TRIPOD: SUSTAINABLE TRAVEL INCENTIVES WITH PREDICTION, OPTIMIZATION AND PERSONALIZATION

Carlos Lima Azevedo[1]
email: cami@mit.edu

Ravi Seshadri[2]
email: ravi@smart.mit.edu

Song Gao[3]
email: sgao@umass.edu

Bilge Atasoy[1]
email: batasoy@mit.edu

Arun P. Akkinepally[1]
email: arunprak@mit.edu

Eleni Christofa[3]
email:echristofa@engin.umass.edu

Fang Zhao[2]
email:fang.zhao@smart.mit.edu

Jessika Trancik[1]
email:trancik@mit.edu

Moshe Ben-Akiva[1]
email: mba@mit.edu

[1] Massachusetts Institute of Technology, Cambridge, MA 02139, US
[2] Singapore-MIT Alliance for Research and Technology, 138602, Singapore
[3] University of Massachusetts, Amherst, MA 01003, US

Word Count: 5984 words + 6 figure(s) x 250 + 1 table(s) x 250 = 7734 words

Submission Date: November 16, 2017

ABSTRACT

In this paper we present Tripod, a smartphone-based system to influence individuals' real-time travel decisions by offering information and incentives with the objective of optimizing system-wide energy performance. When starting a trip, travelers can access Tripod's personalized menu via a smartphone app and are offered incentives in the form of tokens for a variety of energy-reducing travel options in terms of: route, mode, ride-sharing, departure time, driving style and trip making. Options are presented with trip information to help travelers understand the real-time mobility attributes and energy consequences of their choices. By accepting and executing a specific travel option, a traveler earns tokens that depend on the system-wide energy savings created. This encourages them to consider not only their own energy cost, but also the impact of their choice on the system. Earned tokens could then be redeemed for services and goods from participating vendors and transportation agencies.

Little new infrastructure would need to be created - millions of travelers already have smartphones capable of meeting the needs of the proposed system. Our proposed incentive-based system can be implemented for a small fraction of the cost of infrastructure-heavy projects, and at little or no cost to local, state and federal governments. We also showcase the current state of development of Tripod with numerical experiments for the Downtown Boston (Boston Proper) area and for a limited set of incentivized choices, with system-wide energy savings between 3 and 8%.

Keywords: incentives, system optimization, energy, real-time optimization, prediction, personalization

INTRODUCTION

Inefficiency in transportation systems is an issue affecting millions of people every day (see for example (*1*)). It causes increased use of energy and pollution, as well as economic loss, stress, and road rage. State and local transportation departments are highly motivated to manage transportation in a more cost-effective manner, but lack the tools to do so. Real-time demand management has become one of the potential new paths towards more efficient networks, especially with the increased capabilities of information and communication technologies, computation power and mobility-centered devices and software.

From the real-time traffic management perspective, externalities such as congestion and vehicular emissions have been historically addressed with information provision, pricing and, more recently, incentives and quantity control.

Real-time information provision has been widely documented in the literature and both the behavioral, network conditions modeling and optimization and implementation efforts were extensively reported (see for example the review in (*2*) and (*3*)).

Road pricing has received a lot of attention since the seminal paper by (*4*), both from the theory to the actual implementation. In the context of a congested networks, various mathematical models and algorithms have been proposed for optimizing tolls towards better road network performance (see (*5*) or (*6*) for comprehensive reviews). Within pricing strategies, dynamic pricing solutions were studied much more recently. Gallego et al. (*7*) first introduced its optimization formulation and proposed one solution by maximizing revenue under stochastic aggregated demand. Since then, several frameworks have been proposed ((*8*)) but only a few actual implementations have been tested ((*9*)). However, one of the main concerns with road pricing is how it is perceived by the traveler: an unfair additional cost or just another tax. If equity represents one important factor when considering the pricing strategies, their high implementation costs also justify the low applications currently observable worldwide.

Recently, quantity control strategies have been under the spotlight. Congested roads are in the end a scarcity problem and this can be dealt with either a price instrument, a quantity instrument or a combination of both. Quantity control has been implemented most commonly through vehicle quota or usage restrictions. More recently, quantity control has been thought as the provision/distribution of a limited number of credits to be used when traveling on a congested link/path where the capacity is limited. Typically, the equilibrium price of credits is endogenous and clears supply and demand. Given the unpopularity of pricing, the direct redistribution of value that characterizes quality control mechanism is therefore appealing. Yang and Wang (*10*) examine an alternative simple but forceful tradable credit distribution and charging scheme, where credits are universal for all links but link-specific in the amount of credit charge. Lahlou and Wynter (*11*) consider tradable credits scheme for optimizing the performance of a bi-modal network using atomic game framework so as to model explicitly the exchange process across users. Efficient Nash equilibria existence is shown together with the compliance of potential regulator policies. Different market designs for tradable credit schemes have also been analyzed. Nie (*12*) for example, examines the effects of transaction costs on two types of tradable credit schemes: an auction market and a negotiated market. DePalma et al. (*13*) and Akamatsu et al. (*14*) presents a methodology to compare pricing and tradable credit schemes under stochastic demand and prove the increased efficiency of the latter. Similarly, Dogterom et al. (*15*) recently focused on the behavioral responses to different tradable credit schemes that target personal travel and point behavioral economics and cognitive psychology as domains to explore under this research stream. While existing literature clearly identifies the benefits of tradable credit schemes (*6, 16–18*), the extension of such schemes to multi-modal systems with different users and transportation providers has been hampered by its design and technological implementation complexity.

Finally, incentive-based demand management strategies are gaining increasing attention because they are generally considered more acceptable by the traveling public and policy makers. However, while congestion charging has been heavily studied, relatively little literature focus on incentives, how these can be deployed for real-time travel demand management and how the traveler would react to it. Leblanc et al. (*19*) carried out a stated-preference survey in the San Francisco Bay Area to analyze how different incentive schemes can affect commuting decisions. As predicted by behavioral economics, travelers were found to be much more sensitive to charges than to rewards, but also different sensitivities were found towards different incentives, e.g.: cash rewards proved to be more efficient than an HOV pass. Such evaluations are fundamental in the design and optimization of potential real-time incentive schemes which may, on the other hand, benefit from personalization to account for such sensitivities. Similarly, Kumar et al.(*20*) presented a detailed behavior analysis and modeling effort aimed at understanding how incentives affected traveler choices by using data collected from the Spitsmijden reward-based experiment. The Spitsmijden (*21*) experiment was a peak rewarding project in four highway corridors in The Netherlands to investigate the behavioral responses of personal vehicle users towards static incentives targeting departure time behavioral shifts. This peak rewarding experiment highlights the importance to design revenue-neutral incentive-based mechanism to ensure a sustainable outcome. Other experiments followed to explore the behavioral reaction to point-based, lottery-based, personalized or smartphone-based static incentives ((*22*), (*23*), (*24*) and (*25*)). All these proposed schemes are fundamental in capturing different behavioral shifts but are limited in effectively managing demand in real time.

