



HAL
open science

System-level optimization of multi-modal transportation networks for energy efficiency using personalized incentives: formulation, implementation, and performance

Andrea Araldo, Song Gao, Ravi Seshadri, Carlos Lima Azevedo, Hossein Ghafourian, Yihang Sui, Sayeeda Ayaz, David Sukhin, Moshe Ben-Akiva

► To cite this version:

Andrea Araldo, Song Gao, Ravi Seshadri, Carlos Lima Azevedo, Hossein Ghafourian, et al.. System-level optimization of multi-modal transportation networks for energy efficiency using personalized incentives: formulation, implementation, and performance. *Transportation Research Record*, 2019, 2673 (12), pp.425-438. 10.1177/0361198119864906 . hal-02190146

HAL Id: hal-02190146

<https://hal.science/hal-02190146>


Submitted on 28 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

System-Level Optimization of Multi-Modal Transportation Networks for Energy Efficiency using Personalized Incentives: Formulation, Implementation, and Performance

Andrea Araldo^{1,2}, Song Gao³, Ravi Seshadri⁴, Carlos Lima Azevedo⁵, Hossein Ghafourian³, Yihang Sui², Sayeeda Ayaz³, David Sukhin², and Moshe Ben-Akiva²

Transportation Research Record
1–14
© National Academy of Sciences:
Transportation Research Board 2019
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0361198119864906
journals.sagepub.com/home/trr


Abstract

The paper presents the system optimization (SO) framework of Tripod, an integrated bi-level transportation management system aimed at maximizing energy savings of the multi-modal transportation system. From the user's perspective, Tripod is a smartphone app, accessed before performing trips. The app proposes a series of alternatives, consisting of a combination of departure time, mode, and route. Each alternative is rewarded with an amount of tokens which the user can later redeem for goods or services. The role of SO is to compute the optimized set of tokens associated with the available alternatives to minimize the system-wide energy consumption under a limited token budget. To do so, the alternatives that guarantee the largest energy reduction must be rewarded with more tokens. SO is multi-modal, in that it considers private cars, public transit, walking, car pooling, and so forth. Moreover, it is dynamic, predictive, and personalized: the same alternative is rewarded differently, depending on the current and the predicted future condition of the network and on the individual profile. The paper presents a method to solve this complex optimization problem and describe the system architecture, the multi-modal simulation-based optimization model, and the heuristic method for the online computation of the optimized token allocation. Finally it showcases the framework with simulation results.

An efficient, reliable, and sustainable transportation system is vital to the prosperity of society and the well-being of people. Urban transportation networks worldwide, however, are beset by issues of excessive congestion and energy consumption, which are critical obstacles to achieving these goals. Given the limitations in adding capacity, travel demand management has received significant attention from researchers and practitioners as an effective means of achieving a more efficient utilization of existing infrastructure. From the real-time demand management perspective, externalities such as congestion and vehicular emissions have been historically addressed with information provision (1) or pricing strategies (2, 3). Indeed, one of the most widely discussed demand management strategies is congestion pricing (see Lindsey (4) or Tsekeris and Voß (5) for comprehensive reviews). Congestion pricing is based on the idea that transportation users should pay for the full cost of travel, which includes both their own costs and costs imposed on other

users because of congestion. It aims at curbing excessive demand and making efficient use of the existing transportation facilities. Singapore, London, and Stockholm are among the few major cities worldwide that have such a scheme area-wide. For example, in London, one is required to pay £11.50 to drive a personal vehicle within Central London between 7:00 a.m. and 6:00 p.m., with more polluting vehicles paying more. Congestion pricing is controversial for a range of reasons, however,

¹Télécom SudParis - Institut Polytechnique de Paris, UMR CNRS SAMOVAR, Evry, France

²Massachusetts Institute of Technology, Cambridge, MA

³University of Massachusetts, Amherst, MA

⁴Singapore-MIT Alliance for Research and Technology, Singapore

⁵Technical University of Denmark, Lyngby, Denmark

Corresponding Author:

Address correspondence to Andrea Araldo:
andrea.araldo@telecom-sudparis.eu

including the general aversion to charges, as well as equity concerns, in that it is seen as benefiting high income users at the expense of low income users.

Incentive policies are alternative demand management approaches that, instead of charging people for using a congestion-inducing or polluting travel option, reward them for using a less congestion-inducing or polluting travel option. There have been a number of pilot studies at various scales. One of the most notable is the “peak avoidance” (6) experiment conducted in the Netherlands in 2006 on a heavily congested highway. Volunteers participated in a scheme whereby they could receive daily rewards, either monetary or in the form of credits that could be exchanged for a smartphone. Participants could earn a reward by driving at off-peak times, switching to another mode of transportation, such as cycling, public transport or carpooling, or by working from home. Researchers found that between 30% and 40% of participants avoided peak hour driving. Other experiments have explored the behavioral reaction to points-based, lottery-based, personalized, or smartphone based static incentives (7–10).

Intuitively appealing and empirically verified with simple schemes (11–14), the design, implementation, and evaluation of a real-time, personalized incentive scheme that is also optimized at a multi-modal system level remains a challenging problem. The challenge is first of all methodological: which formulation should be used to compute, in real time, the amount of incentives to reward any traveler entering the transportation system? How can the incentive strategy be adapted to a network state that is continuously evolving? How can we consider the impact of the strategy on future time intervals?

This study proposes an ensemble of methods to address these questions and demonstrates their implementation in Tripod, a smartphone based system that provides, in real time, personalized incentives in the form of tokens, with the objective to nudge travelers toward more globally efficient choices of mode, departure time, and route. The primary contribution of this paper is the design and implementation of a framework and an algorithm to perform real-time system-level token optimization in a rolling horizon fashion, based on predictive multi-modal traffic simulation. A novel aspect of the approach is that it reduces this complex optimization problem to one with a single scalar decision variable, termed the token energy efficiency (TEE). Conceptually, the TEE corresponds to the amount of energy a traveler must save to earn one token (15). The framework adjusts this value in real time, adapting to the state of the network, to induce the largest overall energy reduction. Once the TEE for a certain time interval is decided, the tokens awarded to the traveler for choosing a specific alternative are proportional to the TEE. The

contribution of this paper is in the methodology to compute this TEE.

The remainder of the paper is organized as follows. First the overall Tripod architecture is presented, followed by the system optimization architecture within Tripod. Next, the paper presents the multi-modal demand and supply simulators, which are the basis of the simulation-based optimization. The system optimization problem is then formulated and the solution algorithm presented. Next, the paper presents and discusses the results of a simulation using the Boston Central Business District (CBD) network. Finally, there is a conclusion and discussion of directions for future research.

Overview of Tripod

As presented by Lima Azevedo et al. (15), Tripod maximizes in real time the system-wide energy efficiency of multi-modal transportation by offering personalized incentives to encourage travelers to select alternatives with smaller energy impact.

First, Tripod is reviewed from the user’s perspective, to quote Lima Azevedo et al. (15): “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 relation to 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 incentives are provided through a personalized mobility menu, presented to the traveler via Tripod’s smartphone user interface (UI) (see Figure 1).

To achieve system-wide energy efficiency, it is necessary to optimize in real time the incentives offered in the menu, taking into account that the incentive budget is limited. System-wide maximization of energy savings is a challenging problem. It needs to consider system-wide supply and demand interactions as well as individual specific preferences toward different alternatives and token awarding. To tackle this complexity, Lima Azevedo et al. (15) decomposed the energy efficiency maximization into a bi-level structure with two loosely coupled problems: system optimization (SO) and user experience (UE) (see Figure 2). SO is the top level, defining the overall policy optimization, while UE is the lower layer, taking care of individual specific optimization, and thus

