Distributed Stochastic Approximation for Adaptive Frequency Allocation in Subway Networks

Felisa J. Vázquez-Abad[1]
Department of Computer Science and OR, University of Montreal,
P.O. Box 6128, Downtown Station, Montreal, Quebec H3C 3J7, CANADA
also *Principal Fellow*, Department of Electrical and Electronic Engineering, University of Melbourne

Lourdes Zubieta[2]
Department of Business Administration, Bishop's University
Lennoxville, Quebec J1M 1Z7 CANADA

Abstract

We propose to build an intelligent mechanism within a computerized subway system that will adjust frequency of services according to the observed variable demand. Operating costs are generally associated with the number of one way trips per day per line, while social cost is measured as the total waiting time of all passengers per day. The proposed transit system will automatically seek optimal performance minimizing operational costs. The structure of the control variables (headways) and their relationship with other model parameters make the problem suitable for *ersätz* estimation where the gradient w.r.t. a control variable is estimated in terms of local gradients w.r.t. other variables.

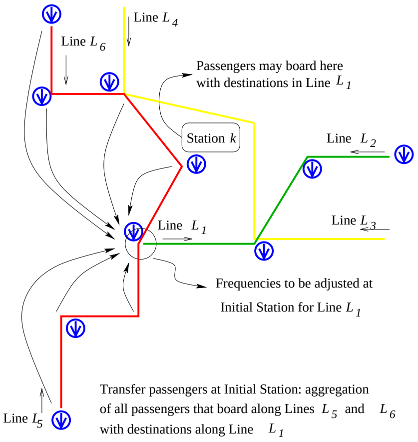

Figure 1: Model of Transfer Station, Eastbound Platform.

1 Introduction

Consider a subway network composed of line itineraires described as a sequence of consecutive subway platforms where passengers board a line for the first time or transfer to other lines. By a "line" we mean here a specific direction. For example, in Figure 1 the initial platform on Line L_1 (say platform p) is in a transfer station where there are three other platforms (belonging to Lines L_2, L_5 and L_6). Passengers arrive at platform p according to three arrival sources: incoming passengers from outside, and transfer passengers arriving into the station from lines L_5 and L_6 (passengers arriving from line L_2 are assumed not to board trains on plaform p). In typical railroad problems where travel times are large and frequencies small, it is important to establish the epochs of the first outgoing train every day for each line, as explained in [2] and [4] . Travel times are asumed deterministic in transfer optimization models, so that timetables are completely determined by the initial departure times and the period of the timetable [1]. In [4] stochastic aproximation is used to establish the optimal threshold for *waiting at a transfer platform*: trains wait for possible transfer passengers from other lines, as long as they do not wait longer than a given threshold that ensures satisfaction of the timetable for future stations. In a subway network trains are frequent enough that synchronization may even hinder performance [7]. In the present paper, it is assumed that there is no synchronization between lines, i.e., trains do not wait for transfer passengers. Each line L_l has a "headway" μ_l (mean times between consecutive trains) associated with it. The model incorporates fluctuations in travel times as well as the time it takes passengers to board the trains.

Given the transit demand, the expected cost $F(\bar{\mu})$ of the network is the expected operating cost per day plus the expected cumulative wait of passengers in platforms. For each line, an increase in frequency will necessarily decrease the average wait of passengers, and $F(\cdot)$ is convex in each component μ_ℓ. Stochastic approximation can therefore be applied to this problem to improve

---

[1] Member of the centre de recherches sur le transport (CRT) and the groupe d'études et de recherche en analyse des décisions (GERAD), University of Montreal. This work was supported by NSERC Grants of the Government of Canada.

[2] Member of the CRT, University of Montreal

performance. Let $\{\varphi_n\}$ be a sequence of consecutive estimators of $\nabla_{\bar{\mu}} F(\bar{\mu})$, each obtained using D consecutive days and adjust the control parameter using:

$$\bar{\mu}(n+1) = \bar{\mu}(n) - \epsilon\, \varphi(n). \tag{1}$$

Suppose that the measurements $\varphi(n)$ were taken but no updates were performed. That is, $(n) \equiv \bar{\mu}$. If for this process $\mathsf{E}[\varphi(n) | \bar{\mu}(n) = \bar{\mu}] = G(\bar{\mu})$, for a bounded and continuous function $G(\cdot)$, then under verifiable conditions (see [5]), as $n \to \infty$ the time-varying control sequence $\{\bar{\mu}(n)\}$ converges in a weak sense to any stable point of the ODE:

$$\frac{d\bar{\mu}(t)}{dt} = -G[\bar{\mu}(t)],$$

that is, points $\bar{\mu}^*$ which satisfy $G(\bar{\mu}^*) = 0$. To minimize F, one usually seeks to build unbiased estimators for which $G(\bar{\mu}) = \nabla_{\bar{\mu}} F(\bar{\mu})$, using data collected from the sensors at each platform. The advantage of such learning algorithms is that once the estimation code is programmed in the sensors at the platforms, the controls themselves adjust their values to track optimal performance even when traffic patterns change due to seasonal or slowly changing environments.

But direct estimation of the gradient of the expected cost per day may prove to be a challenging and computationally impractical goal. When the control parameters μ_ℓ are scalable, we proved in [8] that at any transfer platform, gradient estimation can be performed using the derivatives w.r.t. other control parameters. If parameters are not scalable, we showed how correction factors can be estimated. In this paper we extend our analysis to the whole network and show that the estimation of sensitivities can be simplified by distributing the computation per platform. These are called "ersätz" estimators. While they may be less efficient than some direct estimates, ersätz estimation allows a simple implementation of counters. The local data is then broadcast towards the computers at the initial platforms on each line, which in turn put together the received information to update asynchronously using (1) individually for each component.

In Section 2 we state the model, Sections 3 and 4 introduce the ersätz model and gradients, and Section 5 discusses the implementation of the learning algorithm.

2 Global Model per Day

Suppose that there is a total of L lines denoted $L_1, \ldots, L_L$. A line is defined as a list (or sequence) of platforms $p \in L_\ell$. Notice that given a platform p, there is a unique l such that $p \in L_l$. Denote by $p_1 \,{-}{>}\, p_2$ the relationship that platform p_1 immediately precedes p_2 in the list. To ease notation, it is assumed that there is a unique sequence of platforms (may be changing lines) per origin-destination (o, d) pair, where both o, d are platforms. This assumption can be relaxed to the case where passengers take more than one possible path, assuming that each arriving passenger at (o, d) takes each of the possible paths randomly and independently of the choice of other passengers.

Every day is divided into several time periods (or segments $s = 1, \ldots, S$) that correspond to different travel demand patterns. The length of each segment is T_s, and $\tau_s = T_1 + \ldots + T_{s-1}$ is the epoch of the start of the s-th day segment of length T_s.

Assumption 1 *Passengers with destination platform d arrive at the origin platform o according to a Poisson process on each day segment. The rates are given by $\lambda_{o,d}(s), s = 0, \ldots, S-1$. All Poisson arrival processes are independent and independent of the train departure processes at the platforms.*

At the initial platform p_1 on Line L_l trains depart with an average time $\mu_{l,s}$ (the headway), that is, $V_j(p_1) = V_{j-1}(p_1) + \mu_{\ell,s} + \Delta_j(p_1)$, for $\tau_s \le V_{j-1}(p_1) < \tau_s + T_s$. Fluctuations around $\mu_{l,s}$ model the time required for conductors to prepare their gear. Once a train departs from any platform $p_k \in L_l$, the time until departure of the same train on platform p_{k+1}, with $p_k \,{-}{>}\, p_{k+1}$ is modeled as the travel time plus the time the doors remain open while passengers board the train:

$$V_j(p_{k+1}) = V_j(p_k) + T(p_k, p_{k+1}) + \Gamma_j(p_{k+1}),$$

where $T(p_k, p_{k+1})$ is the average travel time from p_k to p_{k+1} plus the waiting time at platform p_{k+1} prior to departure. Our model assumes that at each platform p the fluctuations $\Gamma_j(p)$ around the mean are i.i.d. In addition, our model assumes that $\mu_{l,s}$ is a scale parameter of the distribution of these fluctuations. This assumption is made in order simplify the form of the ensuing ersätz estimation formulas, as shown in [8], correction factors may sometimes be calculated for other models.

