Delft University of Technology

Real-time predictive control strategy optimization

Gupta, Samarth; Seshadri, Ravi; Atasoy, Bilge; Prakash, A. Arun; Pereira, Francisco; Tan, Gary; Ben-Akiva, Moshe

DOI

[10.1177/0361198120907903](https://doi.org/10.1177/0361198120907903)

Publication date

2020

Document Version

Final published version

Published in

Transportation Research Record

Citation (APA)

Gupta, S., Seshadri, R., Atasoy, B., Prakash, A. A., Pereira, F., Tan, G., & Ben-Akiva, M. (2020). Real-time predictive control strategy optimization. *Transportation Research Record*, 2674(3), 1-11.
<https://doi.org/10.1177/0361198120907903>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Real-Time Predictive Control Strategy Optimization

**Samarth Gupta¹, Ravi Seshadri², Bilge Atasoy³, A. Arun Prakash¹,
Francisco Pereira⁴, Gary Tan⁵, and Moshe Ben-Akiva¹**

Transportation Research Record
2020, Vol. 2674(3) 1–11
© National Academy of Sciences:
Transportation Research Board 2020
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0361198120907903
journals.sagepub.com/home/trr
SAGE

Abstract

Urban traffic congestion has led to an increasing emphasis on management measures for more efficient utilization of existing infrastructure. In this context, this paper proposes a novel framework that integrates real-time optimization of control strategies (tolls, ramp metering rates, etc.) with the generation of traffic guidance information using predicted network states for dynamic traffic assignment systems. The efficacy of the framework is demonstrated through a fixed demand dynamic toll optimization problem, which is formulated as a non-linear program to minimize predicted network travel times. A scalable efficient genetic algorithm that exploits parallel computing is applied to solve this problem. Experiments using a closed-loop approach are conducted on a large-scale road network in Singapore to investigate the performance of the proposed methodology. The results indicate significant improvements in network-wide travel time of up to 9% with real-time computational performance.

Urban transportation networks are subject to a large degree of variability because of fluctuating supply and demand characteristics. These fluctuations result in the pervasive phenomena of recurrent and non-recurrent congestion, which is an escalating problem worldwide. The adverse impacts of the resulting congestion include long travel delays, high travel costs, and significant costs to the economy and the environment. Consequently, there has been an increased emphasis on developing tools to mitigate congestion and utilize existing infrastructure efficiently. In this context, the paper proposes an integrated framework—within a dynamic traffic assignment (DTA) system—to optimize network control strategies in real-time considering network state predictions. Specifically, the generated control strategies are predictive (or proactive) as opposed to being just reactive. The framework also incorporates the generation of consistent guidance—it ensures that the guidance disseminated considers the traveler's response to it, thereby increasing the reliability of the provided information. Further, the effectiveness of the proposed framework is demonstrated through a real-world application to the predictive optimization of network tolls.

The motivation for this study is fourfold. First, the need for decision support tools to facilitate a more efficient utilization of existing infrastructure. Second, most studies on optimal network control do not combine the optimization of network control strategies with the

generation of guidance information. The third motivating factor is the complexity and scale of the problem. As the objective function involves simulation, it tends to be non-linear and non-convex, making it challenging for a real-time application. Finally, the study is also motivated by important applications in real-time traffic management and incident response systems.

In view of the aforementioned motivations, the following objectives are identified: (a) to develop an integrated framework within a real-time DTA system that determines optimal control strategies and consistent guidance information considering traffic state predictions; (b) to apply the framework to the fixed demand dynamic toll optimization problem; and (c) to evaluate the proposed framework using a closed-loop approach (where the DTA system is interfaced with a traffic microsimulator that emulates stochasticity in the real world, thus providing a platform for realistic evaluation) on a large real-world network with link tolls as control strategies.

¹Massachusetts Institute of Technology, Cambridge, MA

²Singapore-MIT Alliance for Research and Technology, Singapore

³Delft University of Technology, Delft, The Netherlands

⁴Technical University of Denmark, Kongens Lyngby, Denmark

⁵National University of Singapore, Singapore

Corresponding Author:

Ravi Seshadri, ravi@smart.mit.edu

There are several salient contributions of this work. First, the proposed simulation-optimization framework simultaneously optimizes network control strategies and computes consistent guidance information based on traffic state predictions. Utilizing traffic state predictions aids in accurately evaluating the effect of control strategies. Furthermore, the control strategy at any location is determined based on global traffic state predictions and not just local predictions, thereby explicitly considering the system-level effects. The consistency in guidance ensures that the information disseminated by the traffic management center is reliable, an important issue that has been overlooked in the literature on control strategy optimization. A parallel genetic algorithm is applied to solve for the optimal control strategy (within the proposed framework) that maintains computational tractability to achieve real-time performance on a large real-world network. Second, the proposed framework is evaluated using a rigorous closed-loop approach which ensures that impacts of the control strategy are not overestimated. The experiments demonstrate the effectiveness of the proposed system which can yield travel time improvements of up to 9%, and average computational times of less than 5 min. Third, a sensitivity analysis is performed with respect to network demand levels and the consistency in guidance information is verified.

Literature Review

Although the framework presented in this paper is applicable to other control strategies, including signal timing and ramp metering, the review here focuses on real-time congestion pricing in view of the application presented. The reader is referred to Chung and Recker (1) for a review of existing toll facilities in the U.S.A. and to de Palma and Lindsey (2) for a discussion of congestion pricing technologies.

