



Real-Time Predictive Control Strategy Optimization

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

Published in:
Transportation Research Record

Link to article, DOI:
[10.1177/0361198120907903](https://doi.org/10.1177/0361198120907903)

Publication date:
2020

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Gupta, S., Seshadri, R., Atasoy, B., Prakash, A. A., Pereira, F. C., Tan, G., & Ben-Akiva, M. (2020). Real-Time Predictive Control Strategy Optimization. *Transportation Research Record*, 2674(3).
<https://doi.org/10.1177/0361198120907903>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Real-Time Predictive Control Strategy Optimization

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

Abstract—Traffic congestion has led to an increasing emphasis on management measures for a 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 guidance generation 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 is applied to solve this problem that exploits parallel computing. 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.

Index Terms—dynamic toll optimization, dynamic traffic assignment (DTA), predictive control optimization, large-scale network, real-time traffic management

I. INTRODUCTION

URBAN transportation networks are subject to large degree of variability due to the 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 high travel delays, high travel costs, and significant costs to the economy and environment. Consequently, there has been an increased emphasis on developing tools to mitigate congestion and efficiently utilize existing infrastructure. In this context, we propose 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 travelers response to it, thereby increasing the reliability of the provided information. Further, we demonstrate the effectiveness of the proposed framework 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

S. Gupta, A. Prakash and M. Ben-Akiva are with Massachusetts Institute of Technology (MIT). {samarthg, arunprak, mba}@mit.edu

R. Seshadri is with Singapore-MIT Alliance for Research and Technology (SMART). ravi@smart.mit.edu

B. Atasoy is with Delft University of Technology (TU Delft). b.atasoy@tudelft.nl

G. Tan is with the National University of Singapore (NUS). gtan@comp.nus.edu.sg

F. Pereira is with the Technical University of Denmark. camara@dtu.dk

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: 1) To develop an integrated framework within a real-time DTA system that determines optimal control strategies and consistent guidance information considering traffic state predictions; 2) To propose a real-time solution methodology to efficiently solve for the optimal strategies under the framework proposed in objective 1; 3) To evaluate the proposed framework using a closed-loop approach (where the DTA system is interfaced with a traffic microsimulator that emulates the stochasticity in real world, thus providing a platform for realistic evaluation) on a large real-world network with link tolls as control strategies.

The salient contributions of this work are, 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. The second contribution is that we apply a highly parallelizable genetic algorithm 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. Third, we evaluate the proposed framework using a rigorous closed-loop approach that 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 minutes. In addition, a sensitivity analysis is performed with respect to network demand levels.

II. LITERATURE REVIEW

Although the framework presented in this paper is applicable to other control strategies including 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 (2011) for a review of existing toll facilities in the US [1] and to de Palma and Lindsey (2011) [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 a 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 the 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 (2005) [3] and Tsekeris and Voss (2009) [4] should be referred 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. (2005) was one of the earliest to study the effect of time-invariant vs. time-dependent pricing using a simulator [5]. Their experiments show that time-dependent tolls can generate twice the welfare gains compared to time-invariant tolls. Xu (2009) presented an optimization framework with the travel time objective and solved the problem using the SPSA (Simultaneous Perturbation Stochastic Approximation) algorithm [6]. Chen et al. (2014) solve the similar problem with the travel time objective [7]. 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. [9] has also studied distance based tolling along with elastic demand.

The studies on the dynamic reactive pricing have predominantly been in the context of managed-lane operations. Yin and Lou (2009) propose two dynamic pricing approaches for managed toll lanes: a feedback-control approach and reactive self-learning approach [10]. 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 [11]–[13]. Morgul (2010) studied dynamic reactive pricing for different tolled links in a network by employing the traffic simulation software Paramics and TransModeler [14]. The algorithm applied was from Zhang et al. (2008) [11]; it is a feedback controller based on speed measurements. It was shown that dynamic tolling results in lower queue lengths and higher speeds.