Optimization mechanisms focusing on incentives have been documented only very recently. Rey et al. (*26*) evaluated a lottery-based revenue-neutral incentive mechanism to reduce the congestion in urban transportation systems by promoting public transit usage during off-peak periods. They "derived the theoretical equilibrium for this decision-making game and test the validity of the proposed mechanism through monetized laboratory experiments". Hu et al. (*27*) on the other hand, proposed a system that considers traffic conditions in a real-time routing guidance app with a point-based scheme. Alternatives with higher points (that can be exchanged for rewards), such as traveling during off-peak times or less congested routes, are presented to the user for network congestion mitigation. A pilot study was deployed in Los Angeles, California, in 2013 and results showed behavior shifts but overall network performance measures were not presented. Gao et al. (*28*) proposed, RoadRunner, an in-vehicle app for quantity control without costly roadside infrastructure. Harnessing vehicle-to-vehicle communications, the number of vehicles in a network is managed through a *token*-based exchange mechanism. Although the exchange and control formulation were not fully explored in this study, both software and hardware were successfully tested in field experiments and network performance assessment was carried out in simulated environment.

In summary, the theory has vastly showed us that shifting travelers to decisions that improve network efficiency, such as off-peak trip making, less congested route choice or transit and shared ride usage, would lead to considerable savings in externalities. Along with pricing and quantity control strategies, incentives have been proposed by the transportation economics field to achieve this goal and theoretical proofs have been presented. However, the design, implementation and operation of such strategies is still a challenging question. Smartphone apps are increasingly used to shift travel behavior for more efficient decision making (*29*). Many apps with static incentives (such as (*30*), (*31*) and (*32*)) can now be found in app stores for different environments, modes and targeted users. Yet, real-time incentive schemes optimization that account for predicted network conditions are still rare, targeting a limited population segment and/or limited set of traveler's choice dimensions ((*33*)).

In this paper we focus our attention on a new real-time incentive schemes, Tripod, to maximize energy savings at the (multi–modal) network level. When starting a trip, travelers can access Tripod's personalized menu via a smartphone app and are offered incentives in the form of tokens for a variety of energy-reducing travel options in terms of: route, mode, ride-sharing, departure time, driving style and actual trip making. Options are presented with information to help travelers understand the energy and emissions consequences of their choices. By accepting and executing a specific travel option, a traveler earns tokens that depend on the system-wide energy savings she or he creates, encouraging them to consider not only their own energy cost, but also the impact of their choice on the system. Tokens can then be redeemed for services and goods from participating vendors and transportation agencies. Tripod's novel design and architecture 1) leverages the existing incentive and tradable credits schemes in the literature and proposes a new solution to optimize incentives in real-time under a fix budget; 2) is a new full system architecture that handles incentive provision and rewarding without new infrastructure; 3) integrates prediction and personalization in the optimization framework 4) optimizes the entire multi-modal system for system-wide energy savings. Tripod is still under development and this paper presents the general architecture and control logic of this innovative system. We also showcase the current state of development with numerical experiments for the Downtown Boston area (Boston Proper) and for a limited set of incentivized choices.

TRIPOD'S RATIONALE

In the previous section a high-level concept of Tripod was introduced. Important premises still need to be introduced before the architecture design, implementation and solution formulation are described in more detail.

Tripod and traveler's decision making

- a traveler is assumed to make mobility choices at the origin and en-route. A mobility choice is defined by four dimensions: mode (inc. shared modes), departure time, route, and driving style (if driving). These are the dimensions incentivized by Tripod along with trip cancellation. The latter is not covered in this paper as both its behavioral and optimization implications require considerable modifications to the proposed approach;
- while en-route, a traveler cannot change mode or departure time, but can change route or/and driving style;
- only opt-in travelers (subset of the travelers population) can access the Tripod app and make a trip request;
- for each user, trip requests can happen both at the pre-trip and en-route level and form control points from whom high fidelity personal data is available through the Tripod app;
- at the Tripod system level, requests can be received at any point in time;
- Tripod incentives are provided in the form of *tokens* for each alternative in the personalized trip menu for a given requested trip;
- if a user decides to select one of the Tripod menu options, the Tripod app will track the

user during the trip and *tokens* are provided only if the realized trip coincides with the selected option;

Tripod tokens and rewards

- *tokens* are awarded to a menu option proportional to the system-wide energy saving, which encourages opt-in travelers to consider not only their own energy savings but also the impact of their choice on the system;
- *tokens* are not allocated directly to specific users, but are awarded to any menu alternative selection and execution that contributes to the optimization of system-wide energy savings;
- the marginal energy saving for each alternative presented to the Tripod user is calculated based on predicted traffic conditions for the requested trip and user specific preferences;
- the internalization of the marginal cost through *tokens* will potentially drive the system towards optimum;
- *tokens* can be redeemed for services and goods (rewards) at participating vendors and agencies. The available services and goods are assumed to be static and can be translated to a set of monetary values of reference for each reward;
- the actual *token* value can be perceived differently by each user;
- *tokens* can be accumulated; when a user has enough tokens to purchase at least one reward, the user can still save or accumulate more *tokens*;
- for practicality we assume that spending the *tokens* on rewards does not increase system energy consumption. This assumption should eventually be relaxed and consider the different types of rewards available.

TRIPOD'S ARCHITECTURE

Tripod maximizes system-wide energy savings by nudging control point travelers through a personalized mobility menu. The menu is presented to the traveler via Tripod's smartphone User Interface (UI) offering *tokens* for certain alternatives and providing predictive information on its attributes. System-wide maximization of energy savings is a challenging problem. It needs to take into account system-wide supply and demand interactions as well as individual specific preferences towards different alternatives and *token* awarding. Since we aim to have a real-time system, the complexity is critical. Therefore the problem is decomposed into two tractable, loosely coupled problems: the System Optimization (SO) and the User Experience (UE). Figure 1 below gives an overview of Tripod's higher-level framework.

SO 1) estimates the current state of the transportation network; 2) predicts the state of the network given different *token* awarding strategies; 3) estimates the energy savings based on predicted conditions for different *token* awarding strategies; 4) optimizes the *token* awarding strategy; 5) provides system-wide *token* energy efficiency value in terms of energy savings per *token* to the UE.