There are two broad categories of tolling strategies: fixed pricing strategies and dynamic pricing strategies. In fixed pricing strategies, the tolls are predetermined; they can be time-invariant or can vary in a predetermined manner during the day (time-of-day tolling). Further, in a fixed pricing strategy, tolls can also vary based on location and vehicle type. In dynamic pricing strategies, the tolls are continually determined based on the current/future traffic conditions and are not predetermined. A dynamic tolling strategy can be either reactive or predictive. In a reactive tolling strategy, the tolls are determined based on the current traffic conditions. In contrast, in predictive tolling, the tolls are determined considering predicted traffic states.

Yang and Huang (3) and Tsekeris and Voß (4) should be referred to for a review of work on static and fixed congestion pricing. Among the studies that determine

time-dependent and fixed pricing, de Palma et al. (5) was one of the earliest to examine the effect of time-invariant versus time-dependent pricing using a simulator. Their experiments show that time-dependent tolls can generate twice the welfare gains compared with time-invariant tolls. Xu (6) presented an optimization framework with the travel time objective and solved the problem using the simultaneous perturbation stochastic approximation (SPSA) algorithm. Chen et al. (7) solved the similar problem for a travel time objective. The problem was solved by statistically modeling the objective function (calculated from the output of DynusT) using Kriging. The same authors later extended the work to objectives of throughput and revenue (8). The tolling scheme was based on the vehicle miles traveled.

Studies on the dynamic reactive pricing have predominantly been in the context of managed-lane operations. Yin and Lou (9) propose two dynamic pricing approaches for managed toll lanes: a feedback control approach and a reactive self-learning approach. The pricing decisions are based on real-time traffic conditions and the objective is to improve the free-flow travel service on the toll lanes while maximizing total throughput. Similar approaches—based on feedback control—have been used to optimize for various other objectives like speed, travel time, delays, and revenues. There have been several studies on dynamic reactive pricing for different tolled links in a network by employing the traffic simulation software Paramics and TransModeler (10–13). The algorithm applied was from Zhang et al. (10); it is a feedback controller based on speed measurements. It was shown that dynamic tolling results in shorter queue lengths and higher speeds.

Dong et al. (14) studied a predictive tolling strategy, where the predicted traffic conditions provided by DYNASMART-X were used to generate the tolls. A feedback control approach was adopted where the toll at a location is determined by adjusting the previous toll based on the deviation of predicted concentration on the corresponding link from the desired level. Hassan et al. (15) also studied predictive tolling to maximize revenue. The toll is optimized based on a formulation where a Greenshields model is embedded to represent traffic dynamics and a binary logit model is incorporated for route choice. A linear approximation is used for the solution of the optimization model and the optimized toll is evaluated through a simulation-based DTA system (DIRECT) with prediction capabilities. They applied the tolling methodology on a synthetic corridor network with two gantries where the tolls need to be optimized. Zhang et al. (16) also study a managed-lane setting and demonstrate the impact of online calibration within a predictive toll optimization framework. More recently Hashemi and Abdelghany. (17) provided a predictive

control framework with an example of timing decisions on signalized intersections. As in (15), the authors use the simulation-based DTA system DIRECT for state estimation and prediction, and for control optimization a genetic algorithm was used, similar to that adopted in this study. They applied the methodology to the US-75 corridor in Dallas. The authors also extended the framework for robust traffic network management on a corridor network by explicitly accounting for uncertainty (18). Finally Hashemi and Abdelghany (19) proposed a deep learning methodology for real-time network management and demonstrated that the method achieves comparable travel time savings on a corridor network to the optimization-based approach they had previously proposed (17).

In summary, many studies adopt a reactive setting, that is, they do not consider the effects in future time periods while determining tolls in the current time period. This myopic tolling policy can result in undesirable and fluctuating tolls and traffic conditions. Additionally, a common approach to determine dynamic tolls is based on feedback control, where the tolls are adjusted based on either observed or predicted characteristics like speed or queues. However, as the characteristics of only the tolled links are used to determine the corresponding tolls, the system-level interactions are ignored, which makes them inefficient for large-scale networks. The approaches that utilize traffic state predictions for the optimization of tolls typically have been applied only to corridor type networks. Furthermore, consistency between the provided guidance and the resulting network conditions is not completely handled in most of the studies. Finally, the evaluation of the optimized tolls is done through the same simulator that is used to optimize the tolls. This may overestimate the network performance improvements. This study addresses these gaps in real-time predictive control systems, more specifically tolling.

Integrated Framework for Real-Time Control Strategy Optimization and Guidance Generation

This section briefly describes the proposed framework for the integrated optimization of control strategies and generation of consistent travel time guidance. For ease of exposition, the framework is illustrated using DynaMIT2.0, a simulation-based DTA system for traffic state estimation and prediction developed at the MIT Intelligent Systems Laboratory (20, 21). Note, however, that the framework is generic and applies to any real-time DTA system. The DynaMIT2.0 system is first briefly introduced, followed by a discussion of the proposed framework.