Dong et al. (2011) studied the predictive tolling strategy, where the predicted traffic conditions provided by DYNASMART-X were used to generate the tolls [15]. 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. (2013) [16] also studied predictive tolling in order 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. More recently, Hashemi and Abdelghany (2016) [17] provided a predictive control framework with an example of timing decisions on signalized intersections. They also used DIRECT for state estimation and prediction and for the optimization of the control they used genetic algorithm similar to our approach. They applied the methodology to the US-75 corridor in Dallas.

In summary, a considerable number of studies have a reactive setting, i.e., 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 the 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 and hence makes them inefficient for large scale networks. Recent studies moved to predictive tolling as some major examples are cited above. However, few studies optimize for predictive dynamic pricing strategies at a network-level, most consider corridor type networks. Furthermore, consistency between the provided guidance and the resulting network conditions is not handled fully 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 [18]. This may overestimate the network performance improvements. This study addresses these gaps in real-time predictive control systems, more specifically tolling.

III. 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 the 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 [19], [20]. However, it is noted that the framework is generic and applies to any real-time DTA system. The DynaMIT2.0 system is first very 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) to first calibrate the demand and supply parameters of the simulator so as to replicate

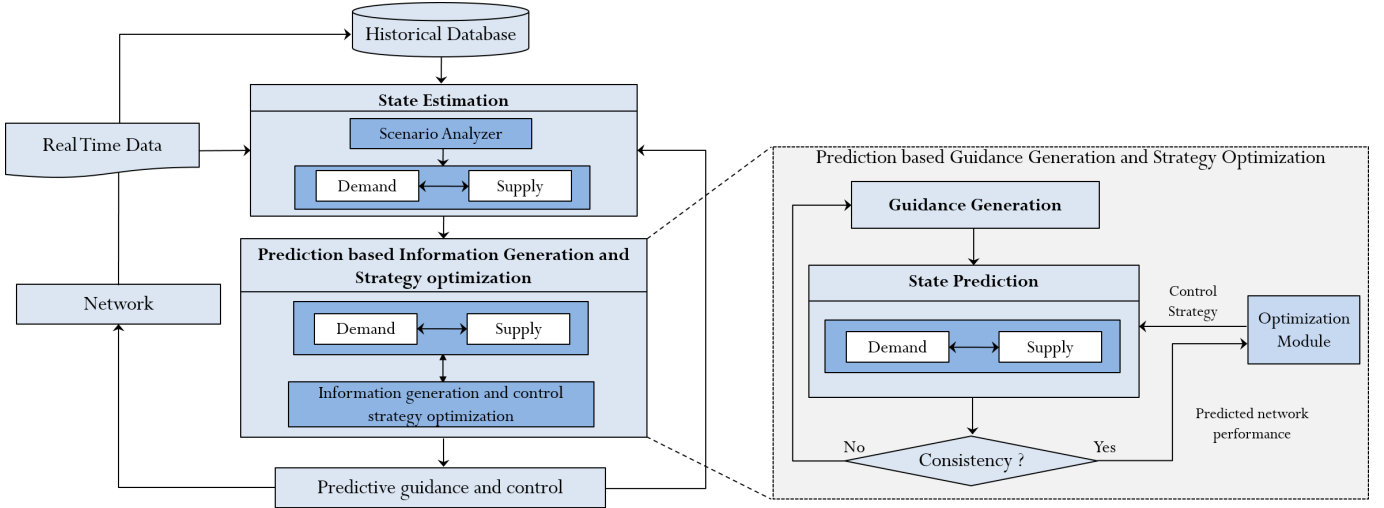


Fig. 1: Framework for Integrated Guidance Generation and Control Strategy Optimization

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 [19] for more details on the DynaMIT) that is disseminated to the 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 or MSA) to obtain a revised travel time guidance solution. This procedure is iteratively performed until convergence, i.e., 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 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.