The second component, UE, includes three modules: User Optimization (UO), User Interface (UI), and a preference updater. The first is responsible for generating a personalized menu of travel options to Tripod users upon request, with updated information and incentives based on the system-wide *token* energy efficiency, the transportation performance predictions and the energy impacts generated by SO. The menu includes alternatives that are attractive to the traveler based on a utility function, where coefficients for explanatory variables that represent personal tastes are

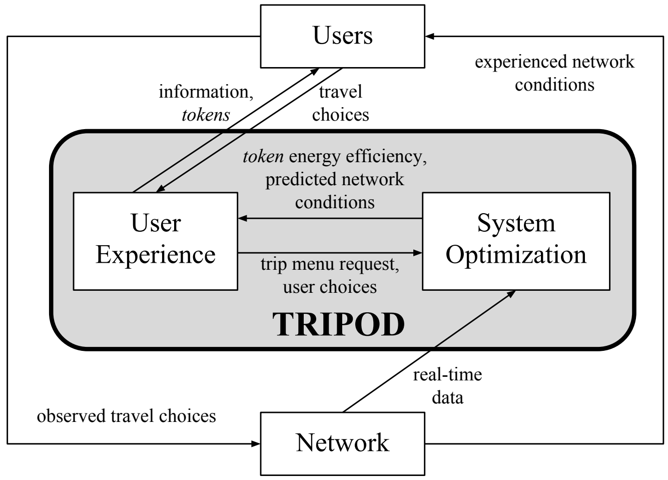

FIGURE 1: Tripod Framework

estimated from historical data, and values of alternative attributes such as travel time and energy cost are calculated based on the consistent, anticipatory information from Tripod's SO. Such a personalized menu aligns with the traveler's interest, and makes the system's architecture sustainable. It encourages energy efficient choices for the traveler by presenting explicit, accurate energy cost information, and notifying the traveler of incidents and providing alternatives. The UO formulation is described in more detail in a dedicated section below.

The second module builds upon the Future Mobility Sensing app (*34*), and extends it with a dedicated set of Tripod UIs (Figure 2) as well as sensing and tracking features. The UI development description is out of the scope of this paper and details on its initial implementation can be found in (*35*).

The preference updater continuously revises user preferences based on her/his observed choices together with observed choices from users of the similar group. These updates are important for the success of Tripod's personalization and are briefly presented in the UO section (a more detailed description of the preference update method formulation can also be found in (*36*)).

System Optimization: Framework

SO builds on a state-of-the-art dynamic traffic prediction model, DynaMIT (*37, 38*), which provides real-time predictions on how users respond to provided information and incentives, and on TripEnergy (*39*), a model that estimates the energy and emission impacts of travel for a variety of vehicle types, driving behaviors an environmental conditions.

DynaMIT is composed of two core modules, *state estimation* (SE) and *state prediction* (SP). The overall framework is depicted in Figure 5 and its modules are summarized below.

Both DynaMIT's SE and SP rely on an integrated demand and supply simulation. The demand simulator consists of choice models to capture habitual route and en-route choice behavior as well as response to prescriptive and descriptive real-time information. The supply simulator combines macroscopic speed-density relationships and deterministic queuing models. For more details, the reader is referred to (*37*).

During each DynaMIT execution cycle, the SE module uses a combination of historical information and real-time data from various sources (surveillance sensors, traffic information feeds,

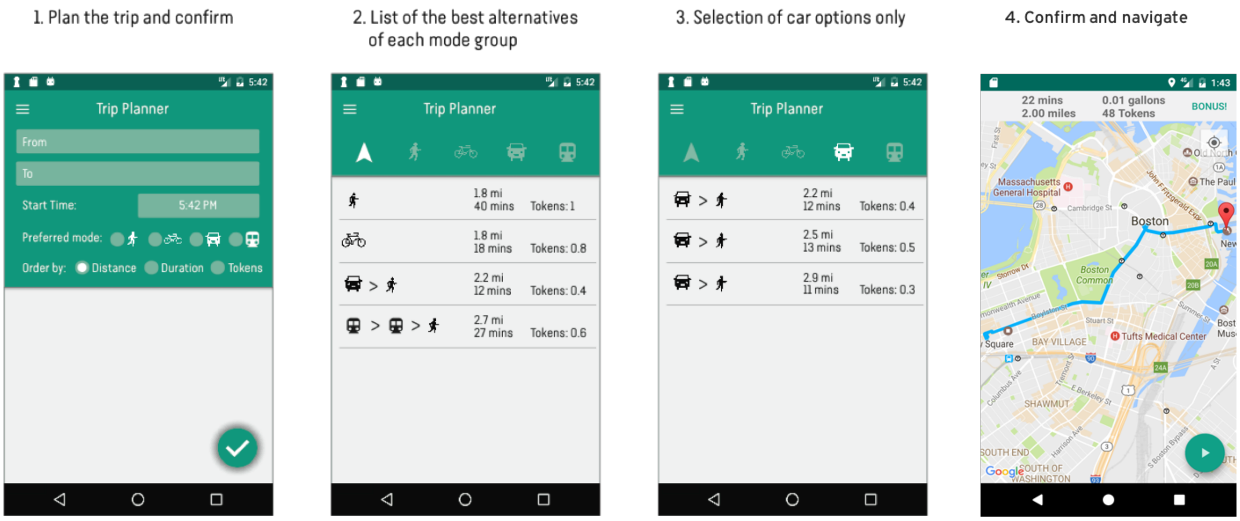

FIGURE 2: Tripod menu UI

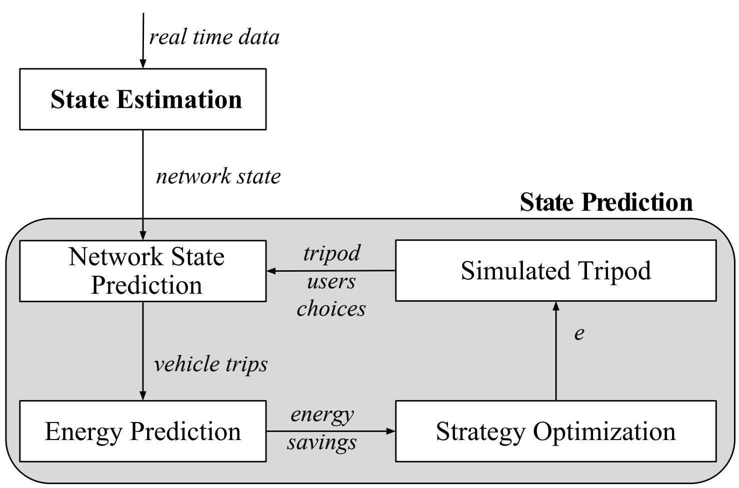

FIGURE 3: SO framework

weather forecasts) to first calibrate the demand (origin-destination matrices, behavioral parameters) and supply parameters (segment capacities, speed-density function parameters) of the simulator so as to replicate prevailing multi-modal network conditions, i.e., estimate the state of the entire network for the current *estimation interval* (e.g.: 5 minutes). Based on this estimate of the current network state, the SP module predicts future detailed traffic conditions on the multi-modal network for a pre-defined *prediction horizon* (e.g.: 30 minutes), consistent with a selected control/incentive strategy (please refer to (*38*) for more details on DynaMIT). Note that the route, departure time and mode choices of a simulated traveler change in response to the incentives and the predicted trip attributes (e.g.: travel time) presented through a simulated Tripod menu, which is personalized given *token* a specific energy efficiency strategy. This yields predicted network travel times which are then combined with the original incentives and trip attributes (using the method of successive averages or MSA) to obtain a revised trip attributes information. This procedure is iteratively performed until consistency is achieved, i.e., the provided travel time and predicted network travel times are within a pre-specified tolerance limit.

