

Measurement Based Optimal Routing on Overlay Architectures for Unicast Sessions

Tuna Güven[†], Richard J. La^{†*},
Mark A. Shayman[†], and Bobby Bhattacharjee[‡]

[†]*Department of Electrical and Computer Engineering,
University of Maryland, College Park*

[‡]*Department of Computer Science,
University of Maryland, College Park*

Abstract

We propose a measurement-based routing algorithm to load balance intradomain traffic along multiple paths for multiple unicast sources.² Multiple paths are established using overlay nodes. The distributed algorithm is derived from simultaneous perturbation stochastic approximation (SPSA) and does not assume that the gradient of an analytical cost function is known. Instead, we rely on (potentially) noisy estimates from local measurements. Using the analytical model presented in the paper, we first show the almost sure convergence of the algorithm to the optimal solution under a decreasing step size policy (as with a standard SPSA model). Motivated by practical concerns, we next consider the constant step size case, for which we establish weak convergence. We provide simulation results to demonstrate the advantages of our proposed algorithm under various network scenarios, and also present a comparative study with MATE (an existing optimal routing algorithm).

Key words: Mathematical Programming/Optimization, Load Balancing, Simulations.

² A shorter version of this work appeared in the Proceedings of the IEEE Conference on Computer Communications (Infocom) 2004.

¹ * - corresponding author: hyongla@isr.umd.edu, A. V. Williams Building, Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742. Phone: (301) 405-4914. Fax: (301) 314-9281

1 Introduction

Rapid growth of the Internet and the emergence of new demanding services have sparked interest in Internet traffic engineering. As defined in [1], traffic engineering deals with the issue of performance evaluation and performance optimization of operational IP networks and encompasses the *measurement, characterization, modeling* and *control* of the Internet traffic.

Due to the evolution of the Internet from ARPANET, traditional routing algorithms for IP networks are mostly based on shortest path routing. However, methods relying on a single path between a source-destination pair often do not utilize network resources efficiently and offer limited flexibility for traffic engineering [1]. Various solutions derived from shortest path routing algorithms have been suggested, mainly by modifying link metrics in accordance with the network dynamics (*e.g.*, [2], [3]). However, these approaches have several shortcomings that have not been addressed satisfactorily. First, they lead to network-wide effects and can result in undesirable and unanticipated traffic shifts [1]. Second, these schemes cannot distribute the load among the paths with different costs.

This paper focuses on the issue of *traffic mapping* (load balancing). More specifically, given some source-destination traffic matrix, we are interested in finding a traffic assignment onto pre-established paths so that the overall system performance is optimized with respect to a selected performance measure (*e.g.*, sum of link costs). In this paper we adopt an overlay architecture proposed in [5] in order to establish multiple paths between a source-destination pair. However, our basic approach can be directly applied to other types of networks, such as MPLS based networks.

We propose a distributed optimal routing algorithm based on stochastic approximation, using local network state information. Our model is similar to that of [4], with the following differences. In [4] the authors mention that the *cost derivatives* cannot be computed and hence, should be estimated by measurements. However, the mathematical analysis given in [4] implicitly assumes that the analytical gradient function is available. In addition, the details on how a cost gradient is estimated are not provided, and the method described in [6] appears to be a variant of a well-known *finite differences* method ([7], [8]). However, the issue of gradient estimate is not clearly or explicitly stated in the aforementioned references. We believe that the gradient estimation method plays a crucial role in the sense that the *convergence* of the optimal routing algorithm strongly depends on the conditions defining this estimation process as described in the stochastic approximation literature (See [8–10]).

In this study we consider the same problem (of mapping traffic onto multiple

end-to-end paths) while relaxing the assumption that the analytical gradient function is available. We derive our proposed *measurement based algorithm* from the idea of simultaneous perturbation stochastic approximation (SPSA). SPSA allows us to greatly reduce the number of measurements required for estimating the gradient, while at the same time it achieves approximately the same level of accuracy as the classical finite differences method at each iteration. By reducing the number of measurements, our algorithm achieves faster convergence. This is because each measurement requires a non-negligible amount of time in a networked environment. We will detail these issues of estimation, measurement, and convergence in the rest of the paper, and we show in Section 5, using simulations, that our SPSA based algorithm outperforms the algorithm proposed in [4].

From a broader perspective, a special case of the proposed algorithm (i.e., single source-destination (SD) case) provides an optimal solution to more general problems that have a *simplex constraint set*. Although applications of SPSA to the constrained optimization problems have generated a certain level of interest in the literature, the simplex constraint set problems have not been handled properly as we will discuss in the following section.

Note that the SPSA algorithm can achieve almost sure convergence, provided that the step size parameter diminishes with the number of iterations. However, such a policy may not be practical under dynamic network conditions as the algorithm may not be able to react to network changes in a timely manner after the step size parameter becomes small. As a result, in practice, we have to reset the step size value after some time interval to ensure that the algorithm is able to react to network dynamics appropriately. An obvious alternative is the use of a constant step size. Even though it is difficult to obtain similar almost sure convergence in a constant step-size case, it has been shown in [11] that weak convergence (i.e., convergence in distribution) is possible under certain conditions. While the weak convergence is not as strong as almost sure convergence, we demonstrate in Section 5 that it does not have a significant effect on the performance of the algorithm in a practical sense.