IV. FORMULATION OF DYNAMIC TOLL OPTIMIZATION PROBLEM




Fig. 2: Illustration of the rolling horizon approach for toll optimization

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 minutes 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 $\boldsymbol{\tau}^h = (\tau_1^h, \tau_2^h \dots \tau_m^h)$ 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 $\boldsymbol{\tau} = (\boldsymbol{\tau}^1, \boldsymbol{\tau}^2, \dots, \boldsymbol{\tau}^H)$.

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 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 our optimization problem are in fact $\boldsymbol{\tau}' = (\boldsymbol{\tau}^2, \dots, \boldsymbol{\tau}^H)$. $\boldsymbol{\tau}^1$ is set to the optimal value for the same prediction interval from the previous execution cycle (denoted by $\boldsymbol{\lambda}$), so that $\boldsymbol{\tau} = (\boldsymbol{\lambda}, \boldsymbol{\tau}')$.

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 $(\boldsymbol{\tau}_{C1}^2, \boldsymbol{\tau}_{C1}^3)$ for the prediction intervals P2 and P3. The toll vector $\boldsymbol{\tau}_{C1}^1$ is set as the optimal value from the previous execution cycle (denoted by $\boldsymbol{\lambda}_1$). Subsequently, in the second execution cycle, the decision vector consists of the toll values $(\boldsymbol{\tau}_{C2}^2, \boldsymbol{\tau}_{C2}^3)$ and $\boldsymbol{\lambda}_2 = \boldsymbol{\tau}_{C1}^{2*}$, where $\boldsymbol{\tau}_{C1}^{2*}$ is the optimal value of $\boldsymbol{\tau}_{C1}^2$ from execution cycle 1.

Furthermore, consider the collection of vehicles $\nu = 1, \dots, V$ 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 $\mathbf{tt}^g = (tt_i^g; \forall i \in A)$, where tt_i^g represents a vector of the 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. Also note the iterative procedure described in Section III 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,

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

where $\tilde{\tau}_k$ is the toll on route k , \bar{tt}_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, 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 [19] for more details.

It should be also be pointed out that since the optimization is performed within a rolling horizon framework and given that the tolls change every five minutes, it is likely that the toll values on which the driver based his pre-trip (or en-route) route choice decision are significantly different from the tolls he pays in reality. To mitigate the public opposition that may arise from this, we impose a limit 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 in our 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.

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

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, we assume that $(\boldsymbol{\tau}^2 = \boldsymbol{\tau}^3 \dots = \boldsymbol{\tau}^H = \bar{\boldsymbol{\tau}})$ which reduces the number of decision variables from $\tilde{m}(H-1)$ to \tilde{m} . In this case, the constraints defined by Equation 3 are replaced by,

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

V. SOLUTION ALGORITHM

As noted earlier, since the objective function of the dynamic toll optimization problem in our 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. Hence, a real-coded

Genetic Algorithm (GA) [21] is applied to solve the DTOP problem formulated in Section IV.

The algorithm starts by randomly generating a set of control strategies or individuals ($CS_i^p, i = 1 \dots N, p : \text{parent}$) collectively known as the initial population or parent population of size N . Each single control strategy CS_i^p comprises the vector of tolls $\tau' = (\tau^2, \dots, \tau^H)$ for the current optimization horizon. These N different control strategies $CS_i^p, i = 1 \dots N$ (or individuals) are evaluated (i.e the objective function value is computed) in parallel by running the state prediction module of DynaMIT2.0 ($PD_i, i = 1 \dots N$) independently for each control strategy CS_i^p . The modularity of DynaMIT2.0 provides the functionality to execute a single state estimation run followed by multiple parallel state predictions with different control strategies and makes real time optimization (within the budget of 5 minutes) computationally possible.

