically occupy cell c, and δ the ratio between the maximum backward shockwave speed and free-flow speed, common expressions for the sending and receiving flow are

$$S^t_c = \min\{x^t_c, Q_c\} \tag{10}$$

and

$$R^t_c = \min\{\delta(N_c - x^t_c), Q_c\}\,. \tag{11}$$

See (Daganzo, 1995*a*,*b*) for additional details.

The update rules for the cell transmission model (the traffic model $\mathcal{M}$) involve the auxiliary variables y^t_{cd}

[5]Particularly, vehicles do not anticipate future changes in travel conditions, nor the impact of receiving future information, cf. (Waller and Ziliaskopoulos, 2002; Boyles, 2009). While such computations are more involved, they may be feasible with CAV technology, and would be an interesting topic for future study.

representing the number of vehicles moving from cell c to cell $d \in C^+(c)$ during time step t. If c and d belong to the same link, we have

$$y_{cd}^t = \min\{S_c^t, R_d^t\}\,. \tag{12}$$

If c and d belong to different links, y_{cd}^t is calculated using various *intersection models* representing traffic behavior at diverges, merges, traffic signals, roundabouts, or other intersection types; see (Tampère et al., 2011) for discussion of intersection model desiderata and some examples. The simulations reported in this section use intersection models reflecting traffic signals.

Once the y^t values are calculated, cell occupancies update in the natural way:

$$x_c^{t+1} = x_c^t + \sum_{b \in C^-(c)} y_{bc}^t - \sum_{d \in C^+(c)} y_{cd}^t \tag{13}$$

with the path-disaggregated $x_{c,\pi}^t$ values updated according to the first-in, first-out principle.

The travel time model for T_{ij}^t calculates the average time spent on link (i,j) by the vehicles which most recently exited the link, based on the difference between their entry and exit times. This results in a slight lag in the travel times used for Δ-*tolling* and for the adaptive routing procedure, since link travel times for vehicles are not computed until they leave the link, even though the routing and tolls are based on decisions made as vehicles enter a link.

5.2 Scenario specification

Every scenario simulated through Model B follows the following principles:

- Demand model - demand is given as the amount of vehicles originating at node n_1 at time step t and are assigned a destination n_2 for any $n_1, n_2 \in N$ and $t > 0$ combination.
- Vehicle model - each vehicle is affiliated with a *value of time*[6] (VOT). Vehicles seek to minimize their generalized cost that is defined as travel time + tolls × value of time.
- Path assignment model - let T_π^{t-1} be the sum of travel times along path π during time step $t-1$ and let τ_π^t be the sum of tolls along π during time step t. When reaching a diverge node n at time step t all paths (Π_{ns}) leading from n to destination s are considered. The vehicle in question is assigned the path $\arg\min_{\pi \in \Pi_{ns}} cost(\pi) = \{\tau_\pi^t + T_\pi^{t-1} \cdot VOT\}$. An additional rule was added to prevent gridlock problems which can arise in dynamic traffic models when a cycle of links are at jam density: if a vehicle is unable to enter a link because its receiving flow is zero for more than 96 seconds,[7] the vehicle is assigned the least cost path to its destination that avoids that link, if such a path exists.

5.3 Experiments and results

For running Model B we used the DTA simulator (Chiu et al., 2011) implemented in Java. Model B was tested on the Sioux Falls and Austin networks also used for the Model A results. A few changes were needed to accommodate the differences in time scale and modeling assumptions between Models A and B. Because the original Sioux Falls demand was specified over 24 hours, the demand was modified to fit a 3-hour experiment more suitable for dynamic traffic assignment, including 28,835 trips over this time period. Both networks also required traffic signal timings. The Austin network data was originally used for dynamic traffic assignment and contained real-world signal data. The Sioux Falls network does not include this data, so we generated an artificial timing based on Webster's (1958) formula for signal timing. The variations in

[6]Value of time represents the monetary value of a single unit of time. It is used to map time into monetary cost.
[7]This value was chosen by trial-and-error and resulted in the best performance.

departure rates over time for these scenarios can be seen as the solid black line in Figure 4. Whenever a vehicle is loaded onto the network, it is assigned a value of time randomly drawn from a Dagum distribution with parameters $\hat{a} = 22020.6$, $\hat{b} = 2.7926$, and $\hat{c} = 0.2977$, reflecting the distribution of personal income in the United States (Łukasiewicza et al., 2012).

The average travel times for Model B for particular values of R and β are shown in Table 1. Since Model B allows heterogeneity in travelers' values of time, in addition to evaluating average travel times, we also evaluate the *social welfare*, defined as the weighted sum of each traveler's travel time according to his or her value of time.[8] If V is the set of vehicles, and vehicle v experiences a travel time of T_v and has a value of time of α_v, social welfare is defined as $-\sum_v T_v \alpha_v$. The tolls are not included in the calculation of social welfare, because we assume that toll revenues are transfer payments which remain internal to society. This assumption was made implicitly in the use of average travel time as the metric for Model A.