The rest of the paper is organized as follows. In Section 2 we define the optimization problem, and present a brief overview of stochastic approximation. In Section 3, we present an optimal routing algorithm, and prove its stability and optimality. We discuss implementation issues in Section 4. Section 5 describes the experimental setup and presents a set of simulation results. We conclude the paper and discuss possible topics of future work in Section 6.

2 The Optimization Problem

2.1 The Routing Model

In this section, we define the optimization problem of interest, describe the network model used for the analysis, and list basic assumptions we make. We will closely follow the formulation in [4] due to the similarity of the problem.

The network is modeled by a set L of unidirectional links. Let $S = \{1, 2, \dots, S\}$ denote the set of SD pairs. An SD pair s has a set $P_s \subseteq 2^L$ of paths available to it, and $N_s = |P_s|$, i.e., N_s is the number of paths available for SD pair s . Denote the set of links that belong to at least one path $p \in P_s$ by $L^s \subset L$. With a little abuse of notation we let $P_s = \{1, 2, \dots, N_s\}$, and define the set of all paths $P = \cup_{s \in S} P_s = \{1, 2, \dots, N\}$, where $N = \sum_{s \in S} N_s$. While by definition, none of the paths can be used by more than one SD pair, any two paths can share a link.

The total input traffic rate of an SD pair s is denoted by r_s , and the SD pair routes x_{sp} amount of traffic on path $p \in P_s$ such that

$$\sum_{p \in P_s} x_{sp} = r_s, \quad \text{for all } s. \quad (1)$$

Let $\mathbf{x}_s = (x_{sp}, p \in P_s)$ be the rate vector of SD pair s , and let $\mathbf{x} = (\mathbf{x}_s, s \in S)$ be the vector of all rates. The flow rate on a link $l \in L$ is given by

$$x^l = \sum_{s \in S} \sum_{p \in P_s : l \in p} x_{sp} \quad (2)$$

For each link l , $C_l(x^l)$ represents the link cost as a function of the link flow rate x^l . We assume that, for all l , $C_l(\cdot)$ is convex and continuously differentiable. The objective is to minimize the total cost $C(\mathbf{x}) = \sum_l C_l(x^l)$ by mapping the traffic on paths in P :

$$\min_{\mathbf{x}} C(\mathbf{x}) = \min_{\mathbf{x}} \sum_l C_l(x^l) \quad (3)$$

$$\text{s. t. } \sum_{p \in P_s} x_{sp} = r_s, \forall s \in S \quad (4)$$

$$x_{sp} \geq \epsilon, \forall p \in P_s, s \in S, \quad (5)$$

where ϵ is an arbitrarily small positive constant. For instance, some of the control packets may be routed along different paths available to an SD pair.

Theoretically if the exact gradient values are known, one can use the well known gradient projection algorithm to solve this constrained optimization problem, where the constraint set Θ is defined by (4) and (5). Each iteration of the algorithm takes the form:

$$\mathbf{x}(k+1) = \Pi_{\Theta}[\mathbf{x}(k) - a(k)\nabla C(k)] \quad (6)$$

where $\nabla C(k)$ is the gradient vector whose $(s, p)^{th}$ element is the first derivative length of path $p \in P_s$ at the k -th iteration ($[\nabla C(k)]_{sp} = \partial C(\mathbf{x}(k))/\partial x_{sp}$), $a(k) > 0$ is the step size, and $\Pi_{\Theta}[\mathbf{v}]$ denotes the projection of a vector \mathbf{v} onto the feasible set Θ with respect to the Euclidean norm.

The above iteration can be carried out in a distributed manner by each pair s without the need to coordinate with other pairs [12], [13]. In other words, the source of each SD pair s updates its rates \mathbf{x}_s independently of other SD pairs:

$$\mathbf{x}_s(k+1) = \Pi_{\Theta_s}[\mathbf{x}_s(k) - a_s(k)\nabla C_s(k)] \quad (7)$$

where $\nabla C_s(k) = (\partial C(\mathbf{x}(k))/\partial x_{sp}, p \in P_s)$ is the vector of first derivative lengths of paths in P_s , i.e., a subvector of $\nabla C(k)$ with elements corresponding to the paths in P_s , and Π_{Θ_s} denotes a projection onto the feasible set Θ_s of SD pair s .

One problem with directly implementing (7) is that the first derivative length of a path, $\partial C/\partial x_{sp}$, may not be available in practice and can only be estimated empirically through *noisy* measurements of the cost function. This is due to the fact that the packet arrival processes are both stochastic and time varying. Therefore, one must resort a gradient approximation method to obtain an estimate to be used in (7). Stochastic approximation methods are natural candidates for such problems.

2.2 Stochastic Approximation

Stochastic Approximation (SA) is a recursive procedure for finding the root(s) of equations using noisy measurements, and is particularly useful for finding extrema of functions [9] (*e.g.*, [14] and [15]).

General constrained SA has the same form as (6) with the gradient vector $\nabla C(k)$ replaced by its approximation $\hat{\mathbf{g}}(k)$. The approximation is obtained through measurements of $C(\mathbf{x})$ around $\mathbf{x}(k)$. Under appropriate conditions, one can show that $\mathbf{x}(k)$ converges to the solution set of (3), which we denote by \mathbf{x}^* .