Different control strategies (or individuals) in the parent population are assigned a rank $R_i, i = 1 \dots N$ based on their respective objective function values $Obj_i, i = 1 \dots N$. The objective function value for each individual is calculated from the output of the state prediction module PD_i . From this set of control strategies (individuals) in the parent population, a new set of control strategies (new individuals), collectively called the child population of size N is generated using genetic operators, i.e., the SBX crossover and polynomial mutation. The newly generated N control strategies (or individuals) of the child population ($CS_i^c, i = 1 \dots N, c : \text{child}$) are evaluated in parallel by running the state prediction module (on the same estimated state used to evaluate parent population) to get their objective function values.

Strategies in the child population $CS_i^c, i = 1 \dots N$ and parent population $CS_i^p, i = 1 \dots N$ are merged together to form a combined set of strategies (mixed population) of size $2N, CS_k^{p+c} (k = 1 \dots 2N)$ and are ranked based on their objective function values. From these $2N$ different strategies (or individuals) CS_k^{p+c} , the best N strategies (or individuals) are selected based on their rank to form the parent population for the next iteration. The procedure of generating a new set of strategies (or child population) continues uptill the termination criteria are met. The termination criteria can be a predefined number of iterations/generations G_{max} , a threshold for the improvement in the objective function value or a computational time budget T_{max} .

In order to facilitate *real time* performance, it is imperative to compute parallel tasks in a computationally efficient manner using a multi-core architecture. In the context of our framework, evaluation of different control strategies CS_i s in a particular iteration are independent of each other, therefore evaluating them in parallel significantly reduces computational time and makes the approach scalable. Different parallel computing architectures and libraries have been proposed in the literature, but for our application we adopt a Master-Slave architecture using the GNU¹ Parallel library [22]. The main benefit of using GNU Parallel is that processor level parallelism can be achieved. To evaluate each control strategy, a new process is launched on a different CPU. Managing and

scheduling of different processes is taken care by the GNU library at the operating system level. In practice, it is observed that even with processor level parallelism the speed-ups are not linear even if no inter-process communication is present. Therefore, the framework is designed so as to allow *Batch-Wise* evaluation of different control strategies. Specifically, during each iteration, all N different control strategies can be launched as different processes on different CPUs, or they can be launched in batches of size $n, n < N$. 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. This framework exploits both parallel as well as distributed computing simultaneously as shown in Figure 3.




Fig. 3: Genetic Algorithm with parallel evaluation of population using parallel & distributed computing techniques.

VI. 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 [23]. 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 (see Figure 4). The effect of the

¹GNU is a recursive acronym for GNU's Not Unix.

Algorithm 1 Control Strategy Optimization Algorithm

- 1) Store the output of the *State Estimation* module for interval $[t_0 - \Delta, t_0]$
 - 2) Replicate the cached estimated state to N different DynaMIT clones
 - 3) Initialize N different control strategies to form initial population
 - 4) Evaluate N different control strategies by running *State Prediction* module iteratively for interval $[t_0, t_0 + H\Delta]$ using N different DynaMIT clones in parallel
 - 5) Assign rank to N individuals on the basis of their Objective Function value.
 - 1: **for all** $g \leftarrow 2$ to $G_{max} \cap \text{clock-time} \leq T_{max}$ **do**
 - a) Generate N new control strategies using tournament selection with SBX crossover to form child population
 - b) Mutate N newly generated control strategies using polynomial mutation
 - c) Evaluate N child strategies by running *State Prediction* module for interval $[t_0, t_0 + H\Delta]$ using N different DynaMIT clones in parallel
 - d) Merge the child population and parent population strategies to form mixed population of size $2N$
 - e) Assign rank to $2N$ individuals on the basis of their Objective Function value
 - f) Select the best N individuals (strategies) to form the parent population
 - 2: **end for**
 - 6) Select the best Control Strategy CS_{best} from the final set of strategies
-