Figure 3 shows the effects of Δ-*tolling* on social welfare as the responsiveness parameter R varies. The red series indicates the no-toll scenario used as a benchmark, while the blue series shows social welfare under the Δ-*tolling* regime. Each data point represents the average of ten scenario runs, and the bands represent 95% confidence intervals around each point. Scenario runs differ from each other in the value of time that is (randomly) assigned to each vehicle. For the experiments in this plot, a fixed value of $\beta = 4$ was used. This value was optimal in Model A for its assumption of BPR-type delay functions, but this choice is not necessarily optimal for Model B which uses a dynamic network loading procedure. (As shown in Table 1, in Austin, a slightly better performance was obtained with $\beta = 2$ compared to $\beta = 4$.) In this plot, note that extreme R values are less effective, whether high or low: when R is near 1 tolls oscillate rapidly (worsening performance over the no-toll baseline), and when R is near zero the tolls have little impact. The best performance occurred in a narrow band around $R \approx 10^{-4}$, yielding increases of 26% and 33% in social welfare for the Sioux Falls and Austin scenarios, respectively. Indeed, for any given β value, near-optimal tolls cluster around a single order of magnitude of R values.

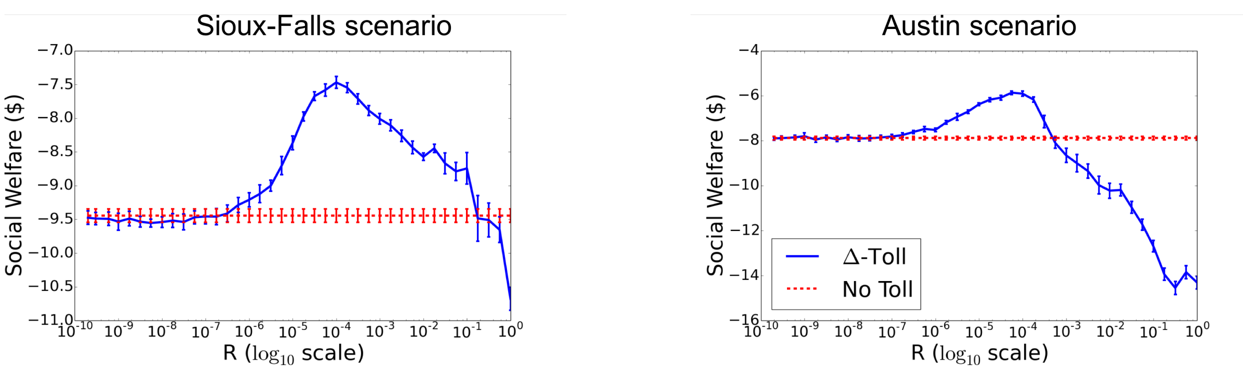

Figure 3: Social welfare in Sioux Falls (left) and Austin (right) as responsiveness parameter R varies using Model B. The legend is identical in both plots.

Figure 4 shows the impacts on social welfare for travelers departing at different times during the simulation for R values in the range 10^{-5}–10^{-3}. All series in this plot correspond to $\beta = 4$. This plot shows the superior performance of $R = 10^{-4}$ in a different way, and indicates that the benefits from tolling increase throughout the peak period, and that the onset of congestion is delayed. In the Sioux Falls scenario, which includes the end of the peak period, we see that the recovery from congestion occurs earlier as well.

Figure 5 shows performance of the tolling scheme as both R and β vary, depicting the difference in social welfare (in percentage) compared to a scenario where no tolls are applied (a value of 126, for instance, correspond to a 26% increase in social welfare). For any fixed value of one parameter, there is a near-optimal

[8]For Model A, where value of time is homogeneous, social welfare is directly proportional to the average travel time.

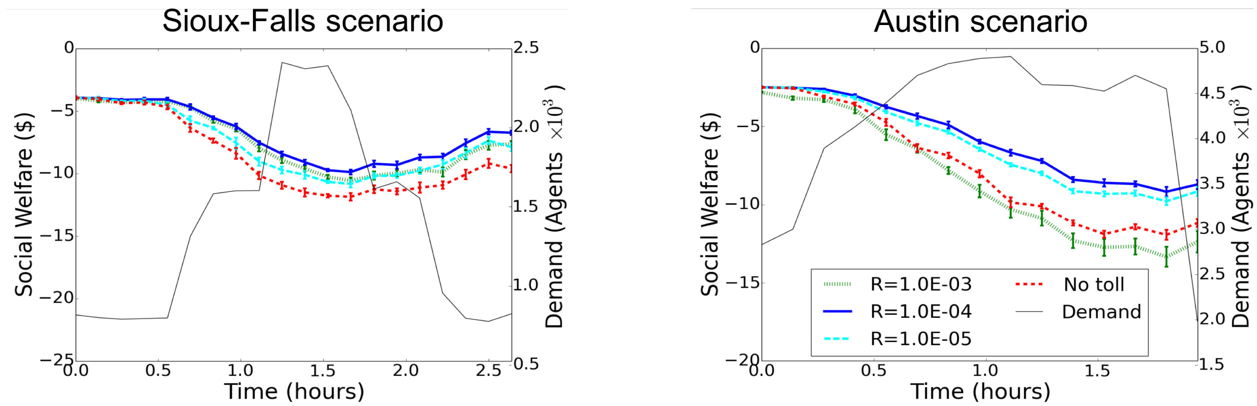

Figure 4: Social welfare over time in Sioux Falls (left) and Austin (right), R varies using Model B. The legend is identical in both plots.

value for the other parameter. This observation leads us to suspect that there are dependencies between the two parameters. We leave exploring such dependencies for future work. Nonetheless, this observation justifies the use of a single β value in the results described earlier.