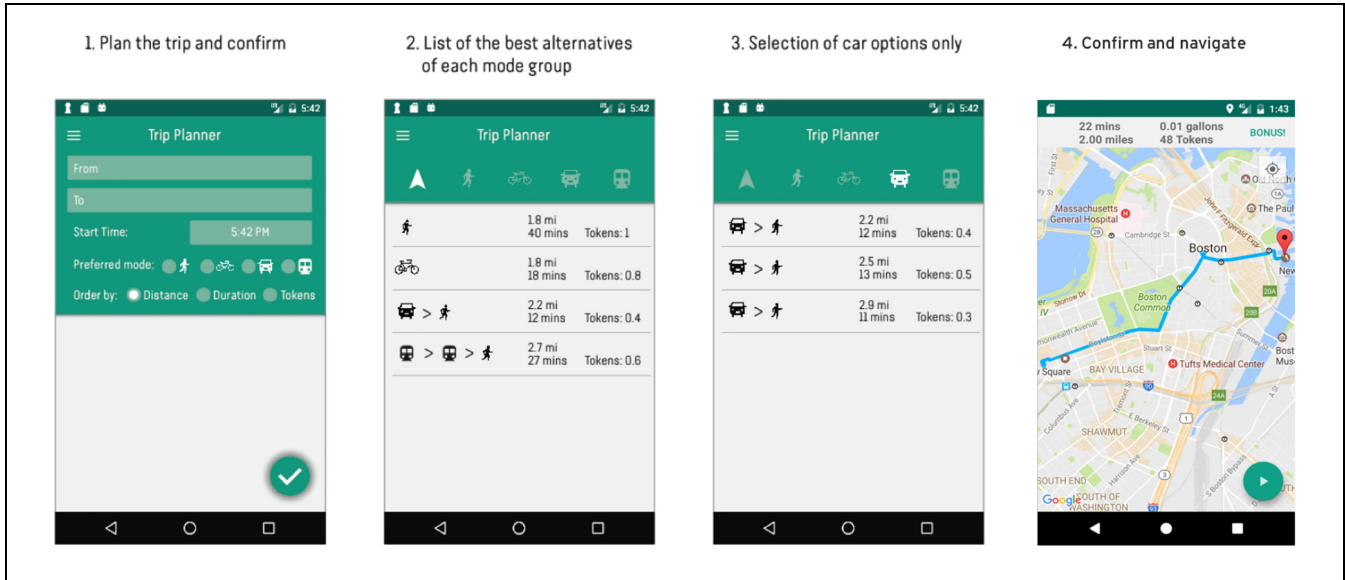


Figure 1. Tripod menu UI (15).

personalization. The link between these two loosely coupled problems consists in the computation in real time of the current TEE, defined as the amount of energy a traveler must save to be rewarded with one token. The TEE is the key decision variable of SO and is used in every menu personalization triggered by each trip request issued by Tripod users on the app (see Figure 2). Along with the TEE, the SO also provides the UE with the full choice set of alternatives (and its policy-consistent predicted attributes) to be considered in the menu personalization. This paper focuses precisely on the SO, detailing in the remaining sections its formulation, implementation, and performance.

The second component, UE, includes three modules: user optimization (UO), UI, and a preference updater. The first is responsible for generating a personalized menu of travel options to the Tripod user upon request, with updated information and incentives based on the system-wide TEE, the transportation performance predictions, and the energy impacts generated by SO. To compute the tokens associated with each menu alternative, UO first computes the energy saving, that is, the amount of energy that this alternative saves compared with the predicted user choice (i.e., the individual predicted choice without incentives). The tokens offered in each alternative are then obtained by just dividing the saving by the current TEE. The UO then selects the alternatives that are attractive to the traveler based on a utility function, where coefficients for explanatory variables that represent personal tastes are estimated from historical choices, and values of alternative attributes such as travel time and energy cost are calculated based on the predicted information from Tripod’s SO. Such a personalized menu aligns with the traveler’s interests and makes the system’s architecture sustainable. It encourages energy efficient choices by presenting users with explicit and accurate energy cost information, notifications of accidents, and alternatives. The UO formulation and preference updater is described in more detail by Song et al. (16).

To summarize, there are two optimization cycles: SO optimizes the entire transportation system at every roll period, that is, every 5 min, whereas the UE optimizes in real time an individual menu for each trip request. UE also keeps track of Tripod users’ preferences from their menu selections. In addition, UE provides the

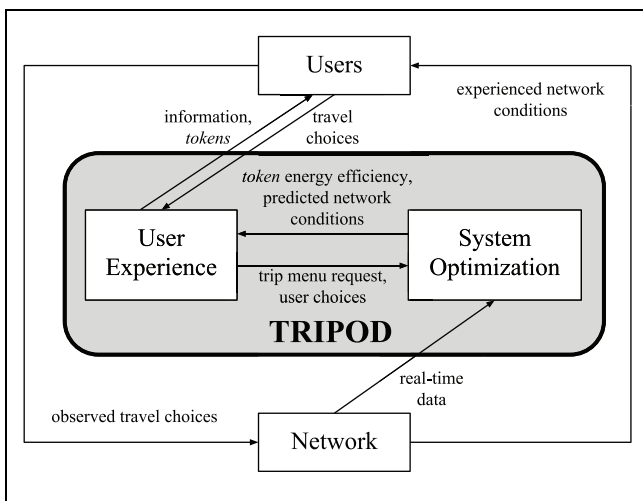


Figure 2. Tripod architecture (15).

information about the updated preferences of Tripod users to SO, for better predictions of SO strategies. For more details on the overall Tripod architecture and the UE optimization framework the reader is referred to Lima Azevedo et al. (15) and Song et al. (16), respectively.

System Optimization Architecture

In summary, SO executes the following operations:

1. It estimates the current state of the multi-modal transportation network.
2. It predicts the state of the network given different token awarding strategies, that is, different TEE values.
3. It estimates the energy savings based on the predicted network conditions for the different token awarding strategies.
4. It provides the UE with a system-wide optimum TEE value in relation to energy savings per token (15).

To do so, SO builds upon a state-of-the-art real-time simulation-based dynamic traffic assignment model, called DynaMIT (Ch.10 of (17)), to provide predictions of the multi-modal network performance, considering how users respond to the information and incentives provided. SO also builds upon TripEnergy (18), a model that estimates the energy impacts of traveling. The next section describes the extensions to DynaMIT to model multiple modes (including transit, carpooling, walk, etc.) and to incorporate the behavioral response of users to information and incentives from Tripod. Carpooling is here a private mode which consists in two travelers with the same origin-destination (OD) and departing in the same 5-min interval choosing to travel with just one private vehicle.

The four steps above are carried out at every roll period (typically 5 min, but it can be longer for more complex networks). The obtained TEE maximizes the system-wide energy savings, based on predicted traffic conditions and energy savings in a future prediction horizon. This is achieved by performing a simulation-based optimization in real time that consists of three major components: the supply simulator, the demand simulator and the system optimizer. The first two components are the supply and demand simulators that interact to simulate the multi-modal system-wide response to different TEE values. The system optimizer searches for the optimal TEE based on the simulated system response. The demand and supply simulators of DynaMIT are extended with new functions. First, modes other than private cars are included. Second, the demand simulator

is extended with simulated user optimization (SUO), which simulates the user optimization of the UE, that is, the generation of the menu of the alternatives shown by the Tripod app, including the tokens allocated to the alternatives. SUO is important to simulate accurately the response of the Tripod users to tokens. Third, the supply simulator is extended with energy estimation, which allows the computation of the energy consumption of the whole system, as well as for each travel alternative. Figure 3 shows the SO architecture and how the three components are integrated to produce the optimal TEE.

At the beginning of a roll period, a state estimation is performed to estimate the current state of the system. The state estimation takes historical demand/supply parameters as starting values, considers real-time events such as accidents, and performs online calibration against real-time measurements such as counts, speed, or travel time measurements. The output from the state estimation is an estimate of the current network state, OD trips, and behavioral parameters governing travelers' choices, including Tripod users' responses to tokens. Within the state estimation, the extended demand and supply simulators interact to produce estimated traffic conditions. In the supply simulator, vehicle trajectories from the supply simulation are fed to the energy estimation module to produce energy consumption estimates. In the demand simulator, SUO receives trip requests from simulated Tripod users and produces personalized menus to simulate users' response. SUO allocates tokens based on TEE generated in the previous roll period. The optimization module is then triggered with the estimated network state as an input. DynaMIT predicts traffic and energy conditions for the future prediction horizon, by making the supply and demand simulators interact for different candidate values of TEE. The SO then chooses the best TEE, which will be employed in the next roll period.

The next section describes the multi-modal extensions of the demand and supply simulators of DynaMIT, including the extensions needed to model the response of the Tripod users.

Multi-Modal Transportation Demand and Supply Models

This section describes the multi-modal demand and supply simulators of DynaMIT that are used in the state estimation and prediction modules described previously.

Multi-Modal Demand Simulator

The multi-modal demand simulator employs disaggregate and aggregate representations of demand in relation to both travelers and vehicles (passenger car equivalents).