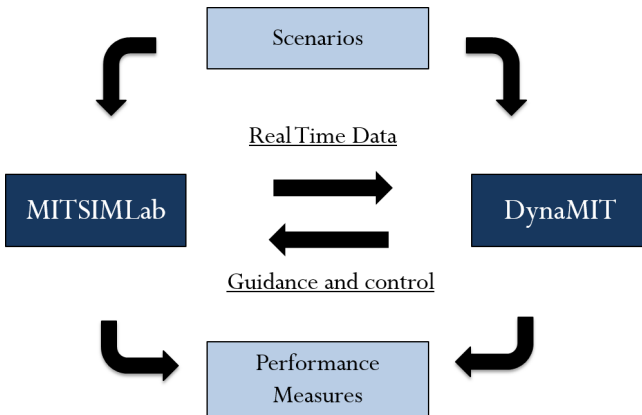


Fig. 4: Closed-Loop Framework

guidance and tolls can then be examined by extracting relevant performance measures from MITSIM avoiding overestimation of the benefits.

The experiments are conducted on the network of major arterials and expressways in Singapore (Figure 5) which consists of 948 nodes, 1150 links, 3891 segments, and 4123 origin-destination (OD) pairs, and 16 tolled links. The labels represent the links where there is a toll gantry.



Fig. 5: Network of Expressways and Major Arterials in Singapore

The section is organized into five parts. The first part 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 terms of travel time savings and the effect of network demand, fourth part discusses the optimal tolls through few gantries, and finally the fifth part discusses computational performance.

A. Closed-Loop Calibration

In order to set up the closed-loop environment, a two stage calibration procedure is adopted using the w-SPSA algorithm [24] (for other approaches see [25]). In the first stage, dynamic OD demand (for a period between 06:30 AM and 12:00 PM), 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 minute time intervals averaged across 30 weekdays in February and March 2015) obtained from the Singapore Land Transport Authority (LTA). The second component is the difference between the parameter values and apriori 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 [24]. The normalized root mean square error 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 normalized root mean square error in the sensor counts before and after the calibration process were 56% and 19% respectively. Further, the RMSN in time-dependent link travel times after calibration was found to be 24%. The results of the second stage of the calibration are summarized in Figures 6a and 6b which show scatter plots of the simulated (DynaMIT) versus actual (MITSIM) sensor counts and link travel times respectively.

B. Experimental Setup

The numerical experiments are conducted using a simulation period from 6:30 AM to 12:00 PM which includes the

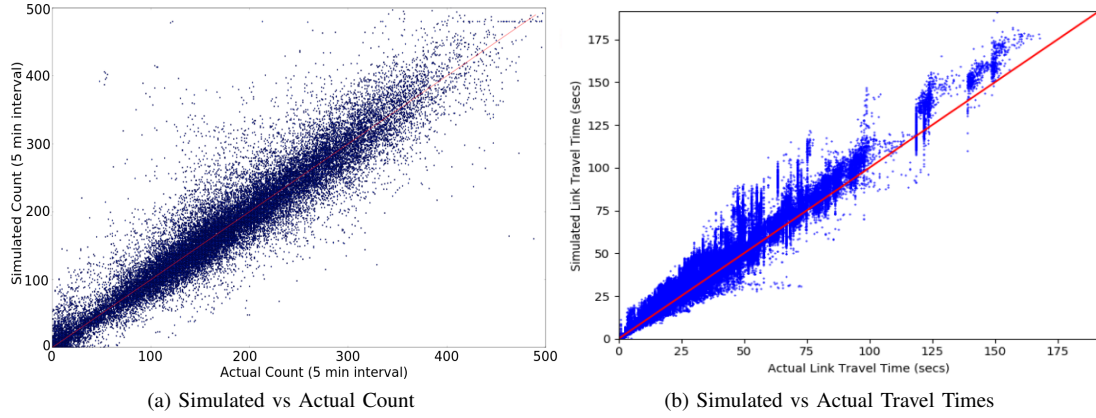


Fig. 6: Closed Loop Calibration

morning peak in Singapore. The state estimation interval (and OD demand interval) is five minutes ($\Delta = 300$ seconds) and the prediction horizon is 15 minutes ($H = 3$). The simulation period is composed of three parts: a *Warm-up* period from 6:30-7:30 AM where no tolls are imposed, a *tolling period* from 7:30 - 11:00 AM, and a *post-tolling* period from 11:00 AM to 12:00 PM where again no tolls are imposed.