Together with the network state prediction, trip and system-wide energy consumptions are computed during the state prediction cycle using TripEnergy. This model accurately reconstructs possible driving behavior (speed and acceleration profiles) and accounts for different vehicle and drive-train characteristics ((*39*)). The energy savings are then computed by subtracting the energy consumption for a candidate incentive strategy and the prediction without incentives.

Within SP, the incentive strategy optimization process is invoked. Within this process, a set of candidate *token* energy efficiency strategies, *e*, is generated and used in the prediction horizons to be evaluated within each prediction in terms of system-wide energy savings. Again, in Tripod's formulation, an incentive strategy is defined by a set of energy efficiency values in terms of energy savings per *token* for the prediction horizon. This strategy definition is further discussed in the formulation section below.

After comparing the set of candidate *token* energy efficiency strategies, SO returns to UE an optimal *token* energy efficiency that is used to award incentives for the next roll period.

In summary, the SO framework extends the DynaMIT system in the following aspects: (1) simulating how users will respond to Tripod information and incentives, so that prediction outcomes are consistent with what users expect if they respond to Tripod; (2) incorporating an energy estimation model, TripEnergy (*39*), within the supply simulator that takes into account the characteristics of the vehicle and trip's low-resolution speed and acceleration profiles; and (3) integrating the optimization of *token* energy efficiency with the state prediction and trip information provision.

System Optimization: Formulation

The transportation network is represented as a directed graph $G(N,A)$ where N represents the set of network nodes ($|N| = n$) and A represents the set of links ($|A| = m$). Consider an arbitrary time interval $[t_0 - \Delta, t_0]$ where Δ is the size of the state estimation interval (e.g.: 5 minutes), also termed as the *roll period*. Assume that the length of the current state prediction horizon is equal to $H\Delta$ and extends from $[t_0, t_0 + H\Delta]$(each Δ interval within the prediction horizon is termed a prediction sub-interval). Further, assume that the *token* energy efficiency is fixed for a time interval equal to the roll period Δ and is aligned with the state estimation intervals of DynaMIT. As noted previously, the decision variables in the optimization problem are the vector of *token* energy efficiencies for the prediction horizon, denoted by $\mathbf{e} = (e_1, e_2, \ldots, e_H)$. Note that the *token* efficiency for a given prediction sub-interval or single roll period h can be formulated as a vector ($\mathbf{e}_h$), with an array of values by origin/destination, driving style, mode or ultimately by individual eventually potentially improving the efficiency of the control system.

The rolling horizon approach is illustrated in Figure 4 for a case where $H = 3$. In the executing cycle C1, the decision variables for the system optimization are the *token* energy efficiencies for the prediction sub-intervals P1–P3. If the optimum solution is denoted by $\mathbf{e}^* = (e_1^{C1}, e_2^{C1}, e_3^{C1})$, then the *token* energy efficiency for the next state estimation interval (execution cycle C2) is given by $e_{E2} = e_1^{C1}$. Similarly, the decision variables in the optimization for execution cycle C2 is the vector of *token* energy efficiencies for the next set of prediction sub-intervals.

Next, define the collection of vehicles $\nu = 1, \ldots V$ on the network during the prediction horizon $[t_0, t_0 + H\Delta]$ and let the predictive travel time guidance be denoted by $\mathbf{tt^g} = (\mathbf{tt_i^g}; \forall i \in A)$,

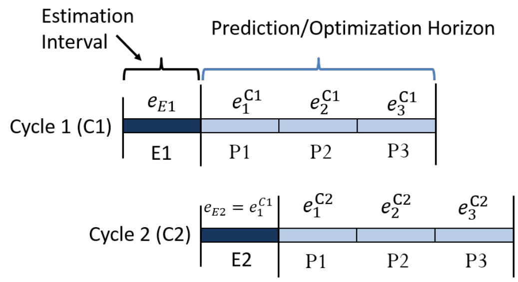

FIGURE 4: Illustration of the rolling horizon approach

where $\mathbf{tt_i^g}$ represents a vector of the time dependent link travel times (guidance) for link i. Further, let the simulated trajectory of vehicle v be denoted by $\boldsymbol{\tau}^v$ and the collection of vehicle trajectories by $\boldsymbol{\tau} = (\boldsymbol{\tau}^1, \dots, \boldsymbol{\tau}^V)$. Note that the vehicle trajectories are a result of the state prediction simulation module of DynaMIT and cannot be written as an explicit function of the *token* energy efficiency and predictive guidance. We characterize the complex relationship through a function $S(.)$ that represents the coupled demand and supply simulators as,

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

where $\mathbf{x}^p, \boldsymbol{\gamma}^p$ represents the forecasted demand and supply parameters for the prediction horizon. Also note that the iterative procedure described in the previous section ensures that the aggregate travel times that result from the vehicle trajectories $\boldsymbol{\tau}$ are consistent with the guidance $\mathbf{tt^g}$.

For a given *token* energy efficiency vector $\mathbf{e}$, the system–wide energy savings across the prediction horizon is defined as the difference between energy consumption under the incentive awarding strategy determined by $\mathbf{e}$ and the energy consumption of a baseline state prediction with no incentive awarding as,

$$E(\boldsymbol{e}, \boldsymbol{\tau}, \bar{\tau}) = \sum_{v=1...V} f(\boldsymbol{\tau}^v(\mathbf{e}), \boldsymbol{\theta}^v) - \sum_{v=1...V} f(\bar{\tau}^v, \boldsymbol{\theta}^v) \tag{2}$$

where $f(.)$ represents the TripEnergy module, $\bar{\tau}^v$ represents the simulated trajectory of vehicle v in the baseline state prediction (without incentives) and $\boldsymbol{\theta}^v$ represent a vector of parameters that characterize vehicle v such as driving style, vehicle type, etc. More details on the TripEnergy module may be found in (*39*).