DynaMIT2.0 is composed of two core modules, state estimation and state prediction, and operates in a rolling horizon mode. During each execution cycle, the state estimation module uses a combination of historical information and real-time data from various sources (surveillance sensors, traffic information feeds, weather forecasts), first, to calibrate the demand and supply parameters of the simulator so as to replicate prevailing traffic conditions as closely as possible. The updated parameters are then utilized to estimate the state of the entire network for the current time interval. Based on this estimate of the current network state, the state prediction module predicts future traffic conditions for a prediction horizon and generates consistent guidance information (refer to Ben-Akiva et al. [20] for more details on DynaMIT) that is disseminated to travelers.

The integrated framework is summarized in Figure 1. During each execution cycle, following state estimation, the prediction based information generation and strategy optimization process is invoked. Within this process, the optimization module generates a series of control strategies (for example, network tolls, signal timings, etc.) for the prediction horizon period which are to be evaluated on the basis of a specific objective. This can include the minimization of total system travel time, maximization of consumer surplus, maximization of operator revenues, and so on. The evaluation of each control strategy involves running the state prediction module iteratively to ensure that the predicted network state is consistent with the provided guidance.

More specifically, the state prediction module (expanded in the right half of Figure 1) begins with the most recently disseminated guidance (for instance, the guidance may be in the form of network link travel times) as a trial solution. The coupled demand and supply simulators are then used to predict the network state based on the given control strategy and assumed guidance as inputs (note that the route choices of drivers change in response to the control strategy and guidance). This yields predicted network travel times which are then combined with the original guidance (using the method of successive averages [MSA]) to obtain a revised travel time guidance solution. This procedure is iteratively performed until convergence, that is, the provided travel time guidance and predicted network travel times are within a pre-specified tolerance limit ϵ_P . Once convergence is achieved, the state prediction and guidance strategy are termed “consistent” and the corresponding network state is then used by the optimization module to evaluate the objective function and search for the optimal control strategy. Following the completion of the optimization procedure, the prediction based information generation and strategy optimization process returns an optimal control strategy that is applied to the network

Figure 1. Framework for integrated guidance generation and control strategy optimization.

and consistent travel time guidance that is disseminated to travelers.

The proposed framework is demonstrated in the subsequent sections through an application to the dynamic toll optimization problem.

Formulation of the Dynamic Toll Optimization Problem

The transportation network of interest is represented as a directed graph $G(N, A)$ where N represents the set of n network nodes and A represents the set of m directed links. Let $\tilde{A} \subseteq A$ represent a subset of network links that are tolled with $\tilde{m} = |\tilde{A}|$. Consider an arbitrary time interval $[t_0 - \Delta, t_0]$ where Δ is the size of the state estimation interval (typically 5 min in real-time DTA systems). Assume that the length of the current state prediction horizon is equal to $H\Delta$ (each Δ interval within the prediction horizon is termed a prediction interval) and extends from $[t_0, t_0 + H\Delta]$. In addition, assume that the link tolls are set for intervals of size Δ (this period is referred to as the tolling interval) and that the tolling intervals are aligned with the state estimation/prediction intervals. Let $\tau^h = (\tau_1^h, \tau_2^h, \dots, \tau_m^h)$ represent the vector of link tolls for the time period $[t_0 + (h-1)\Delta, t_0 + h\Delta]$ where $h = 1 \dots H$. The vector of tolls for the current prediction horizon is thus given by $\tau = (\tau^1, \tau^2, \dots, \tau^H)$.

In real-world applications, given that the state estimation and solution of the optimization problem will require a finite computational time (assume that this is at most equal to the interval length Δ), it will not be possible to implement the optimal toll vector for the first tolling interval within the prediction horizon. Consequently, the

Figure 2. Illustration of the rolling horizon approach for toll optimization.

size of the optimization horizon is assumed to be one tolling interval less than the size of the prediction horizon and the decision variables in the optimization problem are in fact $\tau' = (\tau^2, \dots, \tau^H)$. τ^1 is set to the optimal value for the same prediction interval from the previous execution cycle (denoted by λ), so that $\tau = (\lambda, \tau')$.

This is illustrated in the example in Figure 2 for a case where $H = 3$. In execution cycle 1 (denoted by C1), the decision vector consists of the toll values $(\tau_{C1}^2, \tau_{C1}^3)$ for the prediction intervals P2 and P3. The toll vector τ_{C1}^1 is set as the optimal value from the previous execution cycle (denoted by λ_1). Subsequently, in the second execution cycle, the decision vector consists of the toll values $(\tau_{C2}^2, \tau_{C2}^3)$ and $\lambda_2 = \tau_{C1}^{2*}$, where τ_{C1}^{2*} is the optimal value of τ_{C1}^2 from execution cycle 1.

Furthermore, consider the collection of vehicles $v = 1, \dots, V$ on the network during the prediction

horizon $[t_0, t_0 + H\Delta]$. Let the travel time of vehicle ν be represented by tt^ν and the predictive travel time guidance be denoted by $\mathbf{tt}^g = (\mathbf{tt}_i^g; \forall i \in A)$, where $\mathbf{tt}_i^g$ represents a vector of time-dependent link travel times (guidance) for link i . Note that the vehicle travel times $\mathbf{tt} = (tt^\nu; \nu = 1, \dots, V)$ are a result of the state prediction module of the DTA system and cannot be written as an explicit function of the tolls and predictive guidance. We characterize the complex relationship through a function $S(\cdot)$ that represents the coupled demand and supply simulators as

$$S(\mathbf{x}^p, \boldsymbol{\gamma}^p, \mathbf{tt}^g, \boldsymbol{\tau}) = \mathbf{tt}, \quad (1)$$

where $\mathbf{x}^p, \boldsymbol{\gamma}^p$ represent the forecasted demand and supply parameters for the prediction horizon. Note also that the iterative procedure (described in the previous section) ensures consistency between $\mathbf{tt}^g$ and $\mathbf{tt}$.