The impact of the predictive toll optimization is examined against two benchmarks using the closed-loop framework described earlier. 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 optimized tolls*. In this scenario, we first compute the optimum static tolls 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 in Section III. 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).

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 earlier), *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. For the scenarios with predictive optimization, the DynaMIT2.0 historical demand (obtained from the second stage in the closedloop calibration) remains unchanged for all demand

scenarios. For the scenarios with the 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 to also 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: 1) average travel times (across vehicles) for each departure time interval obtained from MITSIM, 2) 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 is used to run the experiments. For the parameters of GA, we use a population size of 60, probability of cross-over and mutation as 0.7 and 0.1 respectively with a computation budget of 300 seconds. The number of iterations may vary from interval to interval depending on the demand, i.e., peak or off-peak periods.

C. Analysis of Travel Time

In order 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, i.e., starting from 6:30-6:35, 6:35-6:40, ..., 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 I. 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 Figures 7 to 10). This indicates that the static optimum based on historical demands is not robust when the actual demands vary significantly from the historical estimates. Note that the historical demand was scaled up by 20% when computing the static optimum and hence, 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 8:00 am and 9:30 am) 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 $\alpha = 95\%$.

Figures 7 to 10 plot 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/no-toll 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 the toll optimization.

The probability density and cumulative density functions of vehicle travel time are plotted in Figure 11 for the no toll and predictive optimized toll scenarios in the base and high demand cases. The plots highlight the reduction in frequency of trips with higher travel time due to the predictive optimization of tolls.

D. Analysis of Optimized Toll

Here we provide few examples in order to analyze the optimized tolls under predictive optimization with respect to static optimization.

First, we give an example of two gantries on links 45 and 83. We present the optimized tolls under static and predictive strategies in Figure 12 and 13. 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

TABLE I: Travel Time Improvement

Demand Level	% Travel Time improvement				Total Drivers Simulated
	Tolling Period		Peak Period		
	No Toll	Static	No Toll	Static	
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

are optimized at higher values compared to the static case and this indicates that real-time predictive tolls are adjusted better with respect to demand.




Fig. 12: Gantry on link 45




Fig. 13: Gantry on link 83




Fig. 14: Gantry on link 225




Fig. 15: Gantry on link 226

Second, gantries on links 225 and 226 are optimized at lower values during the peak compared to static strategy as shown in Figure 14 and 15. It is observed that these gantries are used towards destinations that have very low demand in the morning peak. Predictive toll optimization is able to lower the tolls during the peak in order to account for lower demand values towards better travel times.