Assumption 2 *For each line L_l, during day segment s, $\Delta_j(p_1) = \mu_{l,s}\, \delta_j(p_1)$ for p_1 the first platform, and for all other platforms $p \in L_l$, $\Gamma_j(p) = \mu_{l,s}\, \delta_j(p)$, where $\{\delta_j(p), j = 1, 2, \ldots\}$ is a sequence of zero mean i.i.d. random variables, also independent amongst platforms.*

The average *operating cost* per day is $K(\bar{\mu})$: it costs κ_ℓ dollars per trip of any train on line L_ℓ, that is:

$$K(\bar{\mu}) = \sum_{s=1}^{S} \kappa_\ell \, \frac{T_s}{\mu_{\ell,s}}$$

where T_s is the length of segment s (in units of fractions of a day) and $\mu_{\ell,s}$ is the corresponding value of the control parameter. The cumulative wait of passengers is the total time that passengers have to wait in the platforms and it represents a *social cost*, penalizing idle time. Passengers wait at any platform $p \in L_\ell$ whenever they arrive (either from the outside or from transfer lines) between two consecutive train departures at p. Call $W_j(p)$ the total wait of passengers on platform p that take the j-th train of line ℓ. If $M_p(\cdot)$ denotes the train departure counting process at platform p, then:

$$F(\bar{\mu}) = K(\bar{\mu}) + \sum_{s=1}^{S} \sum_{l=1}^{L} \sum_{p \in L_\ell} \mathsf{E} \sum_{j=M_p(\tau_s)+1}^{M_p(\tau_s+T_s)} W_j(p) \qquad (2)$$

Notice that $M_p(\cdot)$ are dependent processes: if $p_i, p_k \in L_\ell$ and $p_i \text{->} \cdots \text{->} p_k$, then the j-th train departs platform p_{k+1} at time $V_j(p_k) = V_j(p_i) + T(p_i, p_k) + \gamma_j(p_i, p_k)$, where $T(p_i, p_k)$ is the total average travel and waiting time of trains between the platforms and $\gamma_j(p_i, p_k)$ is the sum of the independent travel time perturbations of train j at the intermediate tracks from p_i up to p_k. However, if $p_i \in L_\ell$ and $p_k \notin L_\ell$, then the correspoding departure processes are assumed to be independent, under no synchronization.

The passenger arrival processes at every platform is compounded by a Poisson process of passengers boarding from outside, plus all the passengers arriving from other transfer lines. The (batch) arrival processes of transfer passengers from line $L_m \neq L_\ell$ at the station where platform $p \in L_\ell$ is located depends on the departure process at the platform prior to the transfer platform on line L_m, as well as possible previous transfers. This high degree of interdependency makes it difficult to estimate the cost function as well as its derivatives with respect to the control parameters $\bar{\mu}$.

3 Ersätz Modeling

This section develops a local model that approximates the global performance meaure (2). From the platform point of view, the interaction with individual processes from other stations in the network can be simplified using the theory of Poisson processes. This allows a decoupling of the whole network where each platform is modeled via incoming passenger processes with known distributions.

3.1 The Ghost Subway Network

Consider any platform p and suppose first that the day segment s is streched out to eternity, that is, the Poisson arrival processes will be assumed stationary with constant rates, and the frequencies are assumed time invariant. In this section we simplify the model by "localizing" the computation of waiting times at each platform. Dropping the explicit dependency in p from the notation of $V_j(p)$, the total waiting time for passengers boarding train j at this platform can be written as:

$$\begin{aligned} W_j(p) &= \sum_{n=N_0(V_{j-1})+1}^{N_0(V_j)} (V_j - S_n^0) + \\ & \sum_{m=1}^{\mathcal{M}} \sum_{k=N_m(V_{j-1})+1}^{N_m(V_j)} P_k^m \, (V_j - S_k^m) \qquad (3) \end{aligned}$$

The first sum is the waiting time of outside passengers, and $N_0(\cdot)$ denotes the aggregate Poisson process of all arrival streams of origin p, with arrival epochs $\{S_n^0\}$ while the second term is the waiting time of transfer passengers, assuming that there are $\mathcal{M}$ different lines that can transfer at this station towards platform p, and we denote by $N_m(\cdot)$ the arrival process of trains from line L_m at the station where platform p is located. Their corresponding arrival epochs are $\{S_k^m\}$. Finally, the quantity P_k^m is the number of passengers in transfer at the arriving train k from line L_m. Whenever convenient, use $P_k^m(o, d)$ for the number of such passengers with origin o and destination d.