A critical issue in SA is the approximation of gradient vector. The standard approach motivated from the definition of gradient is the *Finite Differences* (FD) method, in which each component of $\mathbf{x}(k)$ is perturbed one at a time and corresponding measurements $y(\cdot)$ are obtained. Typically, the i -th component of $\hat{\mathbf{g}}(k)$ ($i = 1, 2, \dots, m$) under the FD method is given by

$$\hat{g}_i(k) = \frac{y(\mathbf{x}(k) + c(k)\mathbf{e}_i) - y(\mathbf{x}(k) - c(k)\mathbf{e}_i)}{2c(k)}$$

where $c(k)$ is some positive number, \mathbf{e}_i denotes a unit vector with one in the i -th position and zeros elsewhere, and $y(\cdot)$ denotes the measured cost function with measurement noise.

An alternative method for estimating the gradient is called the *Simultaneous Perturbation* (SP). In this method, all elements of $\mathbf{x}(k)$ are randomly perturbed together to obtain two measurements. The i -th component of $\hat{\mathbf{g}}(k)$ is computed by

$$\hat{g}_i(k) = \frac{y(\mathbf{x}(k) + c(k)\Delta(k)) - y(\mathbf{x}(k) - c(k)\Delta(k))}{2c(k)\Delta_i(k)}$$

where the vector of the random perturbations for SP, $\Delta(k) = (\Delta_1(k), \Delta_2(k), \dots, \Delta_m(k))$, needs to satisfy certain conditions that will be discussed in the following section. Here m denotes the dimension of the vector \mathbf{x} .

Both of the above approximations have a “two-sided” form in the sense that they use the measurements $y(\mathbf{x}(k) \pm \text{perturbation})$. On the other hand, one-sided gradient approximations require measurements of $y(\mathbf{x}(k))$ and $y(\mathbf{x}(k) + \text{perturbation})$. Although it is known that the standard two-sided form gives more accurate estimates compared to one-sided forms, for real-time applications one-sided gradient approximation may be preferred when the underlying system dynamics change too rapidly to get an accurate gradient estimate with two successive measurements [7]. In this paper we assume that the one-sided form is utilized for the gradient approximation under both methods unless stated otherwise.

An SA algorithm using the FD (resp. SP) gradient approximation method is referred to as a FDSA (resp. SPSA) algorithm. One should note that, in an SPSA algorithm the gradient approximation requires only two cost function measurements, regardless of the value of m . Standard (two-sided) FD approximation requires $2m$ measurements to estimate the gradient. In [9] it is shown that under reasonably general conditions, SPSA and FDSA achieve the same level of statistical accuracy for a given number of iterations even though SPSA uses m times fewer function evaluations than FDSA. This theoretical result has been confirmed in many numerical studies, even in cases where m is on

the order of several hundreds or thousands [7]. This is certainly an important property especially if the measurements are costly and/or time consuming. Clearly, this is the case for the routing problem at hand as measurements require resources and must be collected and reported in a timely manner. In other words, SPSA promises a potential for better statistical accuracy over the same period of “time” due to a much shorter measurement period required at each iteration, even though the two methods have the same statistical accuracy with the same number of “iterations”. This suggests that the algorithm based on SPSA will be able to track and respond to changes in a network much faster than another algorithm based on FDSA and thus improve the overall network performance.

In [9], Spall provides a formal proof of convergence of SPSA algorithm for the “unconstrained” case. The convergence of an SPSA algorithm under inequality constraints are presented in [10,16]. However, these results do not consider the case where $\mathbf{x}(k) \pm c(k)\mathbf{\Delta}(k) \notin \Theta$, which may occur in our routing problem. In [16] Sadegh suggests projecting $\mathbf{x}(k)$ to a point $\mathbf{x}'(k) \in \Theta$ such that $\mathbf{x}'(k) \pm c(k)\mathbf{\Delta}(k) \in \Theta$. If $\mathbf{x}'(k) - \mathbf{x}(k) \rightarrow 0$ as $k \rightarrow \infty$, convergence can still be established. However, when Θ is a simplex, if $c(k) \sum_j \Delta_j(k) \neq 0$ then $\mathbf{x}'(k) \pm c(k)\mathbf{\Delta}(k) \notin \Theta$ for all $\mathbf{x}'(k)$ (as the demand constraint in (5) is violated). Under these conditions, there is no established convergence result of an SPSA algorithm that we can directly apply to our problem. (In [10], although authors claim that they have shown the convergence for the case of a network of queues with similar constraints, they do not consider the aforementioned issue in the proofs.)

In the next section, we will resolve this technical issue by a simple method and present a formal proof of the SPSA algorithm under these constraints.

3 Optimal Routing Using SPSA

3.1 An Optimal Routing Algorithm - Decreasing Step Size

In this section we propose an optimal routing algorithm and prove its stability and optimality. We know from [12] that if each SD pair runs (7) independently and asynchronously,³ the overall algorithm converges. Let us now consider the use of SPSA in place of (7).

³ Here asynchrony refers to the fact there might be a time lag between SD pairs in continuous time operation in practice.