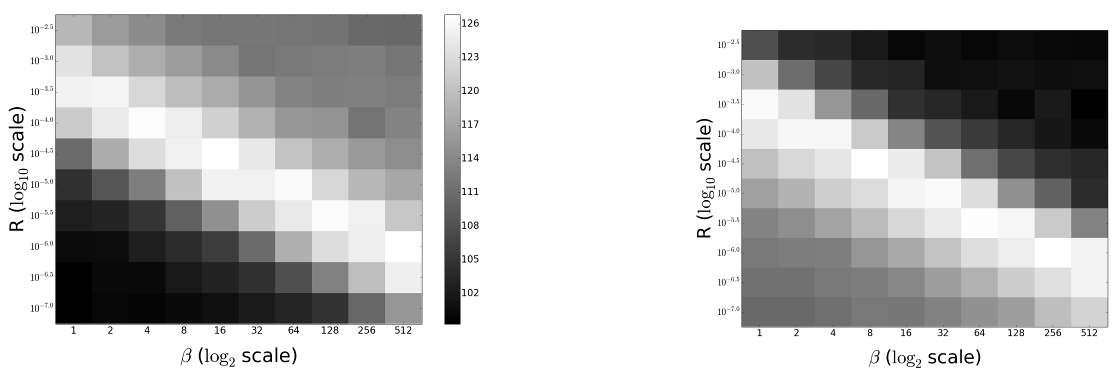

Figure 5: Heat maps showing the difference (in precentage) in social welfare compared to the no tolls scenario for different R and β values in Sioux Falls (left) and Austin (right). (Model B)

6 Model C: Microsimulation

AIM (Autonomous Intersection Manager) is a traffic microsimulator developed by (Dresner and Stone, 2004) to model the flows of CAVs at intersections where priority is granted by reservation, rather than with signals. AIM provides a multiagent framework for simulating autonomous vehicles on a road network grid, and it presents a much more detailed traffic flow model than Models A or B. The AIM simulator uses two types of agents: intersection managers and driver agents. Intersection managers are responsible for directing the vehicles through the intersections, while the driver agents are responsible for controlling the vehicles to which they are assigned. To improve the throughput and efficiency of the system, the driver agents "call ahead" to the intersection manager and request a path reservation (space-time sequence) within the intersection. The intersection manager then determines whether or not this request can be met. If the intersection manager approves a driver agent's request, the driver agent must follow the assigned path through the intersection.

On the other hand, if the intersection manager rejects a driver agent's request, the driver agent may not pass through the intersection but may attempt to request a new reservation. AIM has been used in various studies on reservation-based intersection control: (Dresner and Stone, 2006, 2007) studied variants of the reservation protocol that provided intersection access to human drivers through an occasionally activated traffic signal, and (Fajardo et al., 2011) found that reservations had lower delays than optimized traffic signals for a symmetric intersection. Figure 6 shows a typical snapshot of simultaneous vehicle flow at a congested intersection.

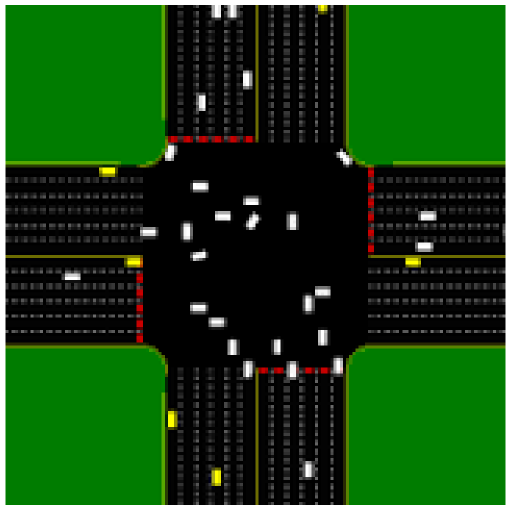

Figure 6: The AIM simulator depicting a reservation-based intersection in operation.

Link travel times T^{t-1} are estimated as an average of the time spent on each link by the vehicles most recently exiting. The scenario specifications for this model are identical to those specified for model B. As with Model B, due to the frequency of updates there is no presumption that an equilibrium is reached when vehicles choose routes.

Unlike Models A and B, the microsimulation environment is not well-suited to explicit mathematical description. The state vector can be defined as $X^t = (V^t, I^t)$, where V^t is the set of vehicles on the network at the start of time step t, including associated information such as their value of time and position, velocity, and acceleration in the network, and I^t is the set of intersections and associated information at time t, such as the trajectories of scheduled reservations. This information is updated according to the rules described above.