Let $T(t_0, \mathbf{e})$ denote the number of *tokens* consumed in the current state prediction horizon $[t_0, t_0 + H\Delta]$. Note that this is a function of the *tokens* energy efficiency vector $\mathbf{e}$ and is computed by the state prediction module based on the simulated choices of Tripod users in DynaMIT. Further, let $B(t_0)$ denote the number of *tokens* available at time t_0 and $W(t_0, \mathbf{e}) = B(t_0) - T(t_0, \mathbf{e})$, the available *token* budget for the rest of the day.

The Tripod system assumes a daily budget of incentives that can be allocated. Consequently, the optimization formulation must, either explicitly or implicitly, account for this budget so as to prevent from myopically allocating too many *tokens* within any given roll period. This is done using a *beyond horizon energy estimator* $g(.)$ which provides an estimate of energy savings from the end of the current prediction horizon ($t_0 + H\Delta$) until the end of the day based on the available budget $W(t_0, \mathbf{e})$, a feature vector $\boldsymbol{\zeta}(t_0)$ that describes characteristics of the current day (e.g.:

the network state in the form of time–dependent link flows, speeds and densities over the period of the day thus far, the network state over the prediction horizon, weather, incidents, special events, etc) and a historical database $\mathbf{Y}$ (containing historical observed and simulated traffic conditions, *token* consumption, energy savings, weather, special–event and incident information). Thus, the beyond–horizon energy savings is given by,

$$\bar{E}(\boldsymbol{e}, t_0) = g(W(t_0, \mathbf{e}), \boldsymbol{\zeta}(t_0), \mathbf{Y}) \tag{3}$$

The beyond horizon energy estimator described above is still under development and as a workaround, for the current implementation, a fixed *token* budget is imposed for each roll period. The collection of simulated vehicle trajectories by $\boldsymbol{\tau}$ results from the reaction of Tripod users to a given *token* energy efficiency strategy. Therefore in DynaMIT, the menu generation and its personalization process needs to be simulated. For that, we have developed the Simulated User Optimization (SUO) module that interacts directly with the demand module of DyaMIT. First, the number of *tokens* awarded to individual n for alternative p is defined as,

$$TK_{np} = max\left(0, \left[\frac{E_{n0} - E_{np}}{e}\right]\right) p \in P_n \tag{4}$$

where E_{np} is the energy consumption of alternative p for individual n (computed based on the predicted network state), E_{n0} is the predicted energy consumption of individual n without *tokens*, and e is the *token* energy efficiency for the current roll period or prediction sub-interval as the case may be. Then the (simulated) individual n -specific UO process described in the next section is ran.

Finally, the response to the simulated Tripod menu is simulated within the SP. For the numerical experiments presented in this paper, it is assumed that the total network demand is fixed (inelastic) and the behavioral response of users to the incentives and predictive travel time is solely through route choice. This is modeled within the demand simulator of DynaMIT using a multinomial logit model. The utility perceived by individual n for path $p \in P_n$ (where P_n is the choice set) is given by,

$$\begin{aligned} U_{np} &= V_{np} + \varepsilon_{np} \\ &= \beta_{TT} TT_{np} + \beta_C (C_{np} - \alpha_{np} \gamma TK_{np}) + \varepsilon_{np} \end{aligned} \tag{5}$$

where V_{np} is the systematic utility, β_{TT} and β_C are model parameters, TT_{np} is the predicted (or historical depending on whether the vehicle has information access) travel time on path p, C_{np} is the monetary cost, γ is the market value of the *token*, α_{np} is a unit-free *token* value inflation/deflation factor, TK_{np} is the number of *tokens* allocated to individual n for using path p (discussed in more detail subsequently), and ε_{np} is a random error component that is *i.i.d.* Gumbel distributed. For the ease of notation, we did not talk about the specification of parameters. Nevertheless, β_{TT} and β_C are defined such that they are introduced with a distribution, i.e., heterogeneity in the population is considered.

It should be noted that a more sophisticated implementation of the Tripod is ongoing which includes a nested logit model that captures travelers'responses to the *tokens* along with the choice dimensions of mode, route, departure time, trip cancellation and driving style. Further, since the current implementation is cars-only we assume a one-to-one correspondence between individuals and vehicles.

With these definitions, the optimization problem to solve for the current prediction horizon is given by,

$$\begin{aligned} \underset{\mathbf{e}}{\text{maximize}} \quad & E(\boldsymbol{e},\boldsymbol{\tau},\bar{\tau}) + \bar{E}(\boldsymbol{e},t_0) \\ \text{subject to} \quad & \boldsymbol{\tau} = S(\mathbf{x}^p,\boldsymbol{\gamma}^p,\mathbf{tt}^g,\boldsymbol{e}) \\ & \mathbf{e} \geq 0 \\ & W(t_0,\mathbf{e}) \geq 0 \end{aligned} \tag{6}$$

User Optimization: Framework

UO generates a personalized menu of options in real-time that maximizes a user-specific objective function (e.g.: consumer surplus) based on the guidance from SO regarding predicted traffic conditions as well as *token* energy efficiencies. The *tokens* associated with each option are determined before running UO based on the formula provided in (eq. 4) with the optimal *token* energy efficiencies, e, given by SO every 5 min. The reference energy value, E_{n0}, is the expected energy consumption without *tokens*, i.e., the preferences are used to compute the choice probabilities of each option and the expected value is computed accordingly when the *tokens* are not present in the utility function (eq. 5). The user specific (menu) objective function - i.e.: selection of the alternatives to present in the menu - is the consumer surplus represented by the log-sum from the choice model.

User Optimization: Formulation

Assume that we are selecting M_n options to be presented on the menu among the set of alternatives $p \in P_n$ for user n. The binary decision variable x_{np} represents if option p is selected to be on the menu or not. We have a set of preference parameters, $i \in I_n$, for each individual n, i.e., we have I_n many draws from the posterior distribution of individual n's preference parameters. This set of parameters represent the distribution of a specific user's preferences which is referred as *intra-consumer heterogeneity*. Note that the preference parameters are intended to be updated as users make new trips, as well as a periodic (e.g., weekly) entire-traveler population-parameter update with information from other travelers *off-line update*. The individual level parameters are updated after every choice that is observed with an efficient Bayesian *on-line update* procedure (see (*36*) for detailed formulations). These two update procedures enable to make use of most up-to-date choice data in an efficient way. The UO model is based on these continuously updated preference parameters and given as follows for a user n:

$$\begin{aligned} \underset{\mathbf{x}}{\text{maximize}} \quad & \sum_{i \in I_n} \frac{1}{\mu} \left[\ln \left(\sum_{p \in P_n} x_{np} exp(\mu V_{npi}) \right) \right] \\ \text{subject to} \quad & \sum_{p \in P_n} x_{np} \leq M_n \\ & x_{np} \in \{0,1\} \quad \forall p \in P_n \end{aligned} \tag{7}$$

where systematic utility, V, is given in (eq. 5). Note that the parameters, β_{TT} and β_C, are represented by I_n draws from the posterior distribution and thus the index i in the utility. The objective function represents the consumer surplus under a set of preference parameters. The constraint maintains that the number of options to be presented on the menu is not greater than M_n. This is included to ensure that the menu is not very large but can be customized for each user. The added value of personalization will be greater when the menu size is limited.