At time k , SD pair s updates its rate according to

$$\mathbf{x}_s(k+1) = \Pi_{\Theta_s}[\mathbf{x}_s(k) - a_s(k)\hat{\mathbf{g}}_s(k)] \quad (8)$$

where $\hat{\mathbf{g}}_s(k)$ is the approximation to $\nabla C_s(k)$ obtained by the SPSA algorithm and is given by

$$\begin{aligned} \hat{g}_{s,i}(k) &:= \frac{N_s}{N_s - 1} \frac{y_s(\mathbf{x}(k) + \mathbf{C}(k)\Delta(k)) - y_s(\mathbf{x}(k))}{c_s(k)\Delta_{s,i}(k)} \\ &= \frac{N_s}{N_s - 1} \frac{(C_s^+(k) + \mu_s^+(k)) - (C_s^-(k) + \mu_s^-(k))}{c_s(k)\Delta_{s,i}(k)}, \text{ for } i = 1, \dots, N_s, \end{aligned} \quad (9)$$

Here $\Delta(k) = (\Delta_s(k), s \in S)$, $\Delta_s(k)$ is the random perturbation vector for source s at iteration k , $\mathbf{C}(k)$ is an $N \times N$ diagonal matrix whose j -th diagonal entry is equal to c_{s_j} (s_j being the SD pair associated with the j -th component of $\Delta(k)$), and $y_s(\cdot)$ are the noisy measurements of the *partial* cost information, which is the summation of the link costs over the set L^s . Specifically, $C_s^-(k) = \Lambda_s(\mathbf{x}(k)) := \sum_{l \in L^s} C_l(x^l(k))$ and $C_s^+(k) = \Lambda_s(\mathbf{x}(k) + \mathbf{C}(k)\Delta(k))$ with $\mu_s^+(k)$ and $\mu_s^-(k)$ being the measurement noise terms.

Note that (9) differs from the standard SA in the following ways. First, each SD pair uses only *partial* cost information (i.e., summation of the costs of the links in L^s) as opposed to the total network cost, which is the summation of the costs of all the links in the network. In addition, the noise terms observed by each SD pair are allowed to be different. Second, while $c(k)$ is a positive scalar in standard SA, in our case $\mathbf{C}(k)$ is an $N \times N$ diagonal matrix. This allows the possibility of having different $c_s(k)$ values for different SD pairs. Also, note that we have an extra multiplicative factor $\frac{N_s}{N_s - 1}$ in (9) compared to the standard SA. This is due to the projection of $\mathbf{x}_s(k) + c_s(k)\Delta_s(k)$ to Θ_s for all $s \in S$ using L_2 projection while calculating $\hat{\mathbf{g}}_s(k)$. This is explained in Appendix II in details. Finally, if $\Pi_{\Theta_s}[\mathbf{x}_s(k) + c_s(k)\Delta_s(k)] = \mathbf{x}_s(k)$, the SD pair draws a new perturbation vector $\Delta_s(k)$ until $\mathbf{x}_s(k) \neq \Pi_{\Theta_s}[\mathbf{x}_s(k) + c_s(k)\Delta_s(k)]$.

Note from (8) that SD pairs may have different step sizes $a_s(k)$. This allows a certain level of asynchrony between SD pairs in the sense that SD pairs can respond to the network changes independently of other to some extent.⁴ However, we assume that SD pairs update their rates once every iteration after they start the algorithm. This assumption is reasonable in our case, for at each iteration SD pairs should make use of the collected information that is already available. This is, however, not to say that the updates take place

⁴ For instance, this formulation covers the case where SD pairs start the algorithm at different times.

simultaneously. The errors due to this asynchrony are assumed to be absorbed into the error terms $\mu_s^+(k)$ in (9).

For the optimality of the new algorithm, we need to show that (8) converges to the same point \mathbf{x}_s^* as (7) for all SD pairs. For the convergence of the algorithm we assume that the following conditions are true:

- A1. $C(\mathbf{x}(k))$ is differentiable for each $\mathbf{x}(k) \in \Theta$, and convex.
- A2. $\Delta_{s,i}(k)$ are (i) mutually independent with zero mean for all $s \in S$ and $i \in P_s$, (ii) uniformly bounded by some finite constant α , and (iii) independent of $(\mathbf{x}(n), n = 0, 1, \dots, k)$. $\mathbf{E}[(\Delta_{s,i}(k))^{-1}]$ and $\mathbf{E}[(\Delta_{s,i}(k))^{-2}]$ are bounded for all k .
- A3. $\mathbf{E}[\mu_s^{(\pm)^2}(k)]$ are bounded and $\mathbf{E}[\mu_s^+(k) - \mu_s^-(k) | \Delta(k), \mathcal{F}_k] = 0$ a. s. for all k , where \mathcal{F}_k is the σ -field generated by $\{\mathbf{x}(0), \dots, \mathbf{x}(k)\}$.
- A4. (i) $\sum_{k=0}^{\infty} a_s(k) = \infty$, (ii) $a_s(k) \rightarrow 0$ as $k \rightarrow \infty$, (iii) $\sum_{k=0}^{\infty} \frac{a_s^2(k)}{c_s^2(k)} < \infty$, (iv) $c_s(k) \rightarrow 0$ as $k \rightarrow \infty$, and (v) $\frac{c_s(k)}{c_{s'}(k)} = O(1)$ for all $s, s' \in S$.
- A5. There exists a finite positive constant M such that

$$\frac{1}{M} \leq \frac{a_s(k)}{a_{s'}(k)} \leq M \quad (10)$$

for all $s, s' \in S$ and for all k .