Let platforms $p \in L_\ell$ and $p_m \in L_m$ be at the same station, that is, lines ℓ and m meet at a transfer station where the corresponding platforms are p and p_m, respectively. Consider a transfer line L_m, and let $p_m \in L_m$ be the platform at the transfer station where p belongs: trains departing platform p_m at times $\{V_k(p_m)\}$ correspond to the transfer arrival process $N_m(\cdot)$ at platform $p \in L_\ell$ and passengers will transfer from p_m to p. Let $\mathcal{D}_o(p_m, p)$ denote the set of destinations that must carry out the transfer from $p_m \in L_m$ to $p \in L_\ell$. Call "first order transfer" passengers those transfer passengers coming from the k-th train on line L_m who transfer for the first time at platform $p \in L_\ell$, illustrated in Figure 1. That is, $\sum_{o \in L_m} P_k^m(o, d) \mathbf{1}_{\{d \in \mathcal{D}_o(p_m, p)\}}$, because these transfer passengers necessarily boarded from the outside at some platform $o \in L_m$.

Proposition 1 *Let p be a platform of line L_ℓ where transfer trains arrive from lines $L_m, m = 1, \ldots, \mathcal{M}$ at times S_k^m, with $T_k^m = S_k^m - S_{k-1}^m$ the interarrival times. Call platform $p_m \in L_m$ the transfer platform belonging to the same station as p. Then the number of transfer passengers P_k^m arriving at S_k^m satisfies $\mathsf{E}[P_k^m | T_k^m] = \rho_m T_k^m$, where:*

$$\rho_m = \sum_{n=1}^{L} \sum_{o \in L_n} \lambda(o, d) \mathbf{1}_{\{d \in \mathcal{D}_o(p_m, p)\}}.$$

Proof: Consider first order transfer passengers first. As shown in Figure 1, the number of passengers that arrive on train k is the sum of the passengers from every possible origin $o \in L_m$ with a destination requiring transfer from p_m to p, that is: $d \in \mathcal{D}_o(p_m, p)$. Working backwards, the time of departure of the k-th train at the origin platform $o \in L_m$ is of the form:

$$\begin{aligned} V_k(o) &= V_k(p_m) - T(o,p_m) - \mu_m\,\delta_k(o,p_m) \\ &= S_k^m - T(o,p_m) - \mu_m\,\delta_k(o,p_m). \end{aligned}$$

Because the number of passengers that board the train at this origin $o \in L_m$ with destination $d \in \mathcal{D}_o(p_m, p)$ is the number of Poisson arrivals with rate $\lambda_{o,d}$ that arrive within $[V_{k-1}(o), V_k(o))$, it follows that this is a Poisson random number whose conditional mean is given by: $\mathsf{E}[\lambda_{o,d}(T_k + \mu_m(\delta_k(o,p_m) - \delta_{k-1}(o,p_m)))|T_k] = \lambda_{o,d}T_k$. This follows because of the assumed independence between passenger arrivals and train departures: the Poisson process with rate $\lambda_{o,d}$ at the origin node on platform L_m is statistically independent of the train departure times and the variables $\{\delta_k(o,p_m)\}$, which are i.i.d.

Next, the total number of first order transfer passegers that board train j from p_m is $\bar{P}_j^m(o,d) = \sum_{k_1+1}^{k_2} P_k^m(o,d)$ is the total number of Poisson arrivals at o between trains $k_1 \equiv N_m(V_{j-1})$ and $k_2 \equiv N_m(V_j)$, that is: $\bar{P}_j^m(o,d) \overset{\mathcal{L}}{=} \eta_{o,d}(V'_j) - \eta_{o,d}(V'_{j-1})$ where $\eta_{o,d}(\cdot) \sim$ Poisson$(\lambda_{o,d})$ and the interval $[V'_{j-1}, V'_j)$ corresponds to the shifted time interval at o between trains k_1 and k_2:

$$\begin{aligned} V'_{j-1} &= V_{k_1}(p_m) - T(o,p_m) - \mu_m\delta_{k_1}(o,p_m), \\ V'_j &= V_{k_2}(p_m) - T(o,p_m) - \mu_m\delta_{k_2}(o,p_m). \end{aligned}$$

Letting $Y'_j = V'_j - V'_{j-1}$, it follows from stationarity of Poisson processes that $\bar{P}_j^m(o,d) \overset{\mathcal{L}}{=} \eta_{o,d}(Y'_j + \mu_m[\delta_{k_2}(o,p_m) - \delta_{k_1}(o,p_m)])$. Finally, because the departure process at o is renewal, $\mathsf{E}[Y'_j|Y_j(p)] = Y_j(p)$ for $Y_j(p) = V_j(p) - V_{j-1}(p)$ and therefore the conditional expectation of $\bar{P}_j^m$ given $Y_j(p)$ is $\lambda_{o,d}\,Y_j(p)$, using independence of the delay functions at line L_m and departure times at platform $p \in L_\ell$ (no synchronization), and the fact that $\delta_k(o,p_m)$ are iid.