6.1 Experiments and results

For running Model C we used the AIM4 microsimulator (http://www.cs.utexas.edu/ aim/aim4sim/aim4-release-1.0.3/aim4-root/docs/install.html). AIM4 is unable to model large networks of the type used for Models A and B, because of the level of detail in its representation of agent behavior (both vehicles and intersection reservations). Hence, the Sioux Falls and Austin scenarios are intractable within AIM. Figure 7 shows the 3×3 grid network used for these experiments. Vehicle agents are generated randomly according to a Poisson process, at a mean rate of 500 vehicles per hour per incoming lane. Each vehicle is assigned either to destination D1 or D2. The network also includes alternative destinations for vehicles headed to either of these destinations. Alternative destinations are used to simulate route choice effects on a network much smaller than the city networks used in Models A and B. These alternative destinations, marked as A1 and A2 in Figure 7, are associated with a time penalty if vehicles leave the network through them instead of their original destination. Vehicles may opt for a path ending at an alternative destination when performing

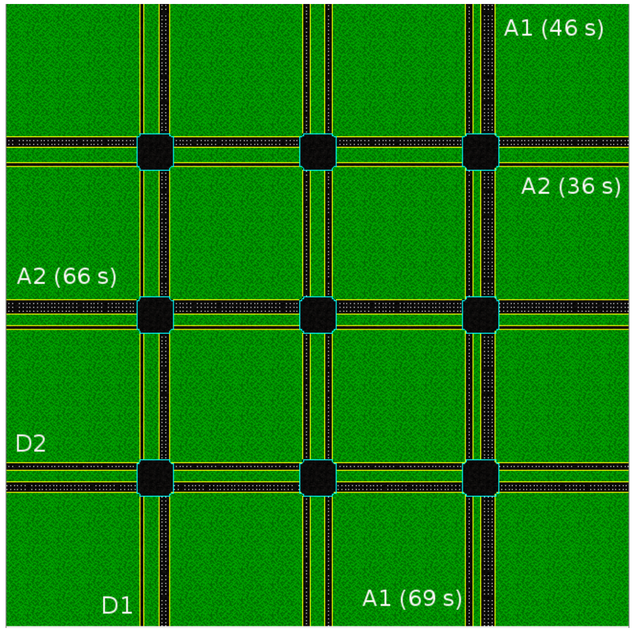

Figure 7: Grid network used for Model C results, with destinations and alternatives marked.

the A^* search when arriving at each intersection.

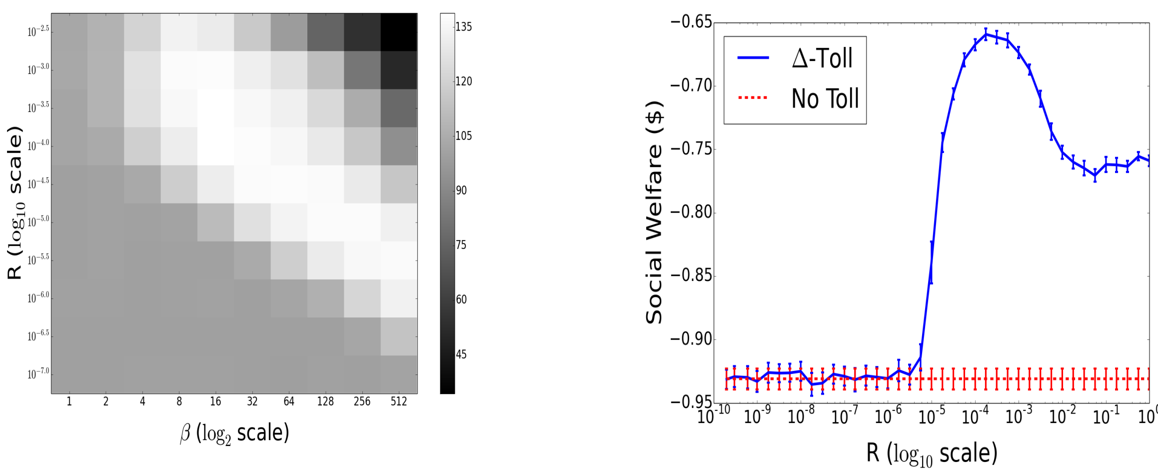

Figure 8: Results from running AIM in the 3×3 grid network. Heat map showing the difference (in precentage) in social welfare compared to the no tolls scenario for different R and β values (left). Social welfare as responsiveness parameter R varies (right). (Model C)

Figure 8 present results that are similar in format to those presented for Model B (Figures 3 and 5). That is, in the left figure, social welfare as responsiveness parameter R varies and β is set (β equals 16). In the right figure, heat map showing social welfare with different R and β values.[9] Each data point in the right figure, and each bracket shade in the left figure, represents the average social welfare over 30 simulation runs, where

[9] Results in a format similar to that in Figure 4 are not presented for this model since, unlike the Model B results, traffic demand is not time varying in this experiment.

social welfare for a single simulation is the average utility over all agents (vehicles). Each run simulates one hour of traffic. Error bars reflecting 95% confidence intervals are shown in the right figure (social welfare vs R).

The general trends that are observed in these results are very similar to those observed in the CTM model: that is, reducing R to approximately 10^{-4} improves system performance (due to mitigation of oscillation and spike effects), and that near-optimal performance can be achieved with most β values by properly tuning the R values . Nonetheless, there are two notable differences between these results and those presented for the CTM:

- Low β values (< 8) do not have a suitable R that yield optimized performance. We believe this discrepancy stems from differences in the congestion accumulation model. Recall that AIM manages intersection in a way that is conceptually different than traffic signals. When traffic signals are considered, the marginal impact of a single vehicle is negligible at low traffic levels (low demand) since vehicles must wait for a green signal regardless of the number of vehicles arriving from other directions. With AIM however, the marginal impact of a vehicle is noticeable even at low traffic levels.
- $R = 1$ (right-most data point in the right figure), presents performance that is better than applying no tolls. We believe this discrepancy also stems from the fact that AIM does not use traffic signals which contribute to spikes and oscillation (see Apendix A).