The UO problem (eq. 7) is a nonlinear integer programming problem. In general such problems are difficult to solve. Currently an approximation of this problem is utilized in the Tripod system such that mean of the posterior distribution for preference parameters is used which removes the summation over the set I_n from the objective function. The solution simply turns to be a sorting algorithm and the menu to be presented will consist of alternatives that provide the top M_n utility values among all alternatives. Furthermore, the approximation of the UO is integrated in DynaMIT (through the above mentioned SUO module) in order to represent UO in SO for evaluating different strategies. As we are dealing with simulated agents in SO, an efficient computation is needed and this approximation serves as an appropriate alternative.

In (*40*) the proposed UO model performance is investigated: the added value of personalized menu optimization compared to non-personalized menu optimization and the performance of the menu optimization based on updated preference parameters. It is observed that in all the cases personalized menu optimization provides better menus, namely the probability that the user will choose an alternative on the menu is consistently higher. Furthermore, the on-line update procedure is shown to provide preference parameters that are close to the estimated parameters with the full set of choices.

NUMERICAL EXPERIMENTS

This section reports results from numerical experiments conducted to evaluate the performance of a preliminary version of Tripod that incentivizes only route choice for private vehicles, and replaces the beyond-horizon energy saving estimation (where *token* budget allocation for the remainder of the day is endogenous) with a roll-period based *token* budget (where *token* budget allocation for the remainder of the day is exogenous). Although this is a simplification that can yield sub-optimal results, it suffices for the purposes of the experiments here, which are intended as a proof-of-concept and to gain some preliminary insights into the performance of Tripod.

Assumptions

For the purposes of the simple experiment, the parameters of the behavioral model (eq. 5) were assumed to be uniformly distributed: $\beta_{TT} = -0.01$ and $\beta_C = -2$ which yields a value of time of 18$ per hour that is consistent with empirical studies. The market value of the *token* γ is assumed to be 0.5$ and $\alpha_{np} = 1 \forall n, p$. Ideally, a stated-preference survey would help in estimating initial model parameters based on Boston commuters' response to hypothetical travel choice scenarios involving *tokens*. In the absence of such information the above values were used in the current experiments. In addition, the *token* budget per five minute interval is assumed to be 1000 and the penetration rate of 50% Tripod op-in travelers.

Setup

The experiments are conducted on the Boston Central Business District (CBD) network which consists of 846 nodes, 1746 links, 3085 segments and 5057 lanes including both highways and arterials. The simulation period was from 6:30 am to 9:00 am with a roll period of 5 minutes

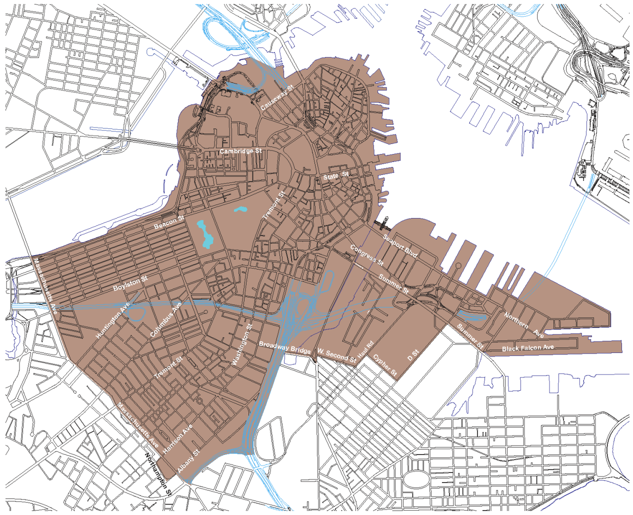

FIGURE 5: Boston CBD Network

and a rolling horizon of 15 minutes. The time-dependent origin-destination demands and supply parameters (segment capacities, free-flow speeds, and traffic dynamics parameters) were obtained from the Boston Region Metropolitan Planning Organization planing model. The effectiveness of Tripod is evaluated by performing two simulations: one in which incentives are awarded based on optimal estimates of the *token* energy efficiency computed by the SO module and a base simulation, involving DynaMIT with no incentives. The performance measures are the average travel time and energy consumption per vehicle and the distribution of travel time and energy consumption.

A sensitivity analysis is also performed by systematically varying the base demand level and the *token* inflation factor α_{np} (which is assumed to be fixed across individuals and alternatives).

Results

The preliminary results show promising potential of tokens in nudging users to more energy efficient choices. The percentage reduction in average energy consumption (across a total of 96000 simulated trips) is 5.94% with a corresponding reduction in average travel times of 4.8%. The reduction in energy consumption and travel times is also evident from the cumulative distribution plots (CDF) in Figure 6.

Effect of Demand

The sensitivity analysis with respect to demand (see Table 1) indicates that as the demand level increases, the percentage reduction in energy and travel time increases, reaching a maximum of 7.45% and 9.59% respectively, following which, the improvement decreases to 3.73% and 0.24% respectively. This may be explained by the fact that at very high levels of demand, the energy savings to be gained through re-routing are minimal given the already congested network state.

Effects of preferences towards incentives

The sensitivity analysis with respect to the token inflation factor is summarized in Table 1 and indicates, as expected, that as the token inflation factor increases (or the sensitivity to tokens increases) the energy and travel time savings increase. The percentage reduction in average energy consumption increases from 3.24% to 6.45% as the token inflation factor increases from 0.5 to 1.2.