The foregoing allows us to establish by induction that passengers from second and higher order transfers arriving to p in train k at epoch S_k^m from $p_m \in L_m$ but which originated at some other line L_n also have a mean conditional expectation proportional to T_k, which establishes the result. ■

Using the above proposition, the objective function of interest in (2) can be expressed in terms of the conditional expectations $\mathsf{E}[P_k^m|T_k^m]$, which in turn are the same as the conditional expectations of *surrogate* passenger numbers that have a Poisson distribution with intensity $\rho_m\,T_k^m$. This is what we call the "ghost" network, because the approximation uses an equivalent model where only arrivals and departures of trains at the platforms are important. Table 1 shows the estimation of waiting times when passengers are generated (real network) and when the expected value of such passengers is used instead (ghost network). Estimations were done for a subway network with 36 platforms and 8 lines. Results are shown only for transfer plarforms.

	Total Waiting Time				
p	Real	σ	Ghost	σ	% Error
2	2940.7	289.5	3023.6	292.4	2.82
5	2890.8	297.7	2985.9	301.2	3.29
7	48918.4	1040.2	49128.8	899.2	0.43
11	7431.5	297.7	7540	249.8	1.46
14	42692.9	1536.9	43342.2	1575.4	1.52
19	14067.8	733.3	14198.7	719.5	0.93
24	33457.8	1909.1	33865.2	1948.2	1.22
29	28177.6	1066.5	28750.9	1113.2	2.03

Table 1: Comparison between real and ghost systems.

3.2 Local Performance Model

Instead of looking at the interdependency of the train processes, the following results establish that the distribution of the processes at each platform can be approximately described locally. The ensuing local models permit a probabilistic representation that "ignores" the outside world and decouples the time-dependencies between lines. Let platform $p \in L_l$ be fixed, as well as the day segment s. Drop the labelling of p, s and call μ the variable $\mu_{l,s}$. At the station where platform p is located, there are $\mathcal{M}$ other platforms from different lines ($\mathcal{M}$ can be zero). The local platform model is built in terms of the following processes, all defined in a common probability space $(\Omega, \mathbb{P})$.

Proposition 2 *The times between consecutive train departures Y_j at platform $p \in L_l$ follow a distribution G_μ where $\mathsf{E}[Y_j] = \mu$ and μ is a scale parameter of G_μ.*

Proof: Let $p_1 \in L_l$ be the initial platform. Then, $Y_j = V_{j+1} - V_j$. By Assumption 2, we can write $Y_j = \mu(1 + \delta_j(p_1) + \delta_j(p_1, p) - \delta_{j+1}(p_1, p)$, which implies that $Y_j = \mu X_j$, where X_j is independent of μ with $\mathsf{E}[X_j] = 1$, which implies the result. ■

Train departure process $M(t)$**.** Let $\{Y_j\}$ be i.i.d. random variables on $(\Omega, \mathbb{P})$ following a distribution G_μ, for which μ is a scale parameter and such that $\mathsf{P}(Y_j = 0) = 0$. Call $V_j = V_{j-1} + Y_{j-1}$ the epochs of train departures at the platform. Assume that time $t = 0$ is the start of the current day segment which is defined to last up to $t = T$. $V_0 \leq 0$ corresponds to the last train departure time from the previous day segment (if any).

Transfer trains arrival processes $N_m(t)$. For $m = 1, \ldots, \mathcal{M}$, $\{T_k^m\}$ are i.i.d. random variables on $(\Omega, \mathbb{P})$, independent of the process $M(\cdot)$, with common distribution F_{μ_m}, for which μ_m is a scale parameter and such that $\mathsf{P}(T_k^m = 0) = 0$. Call $S_k^m = S_{k-1}^m + T_k^m$ the arrival epochs of the trains in transfer process $N_m(\cdot)$. The time $S_0^m \leq 0$ corresponds to the last arrival in the previous day segment.