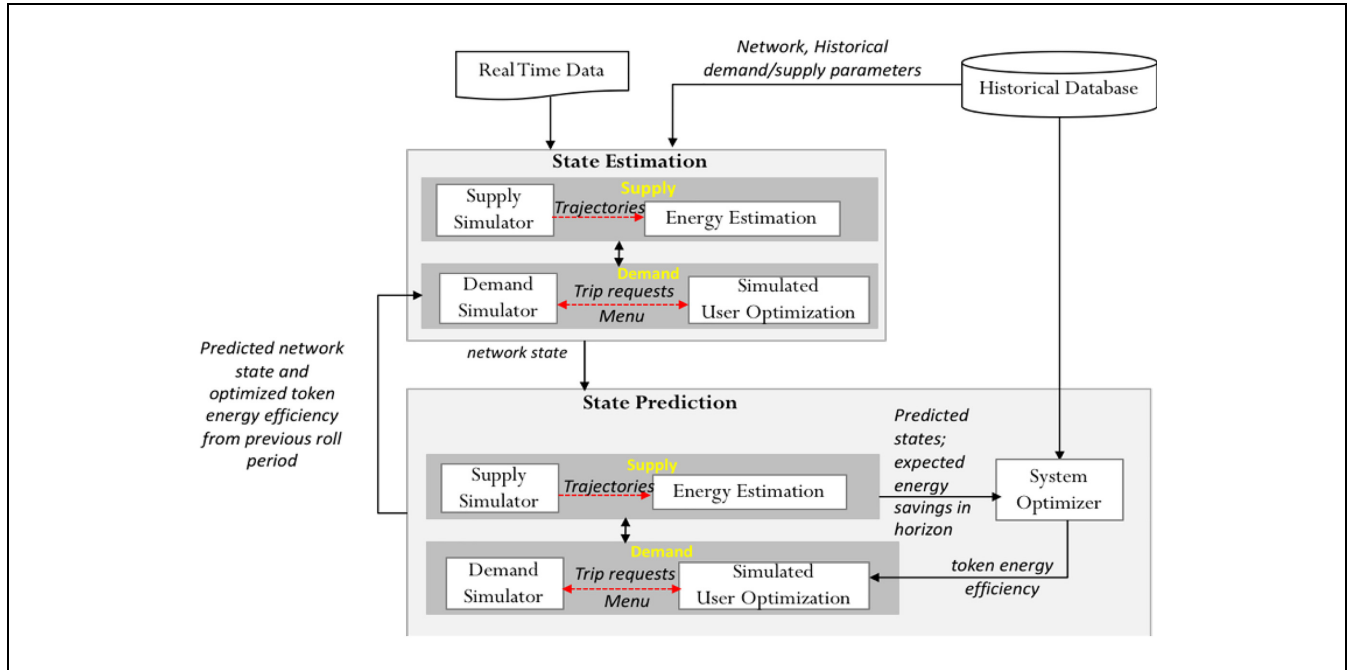


Figure 3. The System Optimization architecture.

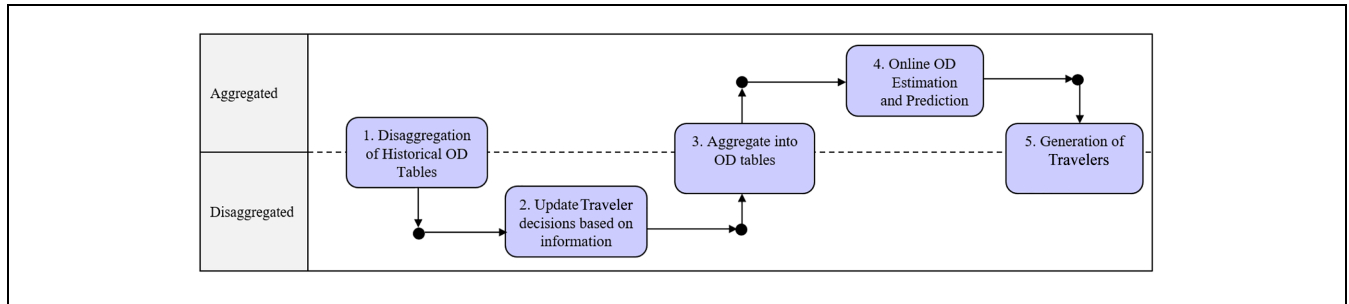


Figure 4. Multi-modal demand simulator.

The disaggregate representation is used to model individual travelers' pre-trip and en-route decisions, including response to information and tokens. An aggregate representation in the form of time-dependent OD matrices (expressed in passenger car equivalents per time interval and travelers per time interval) is also used to estimate and predict multi-modal OD demands.

Figure 4 presents the demand simulator flow diagram. The historical information consists of mode-wise time-dependent OD demand matrices specified in relation to traveler trips.

In the first step, the historical OD matrices are disaggregated to generate a population of travelers who are assigned a habitual route, mode, and departure time. Next, a pre-trip behavioral update is performed, where each traveler updates her choice of mode, route, and departure time based on information of prevailing traffic

conditions and tokens awarded to alternatives (in case of Tripod users). The pre-trip choice is formulated as a nested logit model (19) whose structure is given in Figure 5 (DT refers to departure time interval). The specification of the choice model involves attributes such as travel time, travel cost, and monetary value of tokens awarded, as well as alternative specific constants. For example, the utility of an arbitrary path p under the mode-change (to car) and path-change nest for a habitual transit traveler n with a habitual departure time interval h is given by:

$$U_{np} = \beta_{n-TT} TT_{ph} + \beta_C (C_{ph} - \alpha_{np} \gamma TK_{nph}) + \varepsilon_{np} \quad (1)$$

where β_{n-TT} is the travel time coefficient generated based on a log-normal value-of-time distribution, β_C is the cost coefficient, TT_{ph} is the predicted (or historical, depending

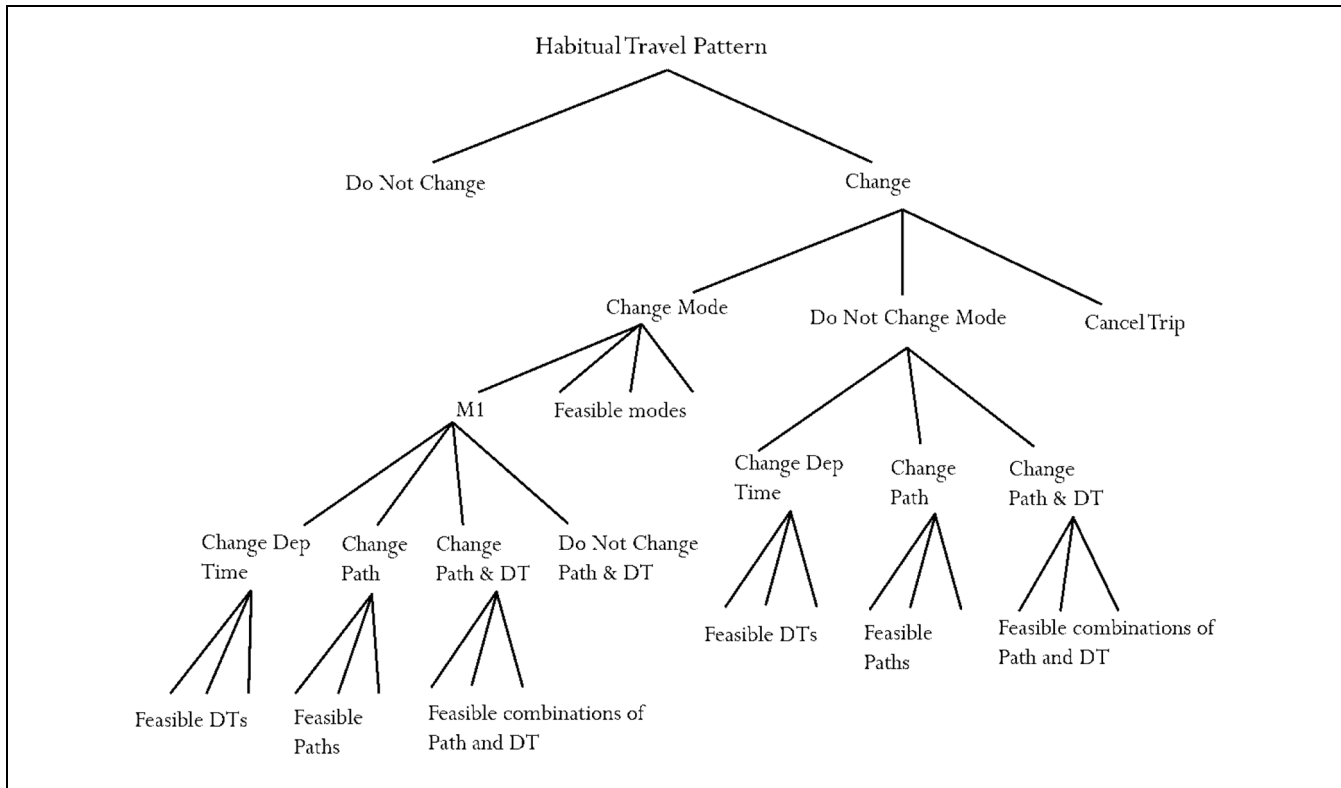


Figure 5. Structure of pre-trip behavior update model.

on whether the traveler has access to information) travel time on path p in time interval h , C_{ph} is the monetary cost, γ is the market value of the *token*, α_{np} is a unit-free *token* value inflation/deflation factor, TK_{nph} is the number of *tokens* allocated to individual n for using path p in interval h , and ε_{np} is a random error component that is *independently and identically distributed* Gumbel distributed.