Though these results are not identical to those obtained by the CTM, their similarity still provides additional evidence of the robustness of Δ-*tolling* across different models and network topologies. On the other hand, the listed discrepancies suggest that the parameters used by Δ-*tolling* need to be re-tuned following changes in traffic flow modeling (such as changes to the intersection management policy).

7 Conclusions

This article presented Δ-*tolling*, a simple road pricing scheme which makes minimal assumptions on the traffic flow model or driver behavior. This scheme involves only two parameters, and only requires link travel time and free flow travel times measurements to set tolls. The flexibility of Δ-*tolling* was demonstrated by applying it in three very different contexts: a day-to-day pricing framework where delay is determined by link performance functions and a static equilibrium model; a within-day adaptive tolling framework using the cell transmission model for dynamic network loading, with adaptive route choice but no equilibrium; and an adaptive tolling application using a new reservation-based intersection scheme for automated vehicles, evaluated in microsimulation. In all of these cases, the Δ-*tolling* scheme was able to achieve significant benefits (measured in average travel time or social welfare) over the no-toll case, even without knowledge of the different traffic models being used, or the different assumptions on driver behavior. Benefits were seen both in small, artificial grid networks with randomized parameters as well as in larger networks representing real-world cities. We also note that Δ-*tolling* does not necessarily require a computer-controlled vehicle; it only requires computer-controlled route choice. Current smart-phone software already provides navigation to human drivers, and such software could be modified to interact with tolling systems.

The Δ-*tolling* scheme represents an advance over previously-suggested toll schemes, by not requiring any of the following assumptions: that demand is known or fixed, that roadway capacity is known or fixed, that the value of time is homogeneous; that the traffic model is known. Furthermore, Δ-*tolling* is applicable across large networks and aims to optimize social welfare. As discussed in Section 2, all previous work we are aware of makes one or more of these assumptions, all of which have significant practical implications. As a few examples, drivers are unlikely to voluntarily report all of their trips to the tolling agency (so that demand is not fully known), all traffic models are approximations to real traffic flow, and what is optimal for a single corridor may not be optimal over a larger network. Our aim in presenting Δ-*tolling* is to show that substantial benefits can be obtained even without knowing all of this information.

Δ-*tolling* is simple to implement since it requires measuring only two variables: current travel time and free flow travel time, both measurements are feasible with now days technology. It is robust to the underlying traffic model, and does admit optimality results under certain assumptions (Model A). Even when optimality is not provable, Δ-*tolling* results in significant average travel time reduction in Models B and C. The fact that Δ-*tolling* gains significant improvements over three different traffic models suggests it may be beneficial in other models as well, including real-life traffic. Our study further found that the optimized β and R values are similar across two, non static, traffic networks. Future work will examine if these values are effective across other models and scenarios. If they are not, techniques for optimizing β and R values should be developed and investigated.

There still remain many assumptions in Δ-*tolling* which future research should address. It is unrealistic to assume 100% responsiveness to tolls. Future work will investigate what ratio of responsive traffic is needed to ensure optimal system performance, and the effect of different responsiveness ratios on the system's performance. Future research will also test the Δ-*tolling* scheme for other combinations of traffic models, network topologies, and driver behavior assumptions such as mode choice, destination choice, and departure time choice.

We also plan to further investigate properties of Δ-*tolling* within Model A, whose analytical framework may determine conditions for provable convergence of Δ-*tolling* to system optimality. Other important topics for future study include developing tolling schemes which may be more politically acceptable than universally tolling network links, including opt-in systems, credit-based systems, or the use of incentives in addition to (or in lieu of) punitive tolls.

Acknowledgements

A portion of this work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by the National Science Foundation (CNS-1330072, CNS-1305287, IIS-1637736, IIS-1651089), ONR (21C184-01), AFOSR (FA9550-14-1-0087), Raytheon, Toyota, AT&T, and Lockheed Martin. Peter Stone serves on the Board of Directors of, Cogitai, Inc. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research. The authors would like to thank the Texas Department of Transportation for supporting this research under project 0-6838, Bringing Smart Transport to Texans: Ensuring the Benefits of a Connected and Autonomous Transport System in Texas. The authors would also like to acknowledge the support of the Data-Supported Transportation Operations & Planning Center and the National Science Foundation under Grant No. 1254921. Finally, the authors would like to thank the four (anonymous) reviewers of this article for helping improve it.

Appendix

A The Need for a Weighting Parameter R

In a road network, agents often plan a full origin to destination route prior to embarking. The route chosen by an agent is optimized according to current traffic conditions and tolls. The lag between decision making and the impact of these decisions might produce oscillation effects, as demonstrated in the following example.

Assume that time is discretized into time steps and that an agent planing its route at time step t will traverse the system at time step $t+1$. Consider a road network with a single origin and destination and two parallel links (highway and shortcut), each leading from origin to destination. The travel time on the highway equals

$\Psi_{hw}(x_{hw}) = 1$, while the travel time on the shortcut is $\Psi_{sc}(x_{sc}) = x_{sc}$. Assume that Δ-*tolling* is used and that at $t = 0$ the network is empty and tolls are zero on both links. Table 2 specifies the flow volume (x), travel time (T) and toll value (τ) for both the highway and shortcut at each time step.