The performance $F(\bar{\mu})$ in (2) is an expectation. Use conditioning to express the performance of the platform (3) in terms of $\rho_m\, T_k^m = \mathsf{E}[P_k^m \,|\, T_k^m]$ and use instead of $W_j(p)$, the surrogate waiting functional:

$$\tilde{\phi}(\bar{\mu},\omega) = \sum_{j=1}^{M(T)} \sum_{m=0}^{\mathcal{M}} \rho_m \sum_{k=N_m(V_{j-1})+1}^{N_m(V_j)} T_k^m\,(V_j - S_k^m).$$

For the first term $m=0$, $\rho_0 = \sum_d \lambda_{p,d}$ is the arrival rate of passengers originating at p and the corresponding expected waiting times of passengers can be calculated theoretically (as shown in [6]) so that instead of the previous pathwise functional, we will consider the partially integrated random functional:

$$\phi(\bar{\mu},\omega) = \frac{\rho_0 T}{2\mu}\mathsf{E}[Y^2] + \sum_{j=1}^{M(T)} \sum_{m=1}^{\mathcal{M}} \rho_m \sum_{k=N_m(V_{j-1})+1}^{N_m(V_j)} T_k^m\,(V_j - S_k^m)$$

It follows from Proposition 1 that $\mathsf{E}[\phi(\bar{\mu})]$ is indeed the expected waiting time at platform p over the time segment s, except for the *end effects* of passengers that boarded at some origin o during a day segment $s' < s$ and transferred at platform p during the current day segment. We state this result in Theorem 1 below.

Theorem 1 *Under stationarity conditions, i.e., if the Poisson arrival processes are stationary with constant rates and the frequencies are time invariant between day segments, then* $\mathsf{E}[\phi(\bar{\mu})] = \mathsf{E}\sum_{j=M(\tau_{s-1})+1}^{M(\tau_s)} W_j(p)$.

The time scales that are realistic for metropolitan transport systems suggest that the end effects will only affect few passengers. However, the net effect may yield a non-negligible bias and we are currently developing correction terms to estimate this bias.

4 Ersätz Derivatives

4.1 The Local Gradients

At platform $p \in L_l$ on day segment s, let $\phi_m(\bar{\mu})$ be the cumulative ersätz waiting time of transfer passengers arriving from transfer stream m within this day segment, for each $m = 1,\ldots,\mathcal{M}$, that is:

$$\phi_m(\bar{\mu},\omega) = \sum_{j=1}^{M(T_s)} \rho_m \sum_{N_m(V_{j-1})+1}^{N_m(V_j)} T_k^m (V_j - S_k^m). \qquad (4)$$

At the start of the segment $t=0$, S_0^m is the epoch of the immediately preceding train arrivals in the previous day segment from line L_m, and $V_0 < 0$ is the epoch of the previous train departure for the current platform. These quantities are assumed known at $t=0$. Denote by E_0 the conditional expectation given $\{S_0^1,\ldots,S_0^{\mathcal{M}},V_0\}$. From the renewal assumptions of the model, this conditioning has the same effect as conditioning on the whole past history of the process. A local sensitivity estimator of the ersätz performance function $\mathsf{E}_0[\phi_m(\bar{\mu})]$ is a function $\varphi_{p,m}(\bar{\mu})$ of the underlying processes of the local model (namely train arrivals at incoming platforms p_m and train departures from p) such that:

$$\mathsf{E}_0[\varphi_{p,m}(\bar{\mu})] = \frac{\partial \mathsf{E}_0[\phi_m(\bar{\mu})]}{\partial \mu_{l,s}} \qquad (5)$$

Building such derivative estimators can be done in a number of ways. One of the authors has developed three estimators using the Likelihood Ratio method, the Smoothed Perturbation Analysis method, and the Weak Derivative method, as explained in [3, 9]. The details of derivative estimation for the local performance will be reported elsewhere. Next we will express all the desired derivatives in terms of the local derivative w.r.t. μ, which is the basis for ersätz estimation.