Step 2 yields an updated list of travelers, which are aggregated in Step 3 back into mode-wise OD matrices in relation to traveler trips. For the private vehicle modes, the ODs in relation to traveler trips are converted to vehicle trips using an average occupancy. The fourth step is OD estimation utilizing the most recent surveillance data from the network. This involves adjusting or estimating OD demands so as to minimize the difference between simulated and observed traffic counts. The OD estimation module makes use of the supply simulator discussed in the next section and results in the estimated number of private vehicle ODs (vehicle trips). These are then used to compute estimates of mode-wise private vehicle ODs (vehicle trips) based on historical modal splits which, in combination with the historical transit ODs, yield the estimated mode-wise OD demands in traveler trips.

These are used to generate the final traveler population for the current estimation interval.

SUO

For the estimation and prediction to be accurate, it is necessary to simulate, within the demand simulator, what will be the options Tripod will provide to the users. This is the role of SUO. The inputs to SUO include all travel options available for a given origin/destination/departure time (from DynaMIT), TEE (from previous roll period or optimization trial value), and Tripod users' characteristics and preference parameters. Three steps are involved in generating a personalized menu of travel options with tokens.

1. For a specific user n , the number of tokens assigned to travel option i ($i = 1, \dots, C_n$) is

$$\max\left(0, \frac{E_{n0} - E_{ni}}{e}\right), \forall n, \forall i \in C_n. \quad (2)$$

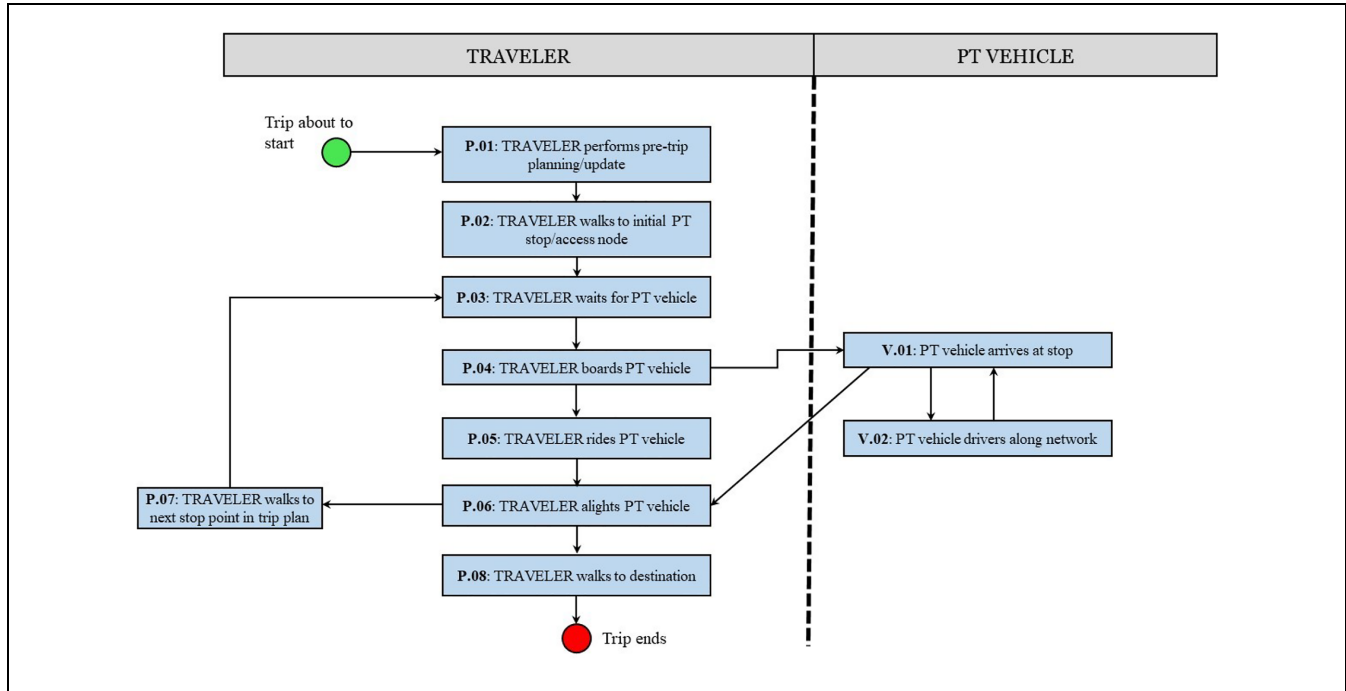


Figure 6. Traveler movement: stages of a transit trip.

E_{ni} is the energy consumption of travel option i for user n , E_{n0} is the expected energy consumption of user n without tokens, $\sum_{i=1}^{C_n} E_{ni}P_{ni}$, where P_{ni} is the probability of user n choosing option i without tokens, and e is the TEE.

2. A personalized menu (a subset of travel alternatives out of all travel alternatives C_n) is generated based on choice probabilities of travel alternatives with tokens assigned in Step 1. SUO maximizes the expected choice probability across the M options on the menu by solving the following problem.

$$\max_{\sum_{i=1}^N x_{ni} \leq M, x_{ni} \in \{0, 1\}} \sum_{i=1}^{C_n} P_{ni}^* x_{ni}. \quad (3)$$

The binary decision variable x_{ni} denotes whether to include option i or not in the menu. P_{ni}^* is the probability of user n choosing option i with tokens. The solution is simply to pick the top M options by sorting P_{ni}^* .

3. Remove tokens assigned to options not on the menu generated in Step 2.

Multi-Modal Supply Simulator

The supply simulator of DynaMIT is mesoscopic (Sec. 1.4.3 of (17)), that is, individual vehicle movement is simulated, but in a simplified manner. The simulator captures traffic dynamics and evaluates the performance of the network, including formation and dissipation of queues, spillback effects, impacts of accidents, and bottlenecks. It represents traffic dynamics using macroscopic speed-density relationships and queuing theory. The multi-modal supply simulator derives largely from the original (Ch.10 of (17)) with two key enhancements: (1) traveler movement: transit travelers agents are introduced and (2) buses: a controller has been developed to manage the fleet of buses.

The various stages of a transit trip are shown in Figure 6 (PT refers to public transit). There are two main actors: (1) traveler and (2) vehicle.

The bus controller operates the fleet of buses on the network (this could involve fleets of multiple operators). It obtains from a database a list of bus lines with the related stops and frequencies/headways.

The existing vehicle movement models are adapted to capture appropriately the dwelling of buses at stops and their impact on the traffic stream. Since DynaMIT naturally models spillback effects and congestion through a

queuing part at the downstream end of each segment, all segments containing a bus stop are split at the location of the stop to capture the queuing caused by bus dwelling.

The movement of buses can be split into two parts: (i) movement between bus stops and (ii) movement into and out of a bus stop. The movement between stops is similar to that of cars: that is, the buses are moved using the speed-density model and the queuing model. Regarding the movement into and out of a stop, when the bus reaches the end of the segment with a bus stop to serve, it moves into the bus stop if the residual capacity is non-zero. Otherwise, it starts queuing and affects the vehicles behind it. When the bus stop's residual capacity allows the bus to enter the stop, the queue starts dissipating. After serving the bus stop, the movement of the bus out of the bus stop depends on the acceptance capacity of the downstream segment. If there is no queue, the bus moves to the downstream segment like any other vehicle. If the downstream segment has zero acceptance capacity, the bus remains in the bus stop until it can move to the downstream segment.