- A6. Let $\hat{a}(k) = \max_{s \in S} a_s(k)$. Then, for all $s \in S$

$$\sum_{k=0}^{\infty} (\hat{a}(k) - a_s(k)) < \infty, \text{ and } \lim_{k \rightarrow \infty} \frac{a_s(k)}{\hat{a}(k)} = 1 \text{ for all } s \in S.$$

Proposition 1 *Under Assumptions A1-A6, the sequence $\mathbf{x}(k) = (\mathbf{x}_s(k), s \in S)$ generated by the algorithm defined by (8) converges to \mathbf{x}^* with probability 1, where \mathbf{x}^* is the solution set of the optimization problem in (3), regardless of the initial vector $(\mathbf{x}_s(0), s \in S) \in \Theta$.*

Proof The proof of the Proposition 1 is given in Appendix I.

3.2 The Optimal Routing Algorithm - Constant Step Size

In the previous subsection we have employed decreasing step sizes $a_s(k)$, $k = 0, 1, \dots$, as defined in A4. Although it is possible to use decreasing step sizes, it is of practical importance to consider the case with constant step size $a > 0$ given in the following form

$$\mathbf{x}_s(k+1) = \Pi_{\Theta_s}[\mathbf{x}_s(k) - a \cdot \hat{\mathbf{g}}_s(k)]. \quad (11)$$

This is because in the case of decreasing step sizes, one has to make sure that each source node resets the step size after a certain period of time when the step size becomes too small to effectively react to the dynamics of the network. Consequently, such a requirement introduces additional complexity which can be avoided by using a constant step size. The main difficulty with constant step size algorithms is that it is difficult to establish almost sure convergence. However, as shown in [11], under certain conditions constant step size SA algorithms can achieve weak convergence (i.e., convergence in distribution), which can be interpreted as convergence to a neighborhood of the optimal operating point(s). Since the performance of the system near the optimal operating point(s) may be comparable to that of the optimal solution(s) in a network problem, the performance degradation, if there is any, due to a constant step size may not be significant. This is supported by simulation results in Section 5.

The following proposition establishes the (weak) convergence of the algorithm with a constant step size.

Proposition 2 *Under Assumptions A1-A3, for any $\delta > 0$, the fraction of time the sequence $\mathbf{x}(k) = (\mathbf{x}_s(k), s \in S)$ generated by the algorithm defined by (11) spends in δ -neighborhood of \mathbf{x}^* on $[0, k]$, goes to one (in probability) as $a \rightarrow 0$ and $k \rightarrow \infty$ regardless of the initial vector $(\mathbf{x}_s(0), s \in S) \in \Theta$.*

Proof The proof of the Proposition 2 is given in Appendix III.

3.3 Measurement Process

In this section, we discuss some of the issues regarding the measurement process and their effect on the overall performance of the proposed algorithm. We will also point out the benefits of SPSA based algorithms over the FDSA alternatives.

By definition, an FDSA based algorithm requires SD pair s to perturb its paths *one at a time*, requiring $N_s + 1$ measurements for the one-sided form and $2N_s$ measurements for the two-sided form in each iteration for the estimation of an $N_s \times 1$ gradient vector. For this reason, FDSA based algorithms need each SD pair to collect measurements (i.e., perturb its paths) at different times. As mentioned in [4], this requires a coordination protocol that determines the order in which paths are perturbed. Moreover, it increases network traffic load from the overhead.

SPSA, on the other hand, allows us to estimate the gradient vector using only two measurements. In the context of our routing problem, this implies that an SD pair can perturb all of its paths simultaneously if an SPSA based algo-

rithm is employed. In addition, it also suggests that all SD pairs can execute the perturbation in parallel without the need for strict coordination needed in FDSA (Proposition 1 and Proposition 2). This enables SD pairs to operate independent of each other to some extent. As a consequence, a potential overhead that would be incurred by a coordination protocol under an FDSA algorithm is avoided. Furthermore, we can significantly reduce the duration of measurement periods by concurrently executing the measurement process at the SD pairs. Since the statistical accuracy of SPSA does not degrade from that of FDSA, we can achieve much faster convergence due to the fact that we significantly reduce the time required for each iteration.

Another issue regarding the measurement process is the effect of asynchronous operation of SD pairs in practice due to the lack of perfect synchronization. It is proved in [4] that, with increasing level of asynchrony, the convergence becomes slower. For each $s \in S$, let t_0^s be the time lag between the time traffic measurements are collected and the time SD pair s carries out a rate update. In other words, t_0^s is the delay between the collection of measurements and execution of an SPSA algorithm by SD pair s . Define $t_0 = \max_{s \in S} t_0^s$. Then, the larger the t_0 is, the slower the convergence will be due to the use of delayed information. On the other hand, in the SPSA case as the level of asynchrony between the SD pairs increases, on the average, the magnitude of the error term in measurements tends to become smaller since the duration of the interval over which the measurements overlap with each other gets shorter, and this may cause a marginal increase in the overall system performance. As we will see in Section 5, these two effects mainly cancel each other and the performance of the algorithm does not degrade significantly if the time lag is not large.

4 Implementation Issues

In this section, we give a brief overview of an overlay architecture used to enable traffic engineering capabilities. The details of the overlay architecture can be found in [5].