It is assumed that the total network demand is fixed (inelastic) and the behavioral response of users to the tolls and predictive travel time guidance is solely through route choice which is modeled within the demand simulator of DynaMIT2.0 using a path size logit model wherein the utility of a vehicle ν on path k is given by,

$$U_k^\nu = \beta_c \tilde{\tau}_k + \beta_t \bar{t}_{t_k}^g + \log(PS_k) + C_k + \epsilon_k^\nu, \quad (2)$$

where $\tilde{\tau}_k$ is the toll on route k , $\bar{t}_{t_k}^g$ is the travel time on route k as per the guidance information (which is the sum of travel times on component links), β_c and β_t represent the cost and travel time coefficients respectively, PS_k represents the path size variable for path k , C_k represents a composite utility pertaining to additional variables including path length, number of left turns and number of signalized intersections, and ϵ_k^ν represents a random error term. Note that, first, for vehicles that do not have access to the guidance information, historical travel times are used and second, similar model structures are used for both the pre-trip and en-route choice models. The reader is referred to Ben-Akiva et al. (20) for more details.

It should also be pointed out that since the optimization is performed within a rolling horizon framework, and given that the tolls change every 5 min, it is likely that the toll values on which the driver based their pre-trip (or en-route) route choice decision are significantly different from the tolls they pay in reality. To mitigate the public opposition that may arise from this, a limit is imposed on how much the tolls can vary across successive tolling intervals on a given gantry. Thus we have,

$$\boldsymbol{\tau}^{h-1} - \boldsymbol{\delta} \leq \boldsymbol{\tau}^h \leq \boldsymbol{\tau}^{h-1} + \boldsymbol{\delta}, \quad h = 2, \dots, H, \quad (3)$$

where $\boldsymbol{\delta} = (\delta_i; \forall i \in \tilde{A})$ represents the vector of limits on the change in tolls across successive intervals.

With this background, the dynamic toll optimization problem (DTOP) in this context is formulated as a non-linear program in Equation 4. The objective function considered here is the total travel time of all vehicles on the network, but can be suitably modified to accommodate other objectives such as consumer surplus, operator revenues, or social welfare, depending on the context. The decision variables are the vector of toll values for the optimization horizon period. The constraints are the DTA system, upper and lower bounds on the toll values (denoted by vectors $\boldsymbol{\tau}_{LB}$ and $\boldsymbol{\tau}_{UB}$), and the constraints on changes in toll values across successive tolling intervals.

$$\begin{aligned} \text{DTOP : } & \min_{\boldsymbol{\tau}} \sum_{\nu=1}^V tt^\nu(\boldsymbol{\tau}') \\ \text{s.t.} & \\ & S(\mathbf{x}^p, \boldsymbol{\gamma}^p, \mathbf{tt}^g, \boldsymbol{\tau}) = \mathbf{tt}, \\ & \boldsymbol{\tau}^{h-1} - \boldsymbol{\delta} \leq \boldsymbol{\tau}^h \leq \boldsymbol{\tau}^{h-1} + \boldsymbol{\delta}, \quad h = 2, \dots, H, \\ & \boldsymbol{\tau}_{LB} \leq \boldsymbol{\tau}^h \leq \boldsymbol{\tau}_{UB}, \quad h = 2, \dots, H. \end{aligned} \quad (4)$$

In case of computational performance constraints, the dimensionality of the DTOP problem above may be significantly reduced by assuming that the vector of tolls does not change across prediction intervals within the optimization horizon. In other words, it is assumed that $(\boldsymbol{\tau}^2 = \boldsymbol{\tau}^3 \dots = \boldsymbol{\tau}^H = \bar{\boldsymbol{\tau}})$ which reduces the number of decision variables from $\bar{m}(H-1)$ to $\bar{m}$. In this case, the constraints defined by Equation 3 are replaced by,

$$\boldsymbol{\lambda} - \boldsymbol{\delta} \leq \bar{\boldsymbol{\tau}} \leq \boldsymbol{\lambda} + \boldsymbol{\delta} \quad (5)$$

Solution Algorithm

As noted earlier, since the objective function of the dynamic toll optimization problem in this context does not have a closed form and is the output of a complex simulator, evolutionary algorithms and meta-heuristics are preferable to classical gradient based approaches. Therefore, a real-coded genetic algorithm (22) is applied to solve the DTOP problem formulated in Equation 4. For more details on the solution algorithm, the user is referred to Gupta et al. (23).

To facilitate *real-time* performance, given that evaluations of different control strategies in a particular iteration are independent of each other, evaluating them in parallel significantly reduces computational time and makes the approach scalable. This study adopts a master-slave architecture using the GNU Parallel library (24). (GNU is a recursive acronym for GNUs Not Unix.) To evaluate each control strategy, a new process is launched on a different CPU. Moreover, the framework is designed to allow batch-wise evaluation of different control strategies. Specifically, during each iteration,

Figure 3. Network of expressways and major arterials in Singapore.

different control strategies can be launched as different processes on different CPUs, or they can be launched in batches of smaller size. In this batch-wise implementation, different batches can either be launched sequentially on a single cluster of CPUs or they can even be launched in parallel on multiple clusters of CPUs.