Optimization Formulation

This section provides a high-level formulation of the SO problem, introduced by Lima Azevedo et al. (15), which will be solved through the heuristic method presented in the next section.

The predictions of network state are performed in discrete time steps with a time interval of Δ , called a roll period. During time interval $[t - \Delta, t]$, the computation is performed to predict what will be the network state in the prediction horizon $[t, t + H\Delta]$, where $H \in \mathbb{N}$. The vector of starting times for the roll periods contained in the prediction horizon for time t is here denoted by $\tau = (t, t + \Delta, t + 2\Delta, \dots, t + H\Delta)$. Alternatively, the notation τ is also used to refer to the prediction horizon $[t, t + H\Delta]$, with the specific use evident from the context. The decision variable for SO is TEE, which represents the amount of network-wide energy savings that must be realized by a user to be awarded one token. The TEE is considered to be constant within each roll period. The TEE values related to a prediction horizon are represented by the vector $e(\tau) = (e(t), e(t + \Delta), \dots, e(t + H\Delta))$. The total energy saving predicted within the horizon is denoted by $ES(e(\tau))$.

DynaMIT state estimation (SE): As described in the previous sections, at any time t the SO starts an execution cycle and performs an estimation of the network state using real-world data collected in the previous roll period $[t - \Delta, t]$ as well as historical real-world data for the same time of day. All the parameters describing demand (OD matrices, behavioral parameters, etc.) and supply (link capacities, speed-density function, etc.) are

calibrated to minimize the discrepancy between simulated and real-world measurements. DynaMIT SE also considers the choices of Tripod users, given their individual menu and behavioral parameters, including token related parameters (e.g., sensitivity to tokens).

DynaMIT state prediction (SP): After SE, the SO loop is initiated by running SP. Given $e(\tau)$, the previous network state and supply-demand parameters, SP predicts how the network performance will evolve during the prediction horizon τ , yielding the predicted network states $x(\tau)$, including user trajectories (or parts of the trajectories that lie within the prediction horizon), $v^n(\tau)$.

Energy estimation: Given the predicted network states $x(\tau)$, the predicted user trajectories $v^n(\tau)$, and token efficiencies $e(\tau)$, the total energy savings for the network during the prediction horizon τ , $ES(e(\tau))$ is calculated by TripEnergy (18). It does so by comparing the predicted energy consumption with the baseline consumption (SP simulation with no tokens), as expressed below.

$$ES(e(\tau)) = \sum_{n=1}^N f(v_e^n, \theta^n) - \sum_{n=1}^N f(v_0^n, \theta^n), \quad (4)$$

where v_e^n, v_0^n are the user trajectories that result from providing tokens based on token efficiency $e(\tau)$ and with no tokens, respectively, θ^n are the travel mode parameters (e.g., car design parameters, bus type, driving style, etc.), N is the number of travelers, and $f()$ is the function that computes energy consumption for each user trajectory.

SO loop: The objective of the SO loop is to determine the optimal TEE for the H roll periods within the prediction horizon. For each given TEE value, the SUO module determines the menu of travel alternatives (with tokens) offered to Tripod users on the network; SP updates the demand and calculates the new network states; the TripEnergy module evaluates the energy savings relative to the no-incentive base case. Based on these inputs, the objective function of maximizing network-wide, entire-day energy savings potential is evaluated. The maximization is performed subject to the constraint that the balance of tokens $W(\tau, e(\tau))$ at the end of the prediction horizon is non-negative. Note that $W(\tau, e(\tau))$ is the token balance at the beginning of the prediction horizon minus token consumption during the prediction horizon for a given vector $e(\tau)$ of TEEs. This optimization can be stated as follows:

$$\begin{aligned} & \max_{e(\tau)} ES(e(\tau)) \\ & \text{subject to :} \\ & e(\tau) \geq 0, W(\tau, e(\tau)) \geq 0. \end{aligned} \quad (5)$$

In the current study the TEE is assumed to be constant in the prediction horizon (still time-varying by roll period). This results in a single decision variable and

allows for a simpler search. The continuous interval (decision space) is discretized using a reasonable step size obtained by trial and error. The objective function value for different TEEs can be evaluated in parallel and the optimal solution obtained in a single iteration of the optimization. This is described in detail in the next section.

Online Optimization

The goal of SO is to find the sequence $e(\tau)$ of TEE values for the next prediction horizon that minimizes the energy consumption of the entire transportation system under the token budget constraint. The optimization is performed online and adapts to the evolution of network state. This implies that it does not compute $e(\tau)$ just when SO is launched, but computes it continuously over time, appending new values to it each time. SO is implemented running M instances of DynaMIT in parallel, all controlled by a coordinator, whose role is to (i) synchronize the instances and ensure they work in real time, (ii) pass the information they need, (iii) orchestrate their operations, and (iv) decide at each roll period t what will be the next TEE $e(t + \Delta)$. The value $e(t + \Delta)$ will be communicated to the SUO module and will be used to compute the incentives proposed to the real users during the interval $[t + \Delta, t + 2\Delta]$. This real-time requirement is equivalent to requiring that $e(t + \Delta)$ be computed before $t + \Delta$.

Assume that tokens are granted to travelers on a first-come-first-served basis. Take into account the constraint on Equation 5 by assigning a maximum per-period token budget. Referring to Equation 2, if TEE is too small, tokens will be given away “too easily,” in return for a small energy reduction, to the first travelers making trips in the roll periods. This would prevent it from rewarding travelers who guarantee greater energy savings but arrive later. On the other hand, if TEE is too high, the amount of tokens given to travelers may be too small to affect their behavior. Therefore, finding the optimal value of TEE is not trivial. It is done by heuristically exploring the impact of a set of TEE values within a certain interval.

At each roll period t , each DynaMIT instance performs an estimation phase (SE), followed by a prediction phase (SP). All instances have identical states during estimation, since they all read the same real-time data and historical data. The instances differ during prediction, more precisely, during the prediction performed in a roll period t , when the coordinator instructs each instance to predict the network state in the prediction horizon. Each instance $m = 1, \dots, M$ predicts the effect of a different future candidate TEE value, which is indicated with $e^m(t)$, assuming it applies to the entire prediction horizon

τ . At the end of the prediction, each instance returns the predicted energy consumption $E^m(t)$. The coordinator chooses the instance m^* that predicted the least energy consumption, that is, $m^* = \arg \min_m E^m(t)$. The respective TEE value $e^{m^*}(t)$ becomes the TEE to be employed in the next roll period, that is, $e(t + \Delta) = e^{m^*}(t)$.

The operations for the computation of the sequence $e(\tau)$ are depicted in Figure 7. Let us suppose a roll period of duration $\Delta = 5$ min and prediction interval 15 min, that is, $H = 3$. Let us start to describe the system when it is at time $t = 8:00$, which is the start of the roll period $[8:00, 8:05]$. Before the end of this period, SO must be able to provide the values of $e(t + \Delta)$, that is, the TEE of the roll period $[8:10, 8:15]$. To do so, the following sequence of operations takes place at 8:00:

1. The coordinator triggers all the DynaMIT instances to execute their estimation phases, based on sensor data related to the previous 5 min, that is, $[7:55, 8:00]$ and $e(t - \Delta)$. The goal of executing these estimation phases is to make the internal simulation model consistent with real data. As discussed previously, all the instances have the same internal state in this phase. Observe that the parallel execution of the estimation phases of the instances corresponds to Step 1 in the section “SO Architecture”.
2. The coordinator assigns to each instance m a *candidate* TEE $e^m(t)$.
3. Each instance predicts the evolution of the network in the interval $[8:05, 8:20]$ and returns the predicted energy consumption $E^m(t)$.
4. The coordinator chooses $e(t + \Delta) = e^{m^*}(t)$, where $m^* = \arg \min_m E^m(t)$ and communicates this value to the SUO module, which will use this value to determine the incentives that will be shown in the menus generated during the next roll period $[8:05, 8:10]$.
5. At 8:05, these operations start again, with estimation based on real data related to $[8:00, 8:05]$.

Results

This section evaluates the impact of Tripod optimization on the multi-modal transportation system in relation to energy consumption, mode share, and travel times. To simplify the analysis, first, static scenarios are analyzed, in which a static TEE allocation is fixed and the penetration rate is varied, that is, the percentage of travelers who are Tripod users. The penetration rate is then fixed and the benefits of the online optimization over the static settings can be studied.