4.1 Path Establishment

Alternative paths between SD pairs are created using overlay nodes. The overlay nodes are located at all source-destination nodes as well as at some core nodes. The basic idea is similar to the ones presented in [17] and [18], with the difference being that the overlay is implemented within a domain (i.e., intra-domain) as opposed to inter-domain. When a packet is sent along the default path (*e.g.*, the shortest path), it is forwarded in the same way as in traditional

IP networks. However, when a packet is to be routed through an alternative path, it is processed at the source overlay node and an additional IP header is attached to the packet. This way the packet is first tunneled to a specific overlay node that lies in the selected alternative path, using the underlying routing protocol in use. When the overlay node receives the packet, it removes the outer IP header and forwards the packet to the final destination (or possibly to another overlay node). It is plain to see that one can utilize as many alternative paths as needed by placing overlay nodes at appropriate locations. Note that using this architecture, we can still employ the simple shortest path routing inside the network, without having to modify the existing traditional routers. The overlay capabilities can be realized by attaching a simple device (*e.g.*, a host with network processor) to the existing routers. This device simply processes the packets, and adds or removes IP headers before the basic forwarding operation is performed at the routers.

As a final remark, we would like to emphasize the point that the basic mechanism of the proposed routing algorithm can be employed in different types of networks. For instance, it can also be deployed in an MPLS based network, where the overlay paths are replaced with LSPs (Label Switched Paths). The proposed use of overlay architecture described in this paper enables us to adopt the algorithm in the traditional IP-based networks.

4.2 Traffic Monitoring

Traffic monitoring is also handled by the overlay architecture. Each link in the network is mapped to the closest overlay node with a tie-breaking rule that gives a unique mapping [5]. Overlay nodes periodically poll the links that they are responsible for, process the data and forward necessary local state information to the SD pairs utilizing the corresponding links. This eliminates the need for *each* SD pair to probe the links used by it. Note that, before forwarding the link cost information to sources nodes of SD pairs, the overlay nodes can aggregate the information gathered from different links. For example, due to the additive cost structure (according to the definition given in (3)), if the overlay nodes are aware of L^s , an overlay node can first compute the sum of the link costs over the links in L^s it is responsible for and then report only the total cost to the source node of SD pair s .

As a consequence, the overhead caused by the distribution of the link state information is minimized.

5 Experimental Setup and Simulation Results

In this section we evaluate the performance of the proposed algorithms through simulations under various network conditions, and study their convergence rates. To this end, we developed a packet level discrete-event simulator to carry out the simulations. Each plot presented in this section is the average of 10 independent simulation runs with different random seeds.

For the simulations we select the cost function of the form

$$C_l(\mathbf{x}) = d_l(\mathbf{x}) + u_l(\mathbf{x})^2 \quad (12)$$

where $d_l(\mathbf{x})$ is the number of packets dropped on link l ,⁵ and $u_l(\mathbf{x})$ is the link utilization. In all simulations, the measurement period is selected to be one second. As a consequence, SD pairs can update their rates at best approximately every 2 seconds since they require two measurements for estimating the gradient vector using the SPSA algorithm.

The requirements on the cost function are stated in Assumption A1, and one can argue that the selected cost function satisfies these requirements as follows. Since we deal with backbone networks, the packet arrival process at a source node is an aggregate of many individual flows. We assume that each individual flow generates packets according to an equilibrium renewal process, i.e., interarrival times of packets from a flow have a fixed distribution, and these equilibrium renewal processes are mutually independent. Then, by the Palm-Khintchine theorem [19], the superposition of these independent renewal processes can be approximated by a Poisson process, where interarrival times of packets are exponentially distributed.

In addition, according to the work presented in [20], the packet size distribution of Internet traffic has two peaks at 500 and 1,500 bytes. Using this observation, we can approximate the packet size as a Bernoulli random variable with values at 500 and 1,500 bytes.

Under the above conditions, we can approximate the links in the network as $M/G/1/K$ queues, where K is the buffer size. Following this assumption we can justify the assumption on convexity of the cost function as follows. One can check that in the regime of interest (*e.g.*, with utilization level being less than 150 percent), the link cost function is convex in the case of $M/M/1/K$ queue. In the case of $M/G/1/K$ queue one can show that the approximation functions

⁵ In the simulation we have used the observed number of packet drops, while in the model this term in the cost function can be interpreted as the expected number of packet losses as a function of the link load.




Fig. 1. Network Topology 1

TABLE I
THE CROSS TRAFFIC DYNAMICS

Link	Load Distribution in time (sec)		
	[0 – 1000)	[1000 – 2500)	[2500 – 3600)
L1	0.77	0.44	0.44
L2	0.33	0.33	0.67
L3	0.33	0.33	0.33

for blocking probability of an $M/G/1/K$ queue (*e.g.*, Gelenbe’s formula [21] and two-moment approximation in [22]), are indeed convex in the regime of interest under various parameter settings.

Simulations are carried out under two different two network topologies. The first topology, which is used in [4], is given in Fig. 1. Due to its simplicity this topology allows us to obtain insights into the fundamental behavior of the proposed algorithm. In addition, it serves us as a base setup so that we can make a comparison with the MATE algorithm presented in [4]. We have three SD pairs (S1-D1, S2-D2 and S3-D3) and each pair has two distinct paths. Note that this creates a considerable amount of interaction between these SD pairs.