Theorem 2 *If* $\mathsf{E}_0[\phi_m(\bar{\mu})], 1,\ldots,\mathcal{M}$ *are continuously differentiable in* $\mu,\mu_1,\ldots,\mu_{\mathcal{M}}$*, then for* $m=1,\ldots,\mathcal{M}$*:*

$$(S_0^m + \mu)\frac{\partial \mathsf{E}_0[\phi_m(\bar{\mu})]}{\partial \mu} = -\mu_m \frac{\partial \mathsf{E}_0[\phi_m(\bar{\mu})]}{\partial \mu_m} + 2\,\mathsf{E}_0[\phi_m(\bar{\mu})].$$

The proof of this result follows from Theorem 1 in [8] and the scalability assumptions in the model of Section 3.1. One of the practical problems for the implementation of gradient-search methods for network optimization is the number of derivatives required. Theorem 2 reduces the number of required estimates: only $\mathcal{M}$ estimators $\varphi_{p,m}$ are required per platform, which implies that the computational effort is independent of the degree of connectivity of the transportation network.

4.2 Global Gradients via Ersätz Estimation

To obtain the desired gradient, the local performance model is used by noticing that the global performance is additive with respect to the platforms. More specifically, from (2), and the results on Section 3:

$$\frac{\partial F(\bar{\mu})}{\partial \mu_{l,\sigma}} = \frac{\partial K(\bar{\mu})}{\partial \mu_{l,\sigma}} - \sum_{p\in L_\ell} \frac{\rho_p(\sigma) T_\sigma\, S^2(\sigma,p)}{2}$$
$$+ \sum_{s=\sigma}^{S} \sum_{p\in L_\ell} \left\{ \sum_{m\in\mathcal{I}(p)} \frac{\partial \mathsf{E}[\phi_{p,m}^s(\bar{\mu})]}{\partial \mu_{l,\sigma}} + \sum_{n\in\mathcal{O}(p)} \frac{\partial \mathsf{E}[\phi_{n,p}^s(\bar{\mu})]}{\partial \mu_{l,\sigma}} \right\}$$

where $S^2(\sigma,p) = \mathsf{E}[Y_j^2(\sigma,p)]/\mu_l^2$ is independent of μ_l and it can be calculated at each platform, $\mathcal{I}(p)$ is the set of all platforms m in the same station as p with passengers transferring to p, and $\mathcal{O}(p)$ is the set of platforms in the same station to which passengers from p may transfer. The rate $\rho_p(\sigma) = \sum_d \lambda_{p,d}(\sigma)$ is the incoming passenger rate from outside. Finally, $\phi_{m,p}^s$ is the ersätz waiting time (4) at platform p of all passengers arriving within the day segment s from the m-th incoming platform, and $\phi^s_{p,n}$ is the waiting time at platform n-th of passengers that left platform p towards n. In expressing the derivative, we used the fact that the average wait of passengers from outside is not affected by any other headway but the one of the line where and when they originate, thus only the term for day segment σ and line l are affected by changes in $\mu_{l,\sigma}$. As well, using the ersätz performance models at each platform, $\mu_{l,\sigma}$ can only affect the surrogate wait at platforms p along line L_l, as well as the wait at platforms $n \in \mathcal{I}(p)$. Furthermore, changes in $\mu_{l,\sigma}$ will not affect *previous* day segments, only future ones, and this only through the initial condition. Finally, for platforms $p \in L_l$ which are not inside transfer stations, both $\mathcal{I}(p)$ and $\mathcal{O}(p)$ are empty. Denote by $\varphi^s_{p,m}$ the local ersätz sensitivity estimator satisfying (5) for each segment s. From the structure of the local models, it follows that $\mathsf{E}[\phi^s_{p,m}(\bar{\mu}) | (V_0(p), S^m) = x]$ is Lipschitz continuous in $\mu_{l,s}$ for every x, which implies:

$$\begin{aligned}
\mathsf{E}[\varphi^s_{p,m}] &= \int \frac{\partial}{\partial \mu_{l,s}} \mathsf{E}[\phi^s_{p,m}(\bar{\mu}) | (V_0(p), S^m) = x] \, \mathsf{P}(dx) \\
&= \frac{\partial}{\partial \mu_{l,s}} \int \mathsf{E}[\phi^s_{p,m}(\bar{\mu}) | (V_0(p), S^m) = x] \, \mathsf{P}(dx) \\
&= \frac{\partial \mathsf{E}[\phi^s_{p,m}(\bar{\mu})]}{\partial \mu_{l,s}},
\end{aligned}$$

where $\mathsf{P}(dx)$ is the distribution of the initial condition of the day segment s and is independent of $\mu_{l,s}$. Using Theorem 2, we propose the ersätz estimation:

$$\begin{aligned}
\frac{\partial F(\bar{\mu})}{\partial \mu_{l,\sigma}} \approx & -\left(\kappa_l \frac{T_\sigma}{\mu^2_{l,\sigma}} + \sum_{p \in L_\ell} \frac{\rho_p(\sigma) T_\sigma \, S^2(\sigma, p)}{2} \right) \\
& + \sum_{p \in L_\ell} \left\{ \sum_{m \in \mathcal{I}(p)} \mathsf{E}[\varphi^\sigma_{p,m}] \right. \\
& \left. + \sum_{n \in \mathcal{O}(p)} \mathsf{E}\left[\frac{2\phi^\sigma_{n,p}(\bar{\mu}) - (S_0^{n,\sigma} - \mu^\sigma_{l(n)}) \varphi^\sigma_{n,p}}{\mu_{l,\sigma}} \right] \right\} \quad (6)
\end{aligned}$$

There are two sources of bias in (6). The first one is due to neglecting end effects in Theorem 1, and the second, to neglecting the propagation effects of the control variable $\mu_{s,l}$ on a given platform $p \in L_l$ for *future* day segments: $\mu_{s,l}$ can affect the performance of day segment $s + 1$ only via the initial conditions. But we are estimating the derivatives conditional on the initial values. We conjecture that the main contribution to the global cost is indeed over the current day segment.

5 The Learning Algorithm

For each segment s, after $D_p(s)$ days, each platform p constructs the sample averages of the local derivative estimators $\varphi^s_{p,m}, m \in \mathcal{I}(p)$. The estimators have been built using the ghost local model and therefore only basic counters are required at each platform, capable of recording train departure times. These averages are broadcast first, to the initial platform on the line L_l where p belongs. Next, each $\varphi^s_{p,m}$ is also broadcasted towards the initial platform along line $L_{l(m)}$ where patform m belongs. At the initial platform of each line L_l, the information received for each segment from all platforms in the network is added to the estimation of $\partial F / \partial \mu_{s,l}$ using (6). Periodically, each line updates its own frequency $\mu_{s,l}$ using (1) for that component.

Under stationary conditions and fixed values of $\bar{\mu}$, all estimators are i.i.d. from day to day. Furthermore, by assumption they are bounded (using the fact that day lengths are bounded). Using the methods in [5], the recursion (1) can be shown to yield sub-optimal policies, due to the sources of bias mentioned before. To track slow changes in traffic patterns, it is necessary to estimate as well the aggregated traffic intensities $\lambda_{o,d}(s)$ of the network. We are currently implementing the scheme to a realistic model with up to third order transfers to assess this sub-optimality numerically.

References

[1] Bookbinder, J.H. and A. Désilets. “Transfer Optimization in a transit network”, *Transp. Sci.* 26, pp. 106-118, 1992.

[2] Fu, M. and S. Hill. “Optimization of Discrete Event Systems via Simultaneous Perturbation Stochastic Approximation,” *IIE Trans.*, Vol. 29, No.3, pp. 233-243, 1997.

[3] Glasserman, P. *Gradient Estimation via Perturbation Analysis.* Kluwer Academic, Boston, 1991.

[4] Heidergott, B. and R. de Vries. “Towards a Control Theory for Transportation Networks”, submitted to *Journal of DEDS*, March 2000.

[5] Kushner, H.J. and G. Yin. *Stochastic Approximation and Applications.* Springer Verlag, New York, 1997.

[6] Ross, S.M. *Introduction to Probability Models, 5th edition.* Academic Press, San Diego, 1993.

[7] Sansó, B. and P. Girard. “Instantaneous Power Peak Reduction and Train Desynchronization in Subway Systems”, *Transp. Sci.*, 31-4, pp. 312-323, 1997.

[8] Vázquez-Abad, F. and L. Zubieta. “Generalizations of the surrogate estimation approach for sensitivity analysis via simple examples”, *Proc. of the 38-th IEEE CDC*, Edited by IEEE Control Systems Society, pp. 1796-1803, 1999.

[9] F.J. Vázquez-Abad. “A Course on Sensitivity Analysis for Gradient Estimation of DES Performance Measures”, Part I of III, in *Discrete Event Systems, Analysis and Control*, R. Boel and G. Stremersch Editors, Kluwer Academic Publishers, Boston, 2000.