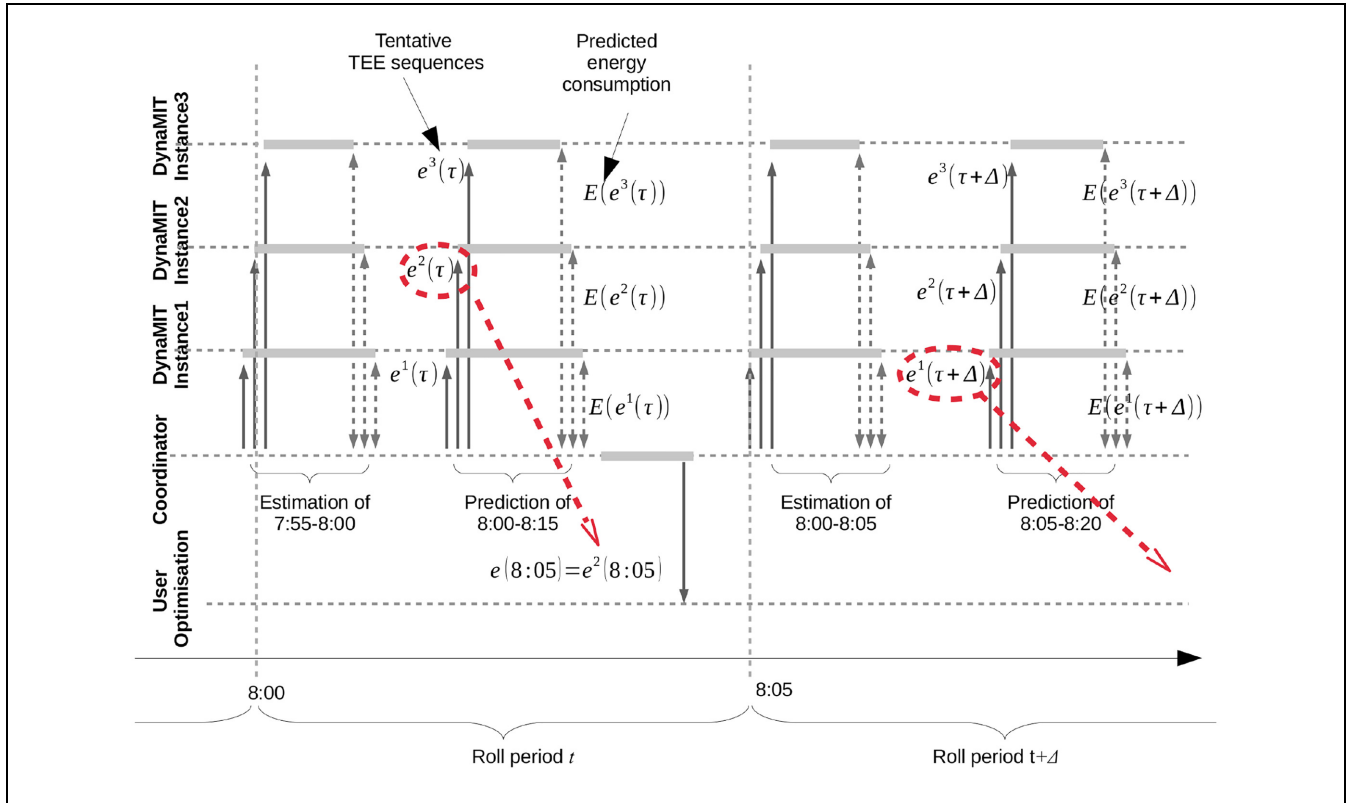


Figure 7. The online optimization procedure.

Simulation Scenario

The experiments are conducted on the Boston Central District (CBD) network with 843 nodes, 1879 links, 3075 segments, and 5034 lanes including both highways and arterials between 6:00 and 9:00 a.m. Note that the focus is restricted to the peak hours when the transportation system energy consumption is maximum. As expected, the energy savings would be lower in other time intervals, which are not shown for lack of space.

The total number of travelers is 47,588. The parameter values in the utility function (Equation 1) are postulated as follows: $\beta^T = -0.01$, based on empirical studies in the literature, the value of time (VOT) is assumed to be log-normal distributed with a mean of \$18 per hour and standard deviation of \$5 per hour, the cost parameter β_n^C of an individual n is calculated based on a sampled VOT from the log-normal distribution. The monetary value of a token is $\gamma = 0.50$ \$. Tokens instead of dollars are used, as the full design of Tripod includes a marketplace where tokens can be exchanged and their monetary value determined by the market. In the current implementation, the marketplace is not in place and thus a fixed value is assumed. The perception parameter $\alpha_n = 1$ for each individual n . As for the parameters of SO, the roll period

length is $\Delta = 5$ min and there is a token budget constraint of 20,000 per roll period.

Impact on Multi-Modal Transportation

In this section, mode shares, average personal energy consumption, average personal travel time, and token consumption with respect to different penetration rates (percentage of travelers using Tripod) are presented. Note that the energy saving of Tripod depends on a myriad of factors, including but not limited to the penetration rate, the sensitivity of travelers to incentives, the spatial-temporal distribution of the demand, and the availability of attractive transit options. The penetration rate is a major factor that is directly related to the investment in the app deployment and thus the focus of the following computational tests. In contrast, other factors are less controllable, for example, the spatial-temporal distribution of demand and the sensitivity of travelers to incentives mainly depend on the broader economic, social, and demographic developments and the availability of attractive transit options requires significant capital investment besides the app. Note that, for the sake of simplicity, the possibility that a Tripod user might opt

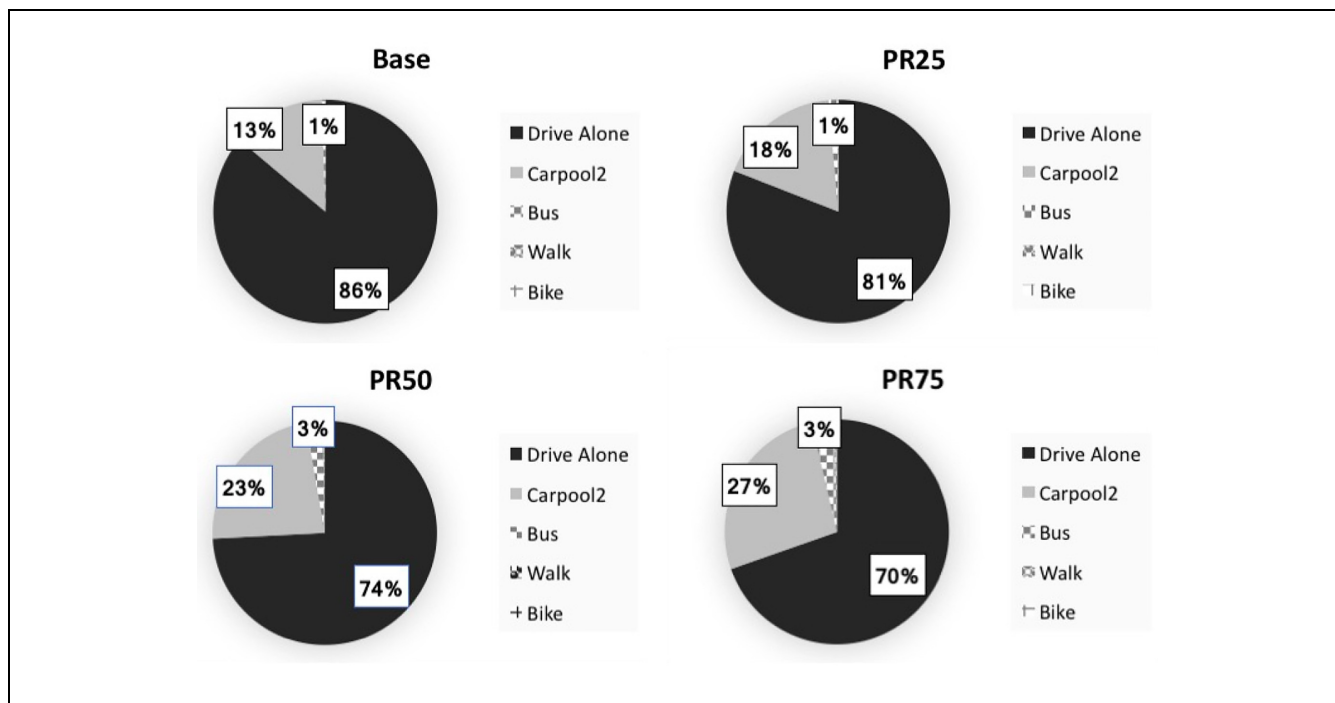


Figure 8. Mode share with various penetration rates (PR %).