Time	Shortcut			Highway		
step	x	T	τ	x	T	τ
0	0	0	0	0	1	0
1	1	1	1	0	1	0
2	0	0	0	1	1	0
3	1	1	1	0	1	0

Table 2: Example demonstrating the need for a weighting parameter to prevent oscillation.

At $t = 0$ the travel time is 0 on the shortcut and 1 on the highway. As a result, all agents planning at $t = 0$ will choose the shortcut. The resulting congestion would cause the shortcut's travel time to increase to 1, along with its toll. All agents planning at $t = 1$ would thus choose to travel via the highway. As a result, the shortcut will remain empty at $t = 2$ and its travel time and toll will reduce to 0. This phenomenon, where all agents choose to go via the shortcut in one time step and the highway in the next, will repeat itself, resulting in sub optimal system performance.

Moreover, if the a link performance function isn't assumed, tolls may unpredictably spike due to traffic irregularities. For instance, whenever a traffic light turns red, traffic stops and the toll set by Δ-*tolling* spikes. On the other hand, when the light turns green traffic flow resumes and the toll drops. Both of these effects can be mitigated by introducing a weighting parameter (specified as R in the model presented in this paper). Effectively, this weighting parameter accounts for the delay between route choice and realized congestion due to such route choices, by introducing a compensatory delay in the tolling system.

References

Arnott, R., De Palma, A. and Lindsey, R. (1993), 'A structural model of peak-period congestion: A traffic bottleneck with elastic demand', *The American Economic Review* pp. 161–179.

Bar-Gera, H. (2014), 'Transportation test problems', Website: `http://www.bgu.ac.il/~bargera/tntp/`. Accessed April 28, 2009.

Beckmann, M. J., McGuire, C. B. and Winston, C. B. (1956), *Studies in the Economics of Transportation*, Yale University Press, New Haven, CT.

Boyles, S. D. (2009), Operational, Supply-Side Uncertainty in Transportation Networks: Causes, Effects, and Mitigation Strategies, PhD thesis, The University of Texas at Austin.

Boyles, S. D., Gardner, L. M. and Bar-Gera, H. (2015), 'Incorporating departure time choice into high-occupancy/toll (hot) algorithm evaluation', *Transportation Research Procedia* **9**, 90 – 105. Papers selected for Poster Sessions at The 21st International Symposium on Transportation and Traffic Theory Kobe, Japan, 5-7 August, 2015.

Boyles, S. D., Kockelman, K. M. and Waller, S. T. (2010), 'Congestion pricing under operational, supply-side uncertainty', *Transportation Research Part C: Emerging Technologies* **18**(4), 519 – 535.

Carey, M. and Srinivasan, A. (1993), 'Externalities, average and marginal costs, and tolls on congested networks with time-varying flows', *Operations research* **41**(1), 217–231.

Cheng, D., Li, W. and Ishak, S. (2014), 'Accounting for travel time reliability and trip purpose in an agent-based approach to toll pricing with dynamic feedback control', *Transportation Research Record: Journal of the Transportation Research Board* **2470**, 131–141.

Chiu, Y.-C., Bottom, J., Mahut, M., Paz, A., Balakrishna, R., Waller, T. and Hicks, J. (2011), 'Dynamic traffic assignment: A primer', *Transportation Research E-Circular* (E-C153).

Daganzo, C. F. (1994), 'The cell transmission model: a dynamic representation of highway traffic consistent with the hydrodynamic theory', *Transportation Research Part B* **28**(4), 269–287.

Daganzo, C. F. (1995*a*), 'The cell transmission model, part II: network traffic', *Transportation Research Part B* **29**(2), 79–93.

Daganzo, C. F. (1995*b*), 'A finite difference approximation of the kinematic wave model of traffic flow', *Transportation Research Part B* **29**(4), 261–276.

Danielis, R. and Marcucci, E. (2002), 'Bottleneck road congestion pricing with a competing railroad service', *Transportation Research Part E: Logistics and Transportation Review* **38**(5), 379–388.

de Palma, A. and Lindsey, R. (2011), 'Traffic congestion pricing methodologies and technologies', *Transportation Research Part C: Emerging Technologies* **19**(6), 1377 – 1399.

Dial, R. B. (2006), 'A path-based user-equilibrium traffic assignment algorithm that obviates path storage and enumeration', *Transportation Research Part B* **40**(10), 917–936.

Dresner, K. M. and Stone, P. (2007), Sharing the road: Autonomous vehicles meet human drivers., *in* 'IJCAI', Vol. 7, pp. 1263–1268.

Dresner, K. and Stone, P. (2004), Multiagent traffic management: A reservation-based intersection control mechanism, *in* 'Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2', IEEE Computer Society, pp. 530–537.

Dresner, K. and Stone, P. (2006), Human-usable and emergency vehicle-aware control policies for autonomous intersection management, *in* 'Fourth International Workshop on Agents in Traffic and Transportation (ATT), Hakodate, Japan'.

Ekstrm, J., Kristoffersson, I. and Quttineh, N.-H. (2016), 'Surrogate-based optimization of cordon toll levels in congested traffic networks', *Journal of Advanced Transportation* **50**(6), 1008–1033.