The setup we use is similar to the one adopted in [4] for comparison purposes. (See [23] for the details of this setup.) The network consists of identical links with a bandwidth of 45 Mbps. The average packet size is 257 bytes. Each SD pair initially uses only the default shortest (minimum hop distance) path. Since all paths have equal length, the default min-hop paths are selected such that L2 is along the default shortest path of S1-D1, while the default shortest paths of S2-D2 and S3-D3 both traverse L3. Each SD pair generates traffic according to a Poisson process with an average rate of 19.8 Mbps (corresponding to 0.44 link utilization). In addition, links L1, L2 and L3 carry uncontrolled cross traffic generated by Poisson processes. The average rate of cross traffic normalized by the link capacity is given in Table I. A random delay is introduced before each SD pair starts running the optimal routing algorithm to ensure that the SD pairs are not synchronized. (The maximum value of this random delay is defined as offset.) In this simulation a decreasing step size policy is adopted that satisfies the conditions in Assumptions A3-A6. Specifically, the step sizes are given by $a_s(k) = 15/(k + 100)^{0.602}$ and $c_s(k) = 75/(k^{0.101})$.⁶

As shown in Figs. 2 and 3, the algorithm quickly eliminates the congestion and successfully balances the traffic. Moreover, these results show that the

⁶ $a_s(k)$ and $c_s(k)$ are reset to their initial values at simulation times 1,001 and 2,501 seconds for faster convergence.

proposed algorithm clearly outperforms the MATE algorithm [4]; while the MATE algorithm requires around 400-500 seconds to converge,⁷ our algorithm takes around 200 seconds. Furthermore, our algorithm quickly (around 50 seconds) eliminates packet drops unlike MATE. (See Figs. 10 and 11 presented in [4].)

⁷ Since simulation code and packet size distributions for the MATE algorithm are proprietary, it was not possible to simulate MATE. Therefore, we base our comparison on the results presented in [4].




Fig. 2. Network topology 1 with an offset of 50 ms




Fig. 3. Network topology 1 with an offset of 50 ms




Fig. 4. Network topology 1 with an offset of 200 ms




Fig. 5. Network topology 1 with an offset of 500 ms

Figs. 4 and 5 illustrate the effect of increased asynchrony between SD pairs. We increase the asynchrony between SD pairs by simply increasing the offset values. From these plots we can conclude that the algorithm is still able to converge in a short time. As we see from Figs. 2 and 4, the performance is almost the same for offset values of 50 ms and 200 ms. However, when we increase the offset to 500 ms, we see that the convergence of our algorithm becomes slightly slower. Thus, these results validate the earlier discussion made in subsection 3.3.




Fig. 6. Network Topology 2

Fig. 6 represents the second topology we consider in this paper. This topology is also used in [24], [25] and closely resembles the MCI Internet topology presented in [26]. Using this topology, we analyze the performance of the proposed algorithm under more realistic network conditions.

Nodes 1, 5, 6, 14 and 18 serve as both source and destination nodes. This gives us a total of 20 SD pairs. Each pair has at least two paths to reach the destination. A total of 78 paths are created between these 20 SD pairs, using overlay architecture. Overlay capability is available at all source/destination nodes as well as nodes 2, 10 and 13. In this experiment, the offset is set to 0.1 sec. The links shown as dashed lines have a capacity of 50 Mbps, while the links represented by solid lines have a capacity of 20 Mbps. The packet size for this scenario is selected to be 500 bytes. All SD pairs initially use only the shortest paths. Each SD pair generates traffic with a rate of 11.5 Mbps. In addition, the cross traffic traverses link (3-12), starting at simulation time 1,600 sec. The cross traffic rate is 18 Mbps and cannot be shifted to any alternative paths as before.

In Fig. 7, we illustrate how the load is distributed by our algorithm.⁸ The links for which we plotted the load are selected in such a way that each of them is located on a distinct alternative path that can be used to divert the traffic sent on link (3-12). The only exception is link (12-16), which is a downstream link of link (3-12) for some of the paths utilized by SD pairs. The plot of traffic on link (12-16) shows how the traffic load is migrated away from the paths that traverse link (3-12). In addition, Fig. 8 shows the total number of packets dropped in the entire network. We observe from both figures that the algorithm quickly removes congestion as the packet drops are eliminated fast and distributes the load among the multiple paths between the SD pairs in a reasonable amount of time. This result indicates that the proposed algorithm successfully converges under the scenarios where many SD pairs operate in an

⁸ $a_s(k)$ is reset to its initial value (0.15) at simulation time 1601 seconds for faster convergence.