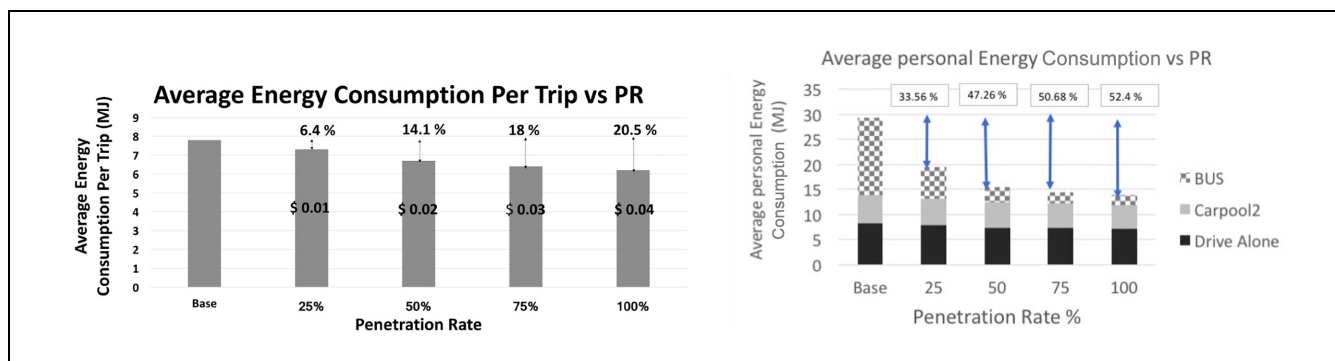


Figure 9. Average energy consumption per trip: (a) overall and (b) mode-specific. Monetary values of energy savings per trip at \$3.00/gallon are shown on the bars (left).

out is not modeled. However, if a user does not find the propositions from Tripod attractive, she will simply ignore them, thus not contributing to the energy savings to be shown later.

Figure 8 shows the mode share at various penetration rates (PRs) of Tripod. Not surprisingly, higher PR results in higher share of greener modes, that is, carpool, bus, walk, and bike. The increase of carpool share is more significant than that for bus, walk, bike, or any combination of these, probably because of the travel time advantage of carpool compared with the other green modes as no pick-up or drop-off travel time is accounted for in carpool.

Figure 9a shows the average personal energy consumption per trip in megajoules (MJ) as a function of the PR of Tripod. Not surprisingly, as more travelers are incentivized (higher PR), the energy saving per person is higher. There is also an indication of the saturation effect, in that the rate of the change decreases with the PR. For example, an additional 4% saving is achieved when the PR increases from 50% to 75%, while an additional 2.5% saving is achieved when the PR increases from 75% to 100%. Figure 9b shows the personal energy consumption, that is, the energy consumed by an individual (different from Figure 9a, in which the energy is per trip). The breakdown by major mode (bus, carpool, and

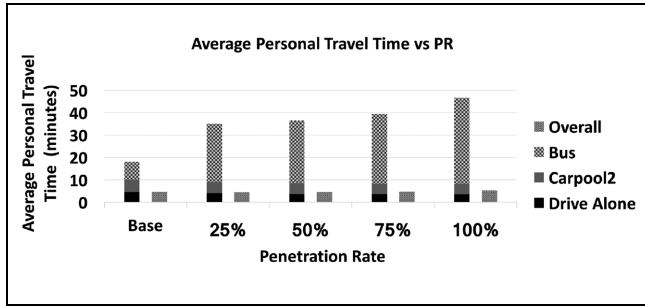


Figure 10. Average personal travel time.

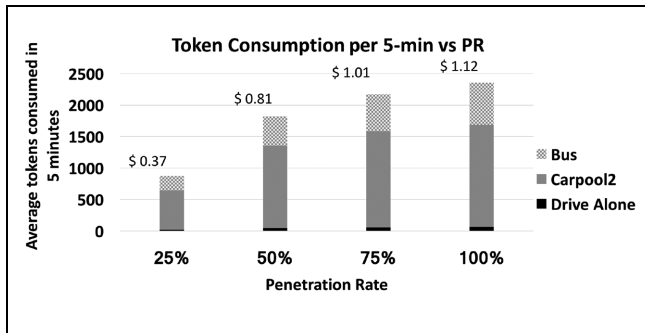


Figure 11. Token consumption per 5 min. The number above each bar is the average perceived monetary value of tokens per trip.

drive-alone) shows that that average personal energy consumption decreases for all three major modes. The personal energy consumption of the two private vehicle based modes (drive-alone and carpool) decreases because of improved traffic conditions, that is, lower travel times (see Figure 10). The personal energy consumption of the bus mode decreases because of higher bus ridership. Note that bus schedule is exogenous in the system and thus bus vehicle energy consumption almost remains the same regardless of incentives. With higher ridership, the bus vehicle energy is shared by more riders, and thus the energy contribution of each decreases. Observe that mode switching is not the only source of energy savings: even the users who drive alone may contribute energy savings by taking more energy efficient routes.

Figure 10 shows the average personal travel time as an increasing function of the PR of Tripod. Note that travel time is not an objective of the optimization, and thus such an increasing trend is not surprising. A breakdown by major mode (bus, carpool, and drive-alone) shows that the travel time of the two private modes (drive-alone and carpool) decreases with the PR, while that of bus riders increases with the PR. Note that travel time of a bus rider includes access, egress, and in-vehicle travel times. The travel time decrease of private modes is because of fewer vehicles on the road, resulting from

mode shift to bus, walk, and bike. The average travel time increase of bus riders is because of more travelers incentivized to take bus, despite its longer access, egress time, or both. The personal travel time averaged over the system is shown beside the bars. Collectively, at lower PRs, the mode shares of drive-alone and carpool are higher and thus the overall average travel time is lower, while at higher PR, the longer travel time of bus riders dominates. However, it should be noted that those who switch to transit and thus have increased personal travel times do so at their own will, that is, the fact that they have switched indicates that they perceive that the incentives more than compensate for their travel time losses.

Figure 11 shows the token consumption by mode as a function of the PR of Tripod. The total consumption increases with PR, as expected. Carpool has the highest token consumption, followed by bus. Both have high energy saving potentials, and yet carpool is in general more attractive than bus because of shorter travel time. Drive-alone has the least token consumption because of the least energy saving potential through route choice.

Average monetary values of the consumed tokens per trip as perceived by the travelers are presented in Figure 11 as numbers above the bars. Note that here the “tokens per trip” are obtained by dividing the total number of distributed tokens by the number of trips. We also compute the “tokens per Tripod-trip” where a Tripod-trip is a trip of a traveler who accepted a Tripod option, thus consuming a positive amount of tokens. The perceived monetary values of tokens per Tripod-trip are, as expected, higher: \$2.45, \$2.68, \$2.76, and \$2.86 for the PRs in the figure. Monetary values of energy savings estimated at an assumed fuel price of \$3/gallon are shown in Figure 9a as numbers on the bars. It should be noted that the perceived monetary value of a token is different from the cost of providing the token, for example, if the tokens are exchanged for goods as in-kind gifts from participating vendors, the cost of the token to the public is in fact zero. Similarly, the cost saving from a consumption reduction of one gallon of fuel is not necessarily the same as the prevailing market price, if the goal is to evaluate the societal cost of consuming one gallon of fuel, especially when the market does not have an adequate mechanism to reflect the external costs of fuel consumption such as environmental costs. Therefore, these monetary values are presented for information purpose and should not be used directly to do a benefit-cost analysis.