Fajardo, D., Au, T.-C., Waller, S., Stone, P. and Yang, D. (2011), 'Automated intersection control: Performance of future innovation versus current traffic signal control', *Transportation Research Record: Journal of the Transportation Research Board* (2259), 223–232.

Farahani, R. Z., Miandoabchi, E., Szeto, W. Y. and Rashidi, H. (2013), 'A review of urban transportation network design problems', *European Journal of Operational Research* **229**(2), 281–302.

Friesz, T. L., Bernstein, D. and Kydes, N. (2004), 'Dynamic congestion pricing in disequilibrium', *Networks and Spatial Economics* **4**(2), 181–202.

Gardner, L., Bar-Gera, H. and Boyles, S. D. (2013), 'Development and comparison of choice models and tolling schemes for high-occupancy/toll (HOT) facilities', *Transportation Research Part B* **55**, 142–153.

Gardner, L., Boyles, S. D., Bar-Gera, H. and Tang, K. (2015), 'Robust tolling schemes for high-occupancy/toll (hot) facilities under variable demand', *Transportation Research Record* **2450**, 152–162.

Gardner, L., Boyles, S. D. and Waller, S. T. (2011), 'Quantifying the benefit of responsive pricing and travel information in the stochastic congestion pricing problem', *Transportation Research Part A* **45**, 204–218.

Gardner, L., Duthie, J., Unnikrishnan, A. and Waller, S. T. (2008), 'Robust pricing for networks with demand uncertainty'. Presented at the 87th Annual Meeting of the Transportation Research Board, Washington, DC.

Gardner, L. M., Unnikrishnan, A. and Waller, S. T. (2010), 'Solution methods for robust pricing of transportation networks under uncertain demand', *Transportation Research Part C: Emerging Technologies* **18**(5), 656 – 667. Applications of Advanced Technologies in Transportation: Selected papers from the 10th {AATT} Conference.

Grether, D., Chen, Y., Rieser, M., Beuck, U. and Nagel, K. (2008), Emergent effects in multi-agent simulations of road pricing, *in* 'Proceedings of the Annual Meeting of the European Regional Science Association (ERSA)', pp. 08–08.

Guo, R.-Y., Yang, H., Huang, H.-J. and Tan, Z. (2015), 'Link-based day-to-day network traffic dynamics and equilibria', *Transportation Research Part B: Methodological* **71**, 248–260.

Guo, R.-Y., Yang, H., Huang, H.-J. and Tan, Z. (2016), 'Day-to-day flow dynamics and congestion control', *Transportation Science* **50**(3), 982–997.

Han, D. and Yang, H. (2009), 'Congestion pricing in the absence of demand functions', *Transportation Research Part E: Logistics and Transportation Review* **45**(1), 159 – 171.

Huang, H.-J. (2002), 'Pricing and logit-based mode choice models of a transit and highway system with elastic demand', *European Journal of Operational Research* **140**(3), 562–570.

Joksimovic, D., Bliemer, M. and Bovy, P. (2005), 'Optimal toll design problem in dynamic traffic networks with joint route and departure time choice', *Transportation Research Record: Journal of the Transportation Research Board* (1923), 61–72.

Joksimovic, D., Bliemer, M. C., Bovy, P. H. and Verwater-Lukszo, Z. (2005), Dynamic road pricing for optimizing network performance with heterogeneous users, *in* 'Proceedings. 2005 IEEE Networking, Sensing and Control, 2005.', IEEE, pp. 407–412.

Levin, M. W., Pool, M., Owens, T., Juri, N. R. and Waller, S. T. (2015), 'Improving the convergence of simulation-based dynamic traffic assignment methodologies', *Networks and Spatial Economics* **15**(3), 655–676.

Lighthill, M. and Whitham, G. (1955), 'On kinematic waves. II. A theory of traffic flow on long crowded roads', *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* pp. 317–345.

Lin, D.-Y., Unnikrishnan, A. and Waller, S. T. (2011), 'A dual variable approximation based heuristic for dynamic congestion pricing', *Networks and Spatial Economics* **11**, 271–293.

Liu, H. X., He, X. and He, B. (2009), 'Method of successive weighted averages (mswa) and self-regulated averaging schemes for solving stochastic user equilibrium problem', *Networks and Spatial Economics* **9**(4), 485–503.

Lo, H. K. and Szeto, W. Y. (2005), 'Road pricing modeling for hyper-congestion', *Transportation Research Part A: Policy and Practice* **39**(7), 705–722.

Lou, Y., Yin, Y. and Laval, J. A. (2011), 'Optimal dynamic pricing strategies for high-occupancy/toll lanes', *Transportation Research Part C: Emerging Technologies* **19**(1), 64–74.

Łukasiewicza, P., Karpioa, K. and Orłowskia, A. (2012), The models of personal incomes in USA, *in* 'Proceedings of the 5th Symposium on Physics in Economics and Social Sciences', Warsaw, Poland.

Merchant, D. K. and Nemhauser, G. L. (1978), 'A model and an algorithm for the dynamic traffic assignment problems', *Transportation science* **12**(3), 183–199.

Papageorgiou, M., Hadj-Salem, H. and Blosseville, J.-M. (1991), 'Alinea: A local feedback control law for on-ramp metering', *Transportation Research Record* (1320), 58–64.