Fig. 7: Travel Time for Low Demand




Fig. 8: Travel Time for Base Demand




Fig. 9: Travel Time for High Demand




Fig. 10: Travel Time for Very High Demand

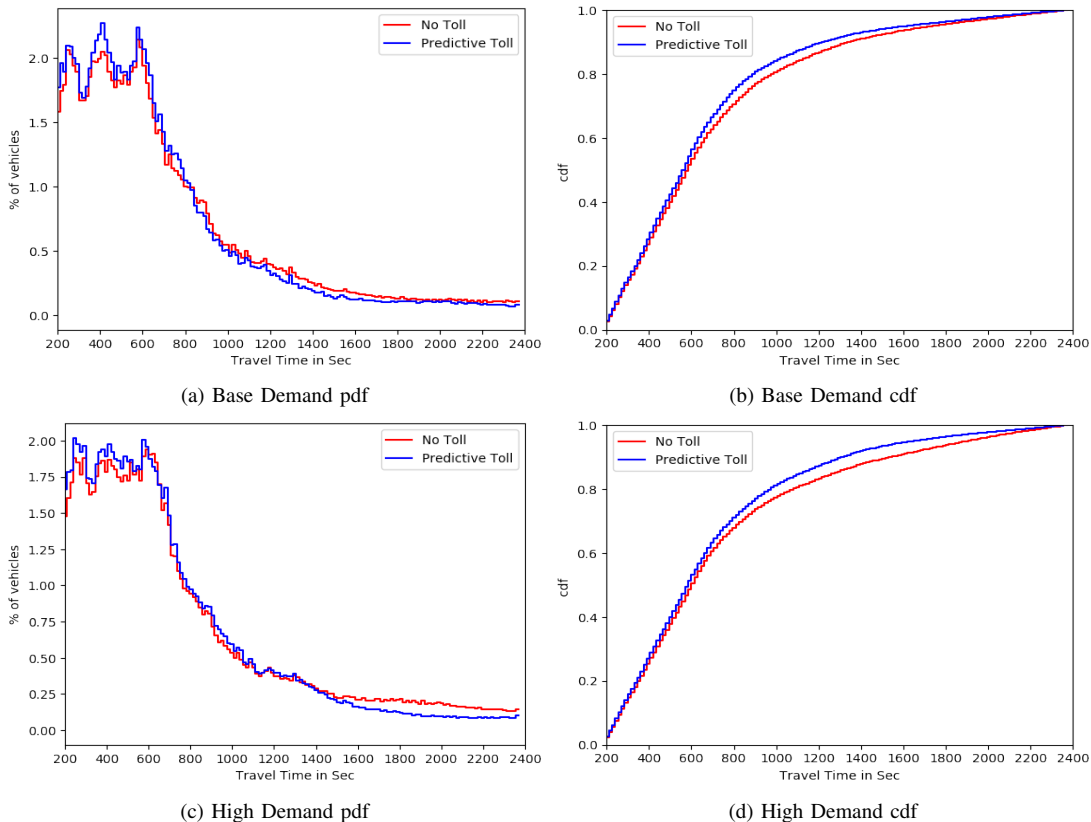


Fig. 11: Peak Period Travel Time Distributions

E. Computational Performance

The results also indicate that the proposed solution algorithm achieves real-time performance, i.e. the average computational time per execution cycle (across all demand levels) is within the five minute time budget (less than a single state estimation interval) discussed in Section IV. The plot of average computational time versus time interval is shown in Figure 16.

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 GA. 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 five minutes and consequently a poor solution 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.

VII. CONCLUSION

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 indicate 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 minutes 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 under non-recurrent scenarios including crisis-management [26], consideration of other objectives such as consumer surplus, operator revenue and multiple objectives; incorporation of traffic state prediction errors [27] and the modeling of elastic demand through trip




Fig. 16: Computational Performance

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.

ACKNOWLEDGMENT

We would like to thank the Singapore Land Transport Authority (LTA) ITS division, for providing the data for this project. This work was supported by the National Research Foundation of Singapore (SMART program).