Performance of the Online Optimization

Figure 12 shows the benefit of the dynamic aspect of the online optimization strategy, which continuously recomputes the TEE $e(t)$ over time to adapt to the network

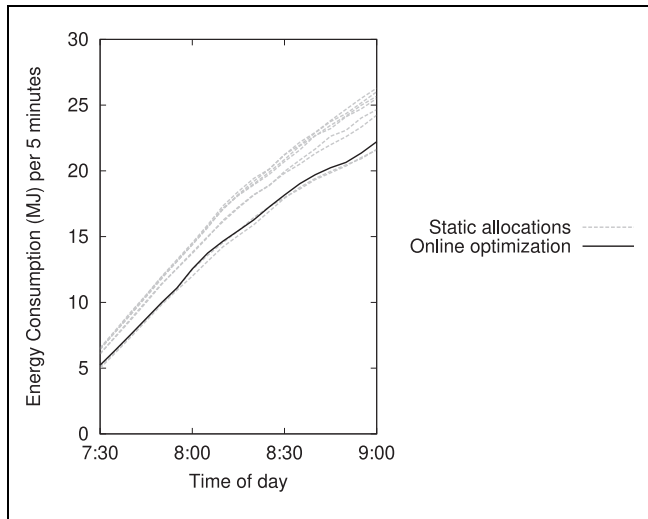


Figure 12. Online optimization versus static allocations.

evolution. For these results $M = 8$ parallel instances are employed. A logarithmic search is implemented in the interval $TEE = 1$ and $TEE = e_{MAX} = 2000$, by assigning to each instance $m = 0, \dots, M - 1$ a candidate $TEE e^m(t)$ such that $\ln e^m(t) = m \cdot (\ln e_{MAX}) / (M - 1)$, which results in the following discrete values: $\{1, 3, 8, 26, 77, 228, 675, 2000\}$. The overall energy consumption is compared with the case of static allocation, in which the $e(t)$ does not change along the time. Different possible values of static $e(t)$ belonging to the same discrete set above are tested. Note that, in reality, if we were to implement a static TEE policy, only one static allocation can be implemented at a time and it is impossible to know in advance what is the best value to apply. On the contrary, online optimization does not require this a-priori knowledge, it adapts automatically to the current conditions of the network, guaranteeing energy reduction.

It should be noted that the optimization is quite demanding in relation to computational resources. However, at least in the scenario considered, with a roll period of 5 min and a prediction horizon of 15 min, the framework has been shown to be scalable, that is, the entire SO operations described in this paper have been done in real time. This means that at each roll period, it is possible to complete the computation of the next TEE before the beginning of the next roll period. The machine used in this study is a PowerEdge T630, equipped with two Intel Xeon E5-2695 v4 2.1GHz processors, 128GB of memory and an SSD disk.

Conclusion

This paper describes the implementation of Tripod's (15) optimization framework. Tripod is a novel demand

management system that incentivizes travelers in real time to reduce the overall energy consumption of a transportation system, under an incentive budget constraint. The optimization tackled in this paper is challenging since it is performed online, it includes several modes of transportation, it computes personalized incentives, and it is guided by the current state of the network and the predicted state. A methodology is proposed to implement a heuristic method that reduces this complex problem to the search of a single value, called token energy value. Predictions are based on multi-modal traffic simulation and models of individual travel decision making, including the response to incentivization. Simulation results show that this system is potentially effective in reducing energy consumption under different scenarios and that large benefits come from the dynamic nature of our optimization.

While this paper has shown Tripod's potential for a specific setting, the analysis was limited to (i) a small network, which does not capture the full extension of travel patterns, network complexity, and computational burden of large networks, (ii) the morning peak period, thus ignoring some behavioral time-dependencies in individual decision making and the budget allocation across longer periods, (iii) a single configuration of Tripod, as one can easily design a system with different user segment participation rates, menu generation constraints, a relaxation in having just a single token energy value or even subsets of choice dimensions to be incentivized, and (iv) a single system objective of energy saving while other viable objectives such as travel time saving and reliability improvement are not accounted for. For this, the authors are working on integrating the proposed framework with an agent-based simulator (20) for impact validation and scenario exploration. Field trials are also being pursued to evaluate the feasibility and the effectiveness of Tripod in realistic settings.

Acknowledgments

This research was carried out under the U.S. Department of the Environment's Advanced Research Projects Agency-Energy (ARPA-E) TRANSNET Program, Award Number DE-AR0000611.

Author Contributions

Study conception and design: SG, CLA, MB-A; simulation setup: AA, RS, YS, SA, DS; analysis and interpretation of results: AA, HG; manuscript preparation: SG, CLA, AA, RS, HG, SA. All authors reviewed the results and approved the final version.

References

1. Ben-Akiva, M., A. De Palma, and K. Isam. Dynamic Network Models and Driver Information Systems.

- Transportation Research Part A: General*, Vol. 25, No. 5, 1991, pp. 251–266.
2. Viti, F., S. F. Catalano, M. Li, C. Lindveld, and H. van Zuylen. An Optimization Problem with Dynamic Route-Departure Time Choice and Pricing. Presented at 82nd Annual Meeting of the Transportation Research Board, Washington, D.C., 2003.
 3. Wang, Y., W. Y. Szeto, K. Han, and T. L. Friesz. Dynamic Traffic Assignment: A Review of the Methodological Advances for Environmentally Sustainable Road Transportation Applications. *Transportation Research Part B: Methodological*, Vol. 111, 2018, pp. 370–394.
 4. Lindsey, R. Reforming Road User Charges: A Research Challenge for Regional Science. *Journal of Regional Science*, Vol. 50, No. 1, 2010, pp. 471–492.
 5. Tsekeris, T., and S. Voß. Design and Evaluation of Road Pricing: State-of-the-Art and Methodological Advances. *Netnomics*, Vol. 10, No. 1, 2009, pp. 5–52.
 6. Kumar, V., C. R. Bhat, R. 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*, 2016. 2543: 166–175.
 7. 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.
 8. Merugu, D., B. S. Prabhakar, and N. Rama. An Incentive Mechanism for Decongesting the Roads: A Pilot Program in Bangalore. *Proc., NetEcon, ACM Workshop on the Economics of Networked Systems*, Stanford, CA, 2009.
 9. Zhu, C., J. S. Yue, C. V. Mandayam, D. Merugu, H. K. Abadi, and B. Prabhakar. Reducing Road Congestion Through Incentives: A Case Study. Presented at 94th Annual Meeting of the Transportation Research Board, Washington, D.C., 2015.
 10. Lu, F. *Framework for a Lottery-Based Incentive Scheme and its Influence on Commuting Behaviors: An MIT Case Study*. PhD thesis. Massachusetts Institute of Technology, 2015.
 11. Meyer, E., J. Chernick, and B. Dalton. *Matching System for Ride Reservation Platforms*. 2012, US Patent 8285570.
 12. Poslad, S., A. Ma, Z. Wang, and H. Mei. Using a Smart City IoT to Incentivise and Target Shifts in Mobility Behaviour – Is it a Piece of Pie? *MDPI Sensors*, Vol. 15, No. 6, 2015, pp. 13069–13096.
 13. *Empower Project*. <http://mobility-apps.eu/>. Accessed 1 November 2018.
 14. 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.
 15. Lima Azevedo, C., R. Seshadri, S. Gao, B. Atasoy, A. P. Akkinipally, E. Christofa, F. Zhao, J. Trancik, and M. Ben-Akiva. Tripod: Sustainable Travel Incentives with Prediction, Optimization, and Personalization. Presented at 97th Annual Meeting of the Transportation Research Board, Washington, D.C., 2018.
 16. Song, X., M. Danaf, B. Atasoy, and M. Ben-Akiva. Personalized Menu Optimization with Preference Updater: A Boston Case Study. *Transportation Research Record: Journal of the Transportation Research Board*, 2018. 2672: 599–607.
 17. Barcelo, J. *Fundamentals of Traffic Simulation*. Springer, New York, 2010.
 18. Needell, Z. A., and J. E. Trancik. Efficiently Simulating Personal Vehicle Energy Consumption in Mesoscopic Transport Models. *Transportation Research Board: Journal of the Transportation Research Board*, 2018. 2672: 163–173.
 19. Ben-Akiva, M., and S. Lerman. *Discrete Choice Analysis: Theory Application to Travel Demand*. MIT Press, Cambridge, MA, 1985.
 20. Chen, S., I. de Lima, Z. Needell, A. Araldo, A. Akkinipally, J. Trancik, M. Ben-Akiva, N. Fournier, A. Deliali, E. Christofa, and C. Lima Azevedo. Integrated Simulation of Activity-Based Demand and Multi-Modal Dynamic Supply for Energy Assessment. *Proc., 21st IEEE International Conference on Intelligent Transportation Systems*, Hawaii, IEEE, New York, 2018.
- The Standing Committee on Transportation Demand Management (ABE50) peer-reviewed this paper (19-03092).*