FIGURE 6: Travel Time and Energy CDFs

(a) Travel Time CDF

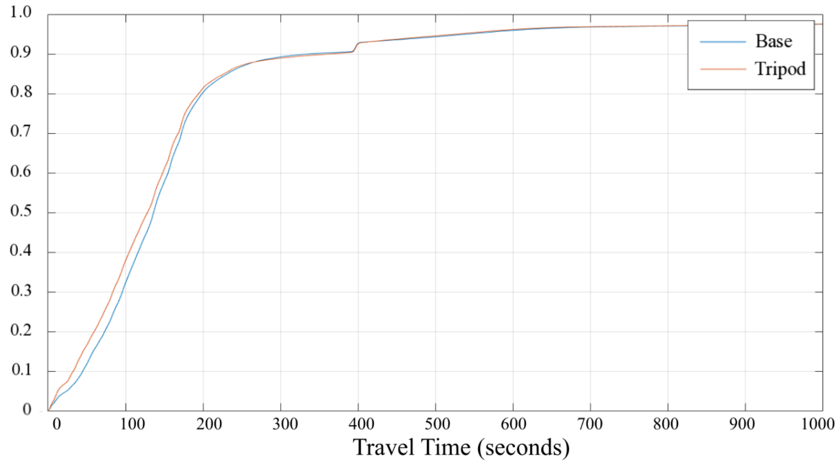

(b) Energy CDF

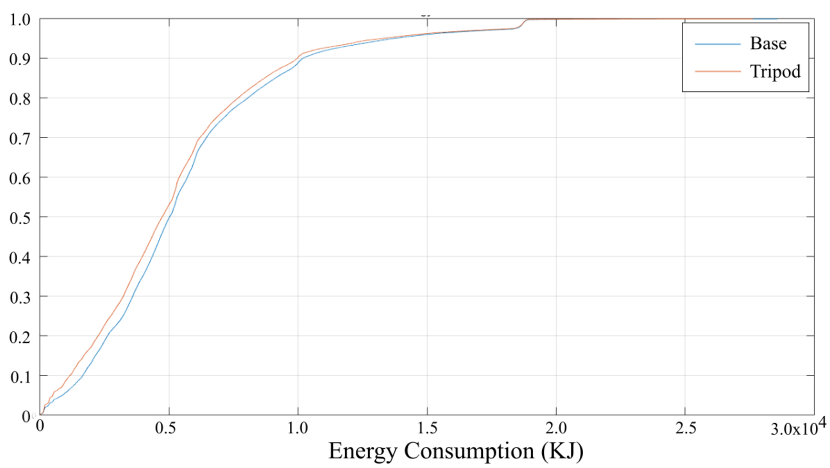

The corresponding improvement in average travel time increases from 0.5% to 5.85%.

CONCLUSION

With the very simple numerical example we were able to illustrate the potential of Tripod, opening a new research stream on real-time optimized incentives. A new solution to optimize incentives in real-time under a fix budget constraint was formulated, implemented and tested, using simulation-based predictions and behavioral models to cope with both reaction to incentives and the personalization features of Tripod. To allow for real-time performance, we decoupled the optimization into two components, which have different time resolutions and objective functions. The System Optimization includes two state-of-the-art models, DynaMIT and TripEnergy, allowing for accurate and flexible estimates and predictions of different network traffic conditions and energy consumption. The User Optimization brings the benefit of personalization and real-time behavioral preference updating to the incentive system. Despite the computational and tractability advantages

TABLE 1: Effect of Demand and Token Inflation Factor Level

Demand (% of base)	Average Energy (KJ)			Average Travel Time (seconds)		
	Tripod	Base	% Diff	Tripod	Base	% Diff
50	4858.72	5235.11	7.19	118.72	129.33	8.21
60	4882.13	5268.14	7.33	119.39	130.19	8.29
65	4956.24	5342.08	7.22	121.91	133.21	8.49
75	4947.68	5329.28	7.16	122.24	133.32	8.31
80	4990.91	5392.62	7.45	127.71	141.26	9.59
90	5054.42	5436.71	7.03	139.15	150.99	7.84
100	5330.04	5666.58	5.94	177.27	186.21	4.80
110	5589.16	5805.96	3.73	193.73	194.19	0.24
Token Inflation factor	Average Energy (KJ)			Average Travel Time (seconds)		
	Tripod	Base	% Diff	Tripod	Base	% Diff
0.5	5483.26	5666.58	3.24	185.29	186.21	0.50
0.6	5435.55	5666.58	4.08	183.87	186.21	1.26
0.7	5409.07	5666.58	4.54	182.40	186.21	2.05
0.8	5370.36	5666.58	5.23	179.46	186.21	3.63
0.9	5343.10	5666.58	5.71	177.73	186.21	4.56
1	5330.04	5666.58	5.94	177.27	186.21	4.80
1.1	5314.15	5666.58	6.22	176.06	186.21	5.46
1.2	5301.32	5666.58	6.45	175.32	186.21	5.85

of the proposed decoupling of the optimization function, it precludes with the optimality condition. Further simulation should be carried out to test the performance of the proposed formulation against others (e.g.: including more decisions variables such as multiple token efficiency values, incorporate equity, mobility or energy constraints in the objective function, etc). Also, the simplifications assumed in the current numerical experiments are currently being relaxed. Finally, the Tripod team is working on the evaluation of the proposed system in a detailed agent- and activity-based simulation environment for the Greater Boston Area to be able to accurately generate a wide range of performance evaluations.

ACKNOWLEDGEMENTS

This research was carried out under the U.S. DoE's Advanced Research Projects Agency-Energy (ARPA-E) TRANSNET Program. We are thankful to the Central Transportation Planning Staff, Boston Region MPO, for providing part of the data used in this study. Finally, we would like to thank Hyoshin Park, Xiang Song, Samarth Gupta, Yihang Sui, David Sukhin, Zachary Needell and Sayeeda Ayaz, Balakumar Marimuthu and Anthony Battista for their important support.

REFERENCES

[1] Winston, C., On the performance of the US transportation system: Caution ahead. *Journal of Economic Literature*, Vol. 51, No. 3, 2013, pp. 773–824.

[2] Balakrishna, R., M. Ben-Akiva, J. Bottom, and S. Gao, *Information Impacts on Traveler Behavior and Network Performance: State of Knowledge and Future Directions*, Springer New York, New York, NY, pp. 193–224, 2013.

[3] Mori, U., A. Mendiburu, M. Álvarez, and J. A. Lozano, A review of travel time estimation and forecasting for Advanced Traveller Information Systems. *Transportmetrica A: Transport Science*, Vol. 11, No. 2, 2015, pp. 119–157.

[4] Pigou, A. C., *The Economics of Welfare*. Macmillan, London, 1920.

[5] Lindsey, R., Reforming road user charges: a research challenge for regional science. *Journal of Regional Science*, Vol. 50, No. 1, 2010, pp. 471–492.

[6] 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.

[7] Gallego, G. and G. van Ryzin, Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons. *Management Science*, Vol. 40, No. 8, 1994, pp. 999–1020.

[8] Zockaie, A., M. Saberi, H. S. Mahmassani, L. Jiang, A. Frei, and T. Hou, Activity-Based Model with Dynamic Traffic Assignment and Consideration of Heterogeneous User Preferences and Reliability Valuation. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2493, 2015, pp. 78–87.

[9] Janson, M. and D. Levinson, HOT or not: Driver elasticity to price on the MnPASS HOT lanes. *Research in Transportation Economics*, Vol. 44, No. Supplement C, 2014, pp. 21 – 32, road Pricing in the United States.