Experiments

This section discusses results from a set of experiments conducted to investigate the performance of the proposed strategy optimization approach using DynaMIT2.0 on the Singapore expressway network. The numerical experiments are conducted using a closed-loop framework, interfacing DynaMIT2.0 and MITSIMLab (MITSIM), a microscopic simulator (25). MITSIM is run concurrently with DynaMIT and mimics the real network, providing sensor counts for the current interval to DynaMIT, which in turn provides predictive guidance and tolls to MITSIM. The effect of the guidance and tolls can then be examined by extracting relevant performance measures from MITSIM, avoiding overestimation of the benefits.

The experiments were conducted on the network of major arterials and expressways in Singapore (Figure 3) which consists of 948 nodes, 1150 links, 3891 segments, 4123 origin-destination (OD) pairs, and 16 tolled links. The labels represent the links, which contain a toll gantry.

This section is organized into six parts. The first subsection discusses the setup of the closed-loop framework and calibration, the second describes the experimental design and inputs. The third section analyzes the results in relation to travel time savings and the effect of network demand, the fourth part discusses the nature of optimal tolls on a few gantries, the fifth describes the consistency of guidance information and finally, the sixth part discusses computational performance.

Closed-Loop Calibration

To set up the closed-loop environment, a two-stage calibration procedure is adopted using the w-SPSA algorithm (26). In the first stage, dynamic OD demand (for a period between 06:30 and noon), driver behavior and route choice parameters of MITSIM are calibrated by minimizing a two-component objective function. The first component is the sum of squared deviations between simulated counts and actual counts (on a set of 325 sensors for 5-min time intervals averaged across 30 weekdays in February and March 2015) obtained from the Singapore Land Transport Authority. The second component is the difference between the parameter values and *a priori* estimates. The inputs for the calibration process is a set of *a priori* parameter values and a seed OD matrix obtained from a prior calibration procedure (26). The normalized root mean square error (RMSN) in the sensor counts before and after the calibration process were 73% and 34%, respectively.

In the second stage, the historical OD matrix, supply and route choice parameters of DynaMIT2.0 are calibrated against the outputs (sensor counts on 650 network segments) generated by MITSIM. The RMSN in the sensor counts before and after the calibration process were 56% and 19%, respectively. Further, a manual calibration of segment capacities and traffic dynamics parameters was also carried out based on link travel times, and the RMSN in time-dependent link travel times after calibration was 24%.

Experimental Setup

The numerical experiments were conducted using a simulation period from 06:30 to noon, which includes the morning peak hours in Singapore. The state estimation interval (and OD demand interval) is 5 min ($\Delta = 300$ s) and the prediction horizon is 15 min ($H = 3$). The simulation period is composed of three parts: a warm-up period from 06:30 to 07:30, when no tolls are imposed, a tolling period from 07:30 to 11:00, and a post-tolling period from 11:00 to noon where again, no tolls are imposed.

The impact of the predictive toll optimization is examined against two benchmarks using the closed-loop framework described above. It is assumed that the base demand (MITSIM OD demand obtained from the closed-loop calibration) represents the historical demand or an “average” day. This demand is then perturbed to reflect day-to-day variability by sampling from a normal distribution with expected value as the base demand and a coefficient of variation of 0.2.

The first benchmark is the *no toll* scenario, where the closed-loop is simulated using the perturbed demand with zero tolls. The second scenario consists of *static*

optimum tolls. In this scenario, the optimum static tolls are first computed, which involves minimizing the total travel times for the entire simulation period (obtained from the state estimation) by implementing a single vector of tolls for the complete tolling period. The closed-loop is now simulated using the perturbed demand with the static optimum tolls. Finally, in the third scenario, the closed-loop is simulated using the perturbed demand and the *predictive optimized tolls* based on the proposed framework for predictive optimization. In all three scenarios, MITSIM receives predictive travel time guidance from DynaMIT2.0 and in turn provides sensor counts to DynaMIT2.0 every estimation interval (or execution cycle). The sensor counts are used by the online calibration process, as noted previously (for more on the online calibration methodology, the reader is referred to Prakash, Seshadri et al. [27–29]).

Further, to investigate the effect of the overall demand level, all the three aforementioned scenarios are simulated for four different demand levels: *low* (base demand reduced by 10%), *base* (closed-loop calibration as noted above), *high* (base demand increased by 10%) and *very high* (base demand increased by 20%). Note that the demands referred to here are the actual MITSIM (real-world) demands. The RMSNs between the sensor counts of DynaMIT2.0 and MITSIM range between 15% and 22.5% across the various demand scenarios (this is an outcome of the online calibration process).

For the scenarios with predictive optimization, the DynaMIT2.0 historical demand (obtained from the second stage in the closed-loop calibration) remains unchanged for all demand scenarios. For the scenarios with static optimum tolls, note that the regulator must perform the determination of the optimum tolls offline, using an estimate of historical demand. Given that different levels of actual demand (unknown to the regulator) are tested, we assume that a single computation of the static optimum tolls is performed by considering a worst case scenario where the calibrated DynaMIT2.0 historical demand is increased by 20%. In addition to the comparison with predictive optimization, this allows us also to test the robustness of the static optimum tolls to both

systematic and random variation in the actual OD demands (from historical estimates).