REFERENCES

- [1] C.-L. Chung and W. Recker, "State-of-the-art assessment of toll rates for high-occupancy and toll lanes," in *Transportation Research Board 90th Annual Meeting*, no. 11-1331, 2011.
- [2] A. de Palma and R. Lindsey, "Traffic congestion pricing methodologies and technologies," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 6, pp. 1377–1399, 2011.
- [3] H. Yang and H.-J. Huang, *Mathematical and economic theory of road pricing*, 2005.
- [4] T. Tsekeris 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, pp. 5–52, 2009.
- [5] A. De Palma, 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, pp. 588–611, 2005.
- [6] S. Xu, "Development and test of dynamic congestion pricing model," Master's thesis, Massachusetts Institute of Technology, 2009.
- [7] X. M. Chen, 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, pp. 359–381, 2014.
- [8] X. M. Chen, 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, pp. 71–95, 2016.
- [9] V.-A. Vu, H. Hashemi, R. Seshadri, S. Gupta, G. Tan, A. A. Prakash, and M. Ben-Akiva, "Predictive distance-based toll optimization under elastic demand for real-time traffic management," *Transportation Research Board 97th Annual Meeting*, 2018.
- [10] Y. Yin and Y. Lou, "Dynamic tolling strategies for managed lanes," *Journal of Transportation Engineering*, vol. 135, no. 2, pp. 45–52, 2009.
- [11] G. Zhang, 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*, no. 2065, pp. 54–63, 2008.
- [12] Y. Lou, 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, pp. 64–74, 2011.
- [13] K. Jang, K. Chung, and H. Yeo, "A dynamic pricing strategy for high occupancy toll lanes," *Transportation Research Part A: Policy and Practice*, vol. 67, pp. 69–80, 2014.
- [14] E. F. Morgul, "Simulation based evaluation of dynamic congestion pricing algorithms and strategies," Ph.D. dissertation, Rutgers University-Graduate School-New Brunswick, 2010.
- [15] J. Dong, 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, pp. 644–657, 2011.
- [16] A. Hassan, K. Abdelghany, and J. Semple, "Dynamic road pricing for revenue maximization: Modeling framework and solution methodology," *Transportation Research Record*, vol. 2345, pp. 100–108, 2013.
- [17] H. Hashemi and K. F. Abdelghany, "Real-time traffic network state estimation and prediction with decision support capabilities: Application to integrated corridor management," *Transportation Research Part C*, vol. 73, pp. 128–146, 2016.
- [18] S. Gupta, R. Seshadri, B. Atasoy, F. C. Pereira, S. Wang, V.-A. Vu, G. Tan, W. Dong, Y. Lu, C. Antoniou, and M. Ben-Akiva, "Real time optimization of network control strategies in dynamit2.0," 2016.
- [19] M. Ben-Akiva, H. N. Koutsopoulos, C. Antoniou, and R. Balakrishna, "Traffic simulation with dynamit," in *Fundamentals of Traffic Simulation*, J. Barceló, Ed. New York, NY, USA: Springer, 2010, pp. 363–398.
- [20] Y. Lu, R. Seshadri, F. 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*, Sept 2015, pp. 2250–2255.
- [21] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 2, pp. 182–197, Apr 2002.
- [22] O. Tange, "Gnu parallel - the command-line power tool," *login: The USENIX Magazine*, vol. 36, no. 1, pp. 42–47, Feb 2011. [Online]. Available: <http://www.gnu.org/s/parallel>
- [23] M. Ben-Akiva, H. Koutsopoulos, T. Toledo, Q. Yang, C. Choudhury, C. Antoniou, and R. Balakrishna, "Traffic simulation with mitsimlab," in *Fundamentals of Traffic Simulation*, ser. International Series in

- Operations Research & Management Science. Springer New York, 2010, vol. 145, pp. 233–268.
- [24] L. Lu, X. Yan, C. Antoniou, and M. Ben-Akiva, “W-spsa: An enhanced spsa algorithm for the calibration of dynamic traffic assignment models,” *Transport Research Part C*, Accepted - Forthcoming.
 - [25] B. Ciuffo and C. L. Azevedo, “A sensitivity-analysis-based approach for the calibration of traffic simulation models,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 3, pp. 1298–1309, 2014.
 - [26] S. N. Hetu, S. Gupta, V.-A. Vu, and G. Tan, “A simulation framework for crisis management: Design and use,” *Simulation Modelling Practice and Theory*, vol. 85, pp. 15 – 32, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1569190X18300273>
 - [27] F. C. Pereira, C. Antoniou, J. A. Fargas, and M. Ben-Akiva, “A metamodel for estimating error bounds in real-time traffic prediction systems,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 3, pp. 1310–1322, 2014.