[10] Yang, H. and X. Wang, Managing network mobility with tradable credits. *Transportation Research Part B: Methodological*, Vol. 45, No. 3, 2011, pp. 580 – 594.

[11] Lahlou, S. and L. Wynter, A Nash equilibrium formulation of a tradable credits scheme for incentivizing transport choices: From next-generation public transport mode choice to HOT lanes. *Transportation Research Part B: Methodological*, Vol. 101, 2017, pp. 185 – 212.

[12] Nie, Y. M., Transaction costs and tradable mobility credits. *Transportation Research Part B: Methodological*, Vol. 46, No. 1, 2012, pp. 189 – 203.

[13] De Palma, A., S. Proost, R. Seshadri, and M. Ben-Akiva, Tolls Versus Mobility Permits: A Comparative Analysis. *TRB 96th Annual Meeting, January, 2017*, 2017.

[14] Akamatsu, T. and K. Wada, Tradable network permits: A new scheme for the most efficient use of network capacity. *Transportation Research Part C: Emerging Technologies*, Vol. 79, 2017, pp. 178 – 195.

[15] Dogterom, N., D. Ettema, and M. Dijst, Tradable credits for managing car travel: a review of empirical research and relevant behavioural approaches. *Transport Reviews*, Vol. 37, No. 3, 2017, pp. 322–343.

[16] Nie, Y. M., A New Tradable Credit Scheme for the Morning Commute Problem. *Networks and Spatial Economics*, Vol. 15, No. 3, 2015, pp. 719–741.

[17] Wu, D., Y. Yin, S. Lawphongpanich, and H. Yang, Design of more equitable congestion pricing and tradable credit schemes for multimodal transportation networks. *Transportation Research Part B: Methodological*, Vol. 46, No. 9, 2012, pp. 1273 – 1287.

[18] Grant-Muller, S. and M. Xu, The Role of Tradable Credit Schemes in Road Traffic Congestion Management. *Transport Reviews*, Vol. 34, No. 2, 2014, pp. 128–149.
[19] Leblanc, R. and J. L. Walker, Which Is the Biggest Carrot? Comparing Nontraditional Incentives for Demand Management. In *Transportation Research Board 92nd Annual Meeting*, 2013, 13-5039.
[20] Kumar, V., C. R. Bhat, R. M. Pendyala, D. You, E. Ben-Elia, and D. Ettema, Impacts of Incentive-Based Intervention on Peak Period Traffic. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2543, 2016, pp. 166–175.
[21] Ettema, D., J. Knockaert, and E. Verhoef, Using incentives as traffic management tool: empirical results of the "peak avoidance" experiment. *Transportation Letters*, Vol. 2, No. 1, 2010, pp. 39–51.
[22] Castellanos, S., Delivering modal-shift incentives by using gamification and smartphones: A field study example in Bogota, Colombia. *Case Studies on Transport Policy*, Vol. 4, No. 4, 2016, pp. 269 – 278.
[23] Merugu, D., B. S. Prabhakar, and N. Rama, An incentive mechanism for decongesting the roads: A pilot program in bangalore. In *Proc. of ACM NetEcon Workshop*, 2009.
[24] Zhu, C., J. S. Yue, C. V. Mandayam, D. Merugu, H. K. Abadi, and B. Prabhakar, Reducing Road Congestion Through Incentives: A Case Study. In *Transportation Research Board 94th Annual Meeting*, 2015.
[25] Lu, F., *Framework for a lottery-based incentive scheme and its influence on commuting behaviors: an MIT case study*. Ph.D. thesis, Massachusetts Institute of Technology, 2015.
[26] Rey, D., V. V. Dixit, J.-L. Ygnace, and S. T. Waller, An endogenous lottery-based incentive mechanism to promote off-peak usage in congested transit systems. *Transport Policy*, Vol. 46, 2016, pp. 46 – 55.
[27] Hu, X., Y.-C. Chiu, and L. Zhu, Behavior Insights for an Incentive-Based Active Demand Management Platform. *International Journal of Transportation Science and Technology*, Vol. 4, No. 2, 2015, pp. 119 – 133.
[28] Gao, J. H. and L.-S. Peh, Roadrunner: Infrastructure-less Vehicular Congestion Control. In *The 21st Intelligent Transport Systems World Congress*, 2014.
[29] Shaheen, S., A. Cohen, I. Zohdy, and B. Kock, *Smartphone applications to influence travel choices: practices and policies*. USDOT, Federal Highway Administration, 2016.
[30] Meyer, E., J. Chernick, and B. Dalton, *Matching system for ride reservation platforms*, 2012, uS Patent 8,285,570.
[31] Poslad, S., A. Ma, Z. Wang, and H. Mei, Using a smart city IoT to incentivise and target shifts in mobility behaviourâĂŤIs it a piece of pie? *Sensors*, Vol. 15, No. 6, 2015, pp. 13069–13096.
[32] *Empower EU project*. `http://mobility-apps.eu/`, ????, accessed: 2017-11-14.
[33] Hu, X., Y.-C. Chiu, and L. Zhu, Behavior Insights for an Incentive-Based Active Demand Management Platform. *International Journal of Transportation Science and Technology*, Vol. 4, No. 2, 2015, pp. 119 – 133.
[34] Cottrill, C., F. Pereira, F. Zhao, I. Dias, H. Lim, M. Ben-Akiva, and P. Zegras, Future mobility survey: Experience in developing a smartphone-based travel survey in singapore. *Transportation Research Record: Journal of the Transportation Research Board*, , No. 2354, 2013, pp. 59–67.
[35] Srinivasan, A., *Leveraging smartphones to incentivize city-Wide, energy efficient transportation*. Massachusetts Institute of Technology, 2017.

[36] Danaf, M., F. Becker, X. Song, B. Atasoy, and M. Ben-Akiva, Personalized Recommendations using Discrete Choice Models with Inter-and Intra-Consumer Heterogeneity. In *International Choice Modeling Conference*, 2017.
[37] Ben-Akiva, M., H. Koutsopoulos, C. Antoniou, and R. Balakrishnan. In *In Fundamentals of Traffic Simulation* (J. Barceló, ed.), Springer, New York, NY, USA, 2010, chap. Traffic Simulation with DynaMIT, pp. 363–398.
[38] Lu, Y., 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 *Proceedings of the IEEE 18th International Conference on ITS*, 2015, pp. 2250–2255.
[39] McNerney, J., Z. Needell, M. T. Chang, M. Miotti, and J. Trancik, Estimating Personal Vehicle Energy Consumption Given Limited Travel Survey Data. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2628, 2017.
[40] Song, X., M. Danaf, B. Atasoy, and M. Ben-Akiva, Personalized Menu Optimization with Preference Updater: A Boston Case Study, 2017, working paper, MIT.