The performance measures are: (i) average travel times (across vehicles) for each departure time interval obtained from MITSIM, and (ii) computational time for each execution cycle of DynaMIT2.0. Note that for each scenario and demand level, the performance measures reported are averages across 10 different runs to account for stochasticity in the overall system.

A high performance computing cluster (HPCC) with 120 CPUs and 256 GB of memory was used to run the experiments. For the genetic algorithm parameters, a population size of 60 was used; probability of cross-over and mutation were 0.7 and 0.1, respectively, with a computation budget of 300 s. The number of generations may vary from interval to interval depending on the demand, that is, peak or off-peak periods.

Analysis of Travel Time

To compute and compare average time-dependent travel times across scenarios, for the entire population, all the drivers departing in a given time interval (e.g., 07:00–07:05) are identified and their average trip travel time is calculated. This process is repeated for each consecutive 5-min interval in the entire simulation period, starting from 06:30–06:35, 06:35–06:40, ..., up to 12:25–12:30. The results indicate that the use of predictive optimized tolls yields significant travel time savings over both the no toll and static optimum scenarios. The percentage improvement in travel times of the predictive optimized toll scenarios over the two benchmark scenarios for the tolling period and peak period (for all demand levels) is summarized in Table 1. The average travel times (over the tolling period) in the case of the predictive optimized tolls are lower than the static optimum and no toll cases by 9.12% and 6.74% in the base demand case. Interestingly, the static optimum is worse than the no toll case for the low, base, and high demand scenarios (see also Figure 4). This indicates that the static optimum based on historical demands is not robust when the actual demands vary significantly from the historical

Table 1. Travel Time Improvement

Demand level	% Travel time improvement				Total demand (vehicle trips)
	Tolling period		Peak period		
	No toll	Static optimum	No toll	Static optimum	
Low	3.71	5.39	7.61	6.25	275,000
Base	6.74	9.12	8.36	7.94	300,000
High	8.24	8.88	9.65	10.74	325,000
Very high	8.38	4.00	8.20	7.01	350,000

Figure 4. Travel time plots for various demand levels: (a) low, (b) base, (c) high, and (d) very high.

estimates. Note that the historical demand was scaled up by 20% when computing the static optimum and therefore, in the very high demand case where the historical estimates are closest to the actual demands, the static outperforms the no toll scenario. The percentage decrease in travel time is 5.39% and 3.71% in the low demand case, 8.88% and 8.24% in the high demand case, and 4.00% and 8.38% in the very high demand case. In addition, the percentage improvement for the peak period (between 08:00 and 09:30) is 7.94% with respect to the no toll scenario and 8.36% with respect to the static optimum scenario for the base demand case. It should be noted that in event of non-recurrent scenarios (like a special event or an incident) one would expect a significantly higher impact of the toll optimization and guidance provision.

Furthermore, for all demand cases, a standard two-sided t-test indicates that the mean travel time (for all departure time intervals within the peak period) of the predictive optimized tolling scenario has a statistically significant difference from that of the no toll/static optimum scenarios at a confidence level of $\alpha = 95\%$.

Figure 4 plots the mean travel times (shaded region represents the standard error in estimate of the mean) versus departure time interval for the three scenarios and each demand level. With regard to the effect of the

overall demand level on the improvement in travel time savings with respect to the static toll/no-toll case, the results indicate the lowest improvements (during the peak period) are attained when the congestion levels are either very low or very high. This occurs because in the *low* demand scenario the relatively uncongested state of the network reduces the impact of toll optimization. On the other hand, the severely congested network state in the very high demand scenario also reduces the possibility of alleviating congestion through the re-routing of vehicles leading once again to smaller benefits of toll optimization.

Analysis of Optimized Tolls

This section provides some examples to analyze the optimized tolls under predictive optimization with respect to static optimization. The first example is of two gantries on links 45 and 83 (see Figure 3). The optimized tolls under static and predictive strategies are presented in Figure 5. The most preferred path for one of the ODs with a very high demand during the morning peak uses these gantries (first 83 and then 45). The predictive tolls are optimized at higher values compared with the static case and this indicates that real-time predictive tolls are adjusted better with respect to demand.

Figure 5. Optimal tolls on selected links: (a) link 45, (b) link 83, (c) link 225, and (d) link 226.

Second, gantries on links 225 and 226 are optimized at lower values during the peak compared with the static strategy, as shown in Figure 5. It is observed that these gantries are used toward destinations that have very low demand in the morning peak. Predictive toll optimization is able to lower the tolls during the peak to account for lower demand values toward better travel times.

Consistency of Guidance Information

A key contribution of the proposed strategy optimization framework is that it ensures that the predicted network states (in relation to link travel times) are consistent with the guidance information provided to travelers. For a given prediction interval, we begin with the historical travel times as a trial guidance solution (in this case, the guidance is in the form of network link travel times). The coupled demand and supply simulators are then used to predict the network state based on the assumed guidance as an input (note that the route choices of drivers change in response to the guidance). A revised travel time guidance solution for the next iteration is then computed using a convex combination (MSA) of the predicted network travel times and the guidance from the current iteration. This procedure is iteratively performed until convergence, that is, the provided travel time guidance and predicted network travel times are within a pre-specified tolerance limit.