Pigou, A. C. (1920), *The Economics of Welfare*, Macmillan and Co., London.

Rambha, T. and Boyles, S. D. (2016), 'Dynamic pricing in discrete time stochastic day-to-day route choice models', *Transportation Research Part B: Methodological* **92, Part A**, 104 – 118. Special issue Day-to-Day Dynamics in Transportation Networks.

Rambha, T., Boyles, S. D., Unnikrishnan, A. and Stone, P. (2017), 'Marginal cost pricing for system optimal traffic assignment with recourse under supply-side uncertainty', In Review in *Transportation Research Part B: Methodological*.

Richards, P. (1956), 'Shock waves on the highway', *Operations Research* **4**(1), 42–51.

Simoni, M., Pel, A., Waraich, R. and Hoogendoorn, S. (2015), 'Marginal cost congestion pricing based on the network fundamental diagram', *Transportation Research Part C: Emerging Technologies* **56**, 221 – 238.

Smith, M. J. (1984), 'The existence of a time-dependent equilibrium distribution of arrivals at a single bottleneck', *Transportation Science* **18**(4), 385–394.

Tampère, C. M., Corthout, R., Cattrysse, D. and Immers, L. H. (2011), 'A generic class of first order node models for dynamic macroscopic simulation of traffic flows', *Transportation Research Part B: Methodological* **45**(1), 289–309.

Tan, Z., Yang, H. and Guo, R.-Y. (2015), 'Dynamic congestion pricing with day-to-day flow evolution and user heterogeneity', *Transportation Research Part C: Emerging Technologies* **61**, 87 – 105.

Tsekeris, T. and Voß, S. (2009), 'Design and evaluation of road pricing: state-of-the-art and methodological advances', *NETNOMICS: Economic Research and Electronic Networking* **10**(1), 5–52.

van den Berg, V. and Verhoef, E. T. (2011), 'Winning or losing from dynamic bottleneck congestion pricing?: The distributional effects of road pricing with heterogeneity in values of time and schedule delay', *Journal of Public Economics* **95**(7), 983–992.

Verhoef, E. T. (2003), 'Inside the queue:: hypercongestion and road pricing in a continuous time–continuous place model of traffic congestion', *Journal of Urban Economics* **54**(3), 531–565.

Vickrey, W. S. (1969), 'Congestion theory and transport investment', *The American Economic Review* pp. 251–260.

Waller, S. T., Kockelman, K. M., Ukkusuri, S. V. and Boyles, S. (2006), Examining the benefit of accounting for traffic dynamics in congestion pricing applications, *in* 'Transportation Research Board 85th Annual Meeting', number 06-2770.

Waller, S. T. and Ziliaskopoulos, A. K. (2002), 'On the online shortest path problem with limited arc cost dependencies', *Networks* **40**(4), 216–227.

Wardrop, J. (1952), 'Some theoretical aspects of road traffic research', *Proceedings of the Institute of Civil Engineers, Part II* pp. 325–378.

Webster, F. V. (1958), Traffic signal settings, Technical report.

Wie, B.-W. (2007), 'Dynamic stackelberg equilibrium congestion pricing', *Transportation Research Part C: Emerging Technologies* **15**(3), 154–174.

Wie, B.-W. and Tobin, R. L. (1998), 'Dynamic congestion pricing models for general traffic networks', *Transportation Research Part B: Methodological* **32**(5), 313–327.

Yang, F., Yin, Y. and Lu, J. (2007), 'Steepest descent day-to-day dynamic toll', *Transportation Research Record: Journal of the Transportation Research Board* **2039**, 83–90.

Yang, H., Meng, Q. and Lee, D.-H. (2004), 'Trial-and-error implementation of marginal-cost pricing on networks in the absence of demand functions', *Transportation Research Part B: Methodological* **38**(6), 477–493.

Yang, H., Xu, W., sheng He, B. and Meng, Q. (2010), 'Road pricing for congestion control with unknown demand and cost functions', *Transportation Research Part C: Emerging Technologies* **18**(2), 157 – 175.

Ye, H., Yang, H. and Tan, Z. (2015), 'Learning marginal-cost pricing via a trial-and-error procedure with day-to-day flow dynamics', *Transportation Research Part B: Methodological* **81, Part 3**, 794 – 807. {ISTTT} 21 for the year 2015SI: {ISTTT21}.

Yin, Y. and Lou, Y. (2009), 'Dynamic tolling strategies for managed lanes', *Journal of Transportation Engineering* **135**(2), 45–52.

Zhang, G., Wang, Y., Wei, H. and Yi, P. (2008), 'A feedback-based dynamic tolling algorithm for high-occupancy toll lane operations', *Transportation Research Record: Journal of the Transportation Research Board* (2065), 54–63.

Zheng, N., Waraich, R. A., Axhausen, K. W. and Geroliminis, N. (2012), 'A dynamic cordon pricing scheme combining the macroscopic fundamental diagram and an agent-based traffic model', *Transportation Research Part A: Policy and Practice* **46**(8), 1291–1303.

Ziliaskopoulos, A. K. and Waller, S. T. (2000), 'An internet-based geographic information system that integrates data, models and users for transportation applications', *Transportation Research Part C: Emerging Technologies* **8**(1), 427–444.