The process of achieving consistency is illustrated in Figure 6 where the mean absolute percentage error between the guidance information and the predicted link travel times are plotted as a function of the prediction iteration. The simulation period is 06:30 to noon, that is, 5.5 h, which involves $5.5 \times 12 = 66$ prediction intervals. The plots show that with as few as three or four iterations of the state prediction, a mean absolute percentage error of less than 5% is achieved in a majority of the 66 prediction intervals in the simulation.

Computational Performance

The results also indicate that the proposed solution algorithm achieves real-time performance, that is, the average computational time per execution cycle (across all demand levels) is within the 5-min time budget (less than a single state estimation interval) discussed above.

The tractable computational times are the result of three contributing factors. The first is the imposition of the constraint on the extent to which tolls on a given gantry can vary across successive tolling intervals, which significantly reduces the search space for the genetic algorithm. This ensures that a population size of 60 suffices to attain a significant reduction in travel times within a low computational time budget. Secondly, the rolling horizon approach implies that the system is re-optimized every 5 min and consequently a poor solution

Figure 6. Consistency between guidance and predicted network travel times.

in one interval can be quickly rectified or improved in subsequent intervals. This, along with the feedback from the real network to the DTA system (through the online calibration) makes the control strategy optimization framework more robust. Finally, and most importantly, the synchronous parallel evaluation of strategies in each iteration of the optimization procedure allows for evaluation of a sufficiently large number of candidate solutions.

Conclusions

This paper proposes an integrated framework that combines the optimization of network control strategies with the generation of consistent guidance information for real-time DTA systems. The efficacy of the proposed framework is demonstrated through a fixed demand dynamic toll optimization problem. Furthermore, a highly parallelizable genetic algorithm-based solution approach is adopted. Numerical experiments conducted on a large-scale real-world network (expressways and major arterials in Singapore) indicate that use of the proposed framework can yield significant network-wide travel time savings of up to 8.36% and 7.94% over the no toll and static optimum scenarios, respectively. A sensitivity analysis of demand levels further indicates that the highest improvements are attained at moderate and high demand levels. Finally, the proposed solution algorithm achieves real-time performance with a computational time of less than 5 min for each execution cycle within the rolling horizon scheme. The proposed framework and solution approach have important applications

for real-time traffic management and advanced traveler information systems.

Some directions for future research include the application of the strategy optimization framework to non-recurrent scenarios, consideration of other objectives such as consumer surplus, operator revenue, and multiple objectives; incorporation of traffic state prediction errors, and the modeling of elastic demand through trip cancellation and departure time shifts in response to tolls. The application to other network control strategies and examination of the suitability of alternative solution algorithms also promise to be interesting areas for future research.

Author Contributions

The authors confirm contribution to the paper as follows: study conception and design: S. Gupta, R. Seshadri, B. Atasoy, A. A. Prakash, G. Tan, F. Pereira, M. Ben-Akiva; analysis and interpretation of results: S. Gupta, R. Seshadri, B. Atasoy, A. A. Prakash; draft manuscript preparation: S. Gupta, R. Seshadri, B. Atasoy, A. A. Prakash. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the National Research Foundation of Singapore through the Singapore-MIT Alliance for Research and Technology's FM IRG research programme.

References

1. Chung, C. L., and W. Recker. State-of-the-Art Assessment of Toll Rates for High-Occupancy and Toll Lanes. Presented at 90th Annual Meeting of the Transportation Research Board, Washington, D.C., 2011.
2. De Palma, A., and R. Lindsey. Traffic Congestion Pricing Methodologies and Technologies. *Transportation Research Part C: Emerging Technologies*, Vol. 19, No. 6, 2011.
3. Yang, H., and H. J. Huang. *Mathematical and Economic Theory of Road Pricing*. Elsevier, Amsterdam, The Netherlands, 2005.
4. Tsekeris, T., and S. Voß. Design and Evaluation of Road Pricing: State-of-the-Art and Methodological Advances. *NETNOMICS: Economic Research and Electronic Networking*, Vol. 10, No. 1, 2009, pp. 5–52.
5. De Palma, A., M. Kilani, and R. Lindsey. Congestion Pricing on a Road Network: A Study using the Dynamic Equilibrium Simulator Metropolis. *Transportation*

- Research Part A: Policy and Practice*, Vol. 39, No. 7, 2005, pp. 588–611.
6. Xu, S. *Development and Test of Dynamic Congestion Pricing Model*. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA, 2009.
 7. Chen, X. M., L. Zhang, X. He, C. Xiong, and Z. Li. Surrogate-Based Optimization of Expensive-to-Evaluate Objective for Optimal Highway Toll Charges in Transportation Network. *Computer-Aided Civil and Infrastructure Engineering*, Vol. 29, No. 5, 2014, pp. 359–381.
 8. Chen, X. M., C. Xiong, X. He, Z. Zhu, and L. Zhang. Time-of-Day Vehicle Mileage Fees for Congestion Mitigation and Revenue Generation: A Simulation-Based Optimization Method and Its Real-World Application. *Transportation Research Part C: Emerging Technologies*, Vol. 63, 2016, pp. 71–95.
 9. Yin, Y., and Y. Lou. Dynamic Tolling Strategies for Managed Lanes. *Journal of Transportation Engineering*, Vol. 135, No. 2, 2009, pp. 45–52.
 10. Zhang, G., Y. Wang, H. Wei, and P. Yi. A Feedback-Based Dynamic Tolling Algorithm for High-Occupancy Toll Lane Operations. *Transportation Research Record: Journal of the Transportation Research Board*, 2008. 2065: 54–63.
 11. Lou, Y., Y. Yin, and J. A. Laval. Optimal Dynamic Pricing Strategies for High-Occupancy/Toll Lanes. *Transportation Research Part C: Emerging Technologies*, Vol. 19, No. 1, 2011, pp. 64–74.
 12. Jang, K., K. Chung, and H. Yeo. A Dynamic Pricing Strategy for High Occupancy Toll Lanes. *Transportation Research Part A: Policy and Practice*, Vol. 67, 2014, pp. 69–80.
 13. Morgul, E. F. *Simulation Based Evaluation of Dynamic Congestion Pricing Algorithms and Strategies*. PhD Thesis, Rutgers University Graduate School, New Brunswick, 2010.
 14. Dong, J., H. S. Mahmassani, S. Erdoğan, and C.-C. Lu. State-Dependent Pricing for Real-Time Freeway Management: Anticipatory versus Reactive Strategies. *Transportation Research Part C: Emerging Technologies*, Vol. 19, No. 4, 2011, pp. 644–657.
 15. Hassan, A., K. Abdelghany, and J. Semple. Dynamic Road Pricing for Revenue Maximization: Modeling Framework and Solution Methodology. *Transportation Research Record: Journal of the Transportation Research Board*, 2013. 2345: 100–108.
 16. Zhang, Y., B. Atasoy, A. Akkinapally, and M. Ben-Akiva. Dynamic Toll Pricing using Dynamic Traffic Assignment System with Online Calibration. *Transportation Research Record: Journal of the Transportation Research Board*, 2019. 2673: 532–546.
 17. Hashemi, H., and K. Abdelghany. Real-Time Traffic Network State Estimation and Prediction with Decision Support Capabilities: Application to Integrated Corridor Management. *Transportation Research Part C: Emerging Technologies*, Vol. 73, 2016, pp. 128–146.
 18. Abdelghany, K., H. Hashemi, and M. E. Khodayar. A Decision Support System for Proactive-Robust Traffic Network Management. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 20, No. 1, 2018.
 19. Hashemi, H., and K. Abdelghany. End-to-End Deep Learning Methodology for Real-Time Traffic Network Management. *Computer-Aided Civil and Infrastructure Engineering*, Vol. 33, No. 10, 2018, pp. 849–863.
 20. Ben-Akiva, M., H. N. Koutsopoulos, C. Antoniou, and R. Balakrishna. Traffic Simulation with DynaMIT. In *Fundamentals of Traffic Simulation* (J. Barceló, ed.), Springer, New York, 2010, pp. 363–398.
 21. Lu, Y., R. Seshadri, F. C. Pereira, A. O'Sullivan, C. Antoniou, and M. Ben-Akiva. DynaMIT2.0: Architecture Design and Preliminary Results on Real-Time Data Fusion for Traffic Prediction and Crisis Management. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, Las Palmas, Spain, pp. 2250–2255. <https://doi.org/10.1109/ITSC.2015.363>.
 22. Deb, K., A. Pratap, S. Agarwal, and T. Meyarivan. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, 2002, pp. 182–197.
 23. Gupta, S., R. Seshadri, B. Atasoy, F. Pereira, S. Wang, V. Vu, G. Tan, W. Dong, Y. Lu, C. Antoniou, and M. Ben-Akiva. Real Time Optimization of Network Control Strategies in DynaMIT2. 0. Presented at 95th Annual Meeting Transportation Research Board, Washington, D.C., 2016.
 24. Tange, O. GNU Parallel – The Command-Line Power Tool. *login: The USENIX Magazine*, Vol. 36, No. 1, 2011, pp. 42–47. <http://www.gnu.org/s/parallel>.
 25. Ben-Akiva, M., H. N. Koutsopoulos, T. Toledo, Q. Yang, C. F. Choudhury, C. Antoniou, and R. Balakrishna. Traffic Simulation with MITSIMLab. In *Fundamentals of Traffic Simulation, International Series in Operations Research & Management Science*, Vol. 145, Springer, New York, 2010, pp. 233–268.
 26. Lu, L., Y. Xu, C. Antoniou, and M. Ben-Akiva. An Enhanced SPSA Algorithm for the Calibration of Dynamic Traffic Assignment Models. *Transportation Research Part C: Emerging Technologies*, Vol. 51, 2015, pp. 149–166.
 27. Prakash, A. A., R. Seshadri, C. Antoniou, F. C. Pereira, and M. Ben-Akiva. Improving Scalability of Generic Online Calibration for Real-Time Dynamic Traffic Assignment Systems. *Transportation Research Record: Journal of the Transportation Research Board*, 2018. 2672: 79–92.
 28. Prakash, A. A., R. Seshadri, C. Antoniou, F. C. Pereira, and M. Ben-Akiva. Reducing the Dimension of Online Calibration in Dynamic Traffic Assignment Systems. *Transportation Research Record: Journal of the Transportation Research Board*, 2017. 2667: 96–107.
 29. Zhang, H., R. Seshadri, A. A. Prakash, F. C. Pereira, C. Antoniou, and M. Ben-Akiva. Improved Calibration Method for Dynamic Traffic Assignment Models: Constrained Extended Kalman Filter. *Transportation Research Record: Journal of the Transportation Research Board*, 2017. 2667: 142–153.