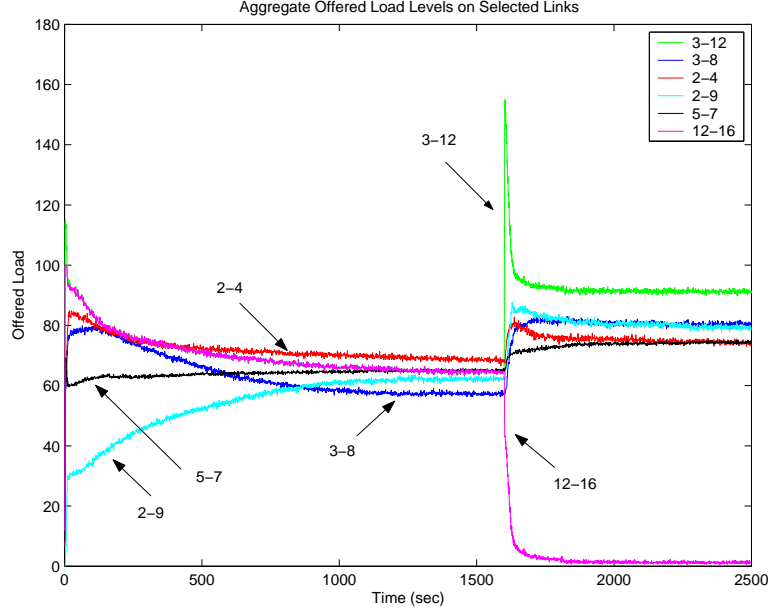


Fig. 7. Offered Load on Network Topology 2




Fig. 8. Total packet drops on Network Topology 2

independent and asynchronous fashion.

Fig. 9 demonstrates the effect of using a constant step size. We plot the variation of traffic load on link (3-8) using two different values of fixed step size $a_s(k) = a = 0.15$ and $a_s(k) = a = 0.225$, and compare their performance with the decreasing step size case where the step size for each SD pair is given by $a_s(k) = 15/(k+100)^{0.602}$. Even though we do not have almost sure convergence with constant step size, our results suggest that the convergence in distribution is sufficient for practical purposes. Moreover, the initial convergence rate

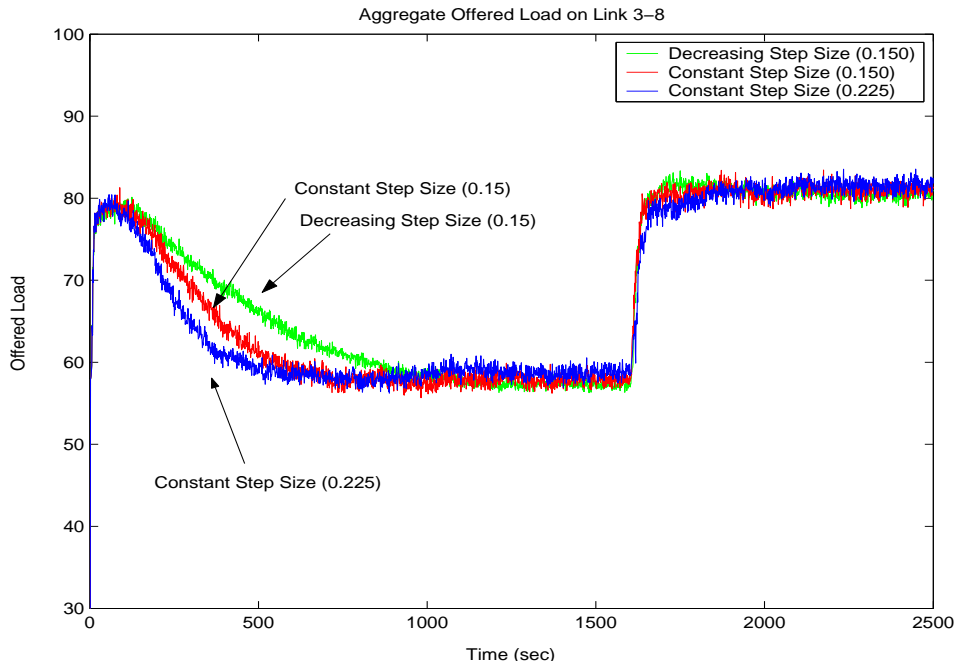


Fig. 9. Offered Load on Link 3-8

improves slightly as the algorithm is able to reach a small neighborhood of the optimal operating point faster compared to the decreasing step size case. Since this neighborhood appears to be small and a constant step size policy gives us the robustness to track dynamical changes in the network, this result supports the use of a constant step size algorithm in a practical implementation.

6 Conclusion

In this paper, we have considered optimal multi-path routing in environments where the link cost derivatives can be estimated (but for which an analytic expression may not exist). We have mathematically proven the optimality and stability of our proposed scheme, which is based on simultaneous perturbations. We have demonstrated that an SPSA algorithm provides significant improvements over an algorithm based on traditional finite-difference methods. Specifically, we have shown that our scheme results in much shorter measurement periods during the gradient estimation phase, and as a result, converges faster. Our simulation results show that our scheme can quickly alleviate network congestion and distribute load efficiently under dynamic network conditions.

Our work provides a basis for an architecture for effective traffic engineering in IP Networks. A promising extension to this work would be the adaption of our scheme to a Differentiated Services (DiffServ) capable network.

A Appendix I

Proof of Proposition 1

Collecting the terms of (8) for all SD pairs we have

$$\mathbf{x}(k+1) = \Pi_{\Theta}[\mathbf{x}(k) - \mathbf{A}(k)\hat{\mathbf{g}}(k)], \quad (\text{A.1})$$

where $\hat{\mathbf{g}}(k) := (\hat{\mathbf{g}}_s(k), s \in S)$, and $\mathbf{A}(k)$ is an $N \times N$ diagonal matrix and the diagonal entries of $\mathbf{A}(k)$ are given by the corresponding step sizes of SD pairs, $a_s(k)$. This is different from the general stochastic approximation algorithms in the sense that the step size $a(k)$ is no longer a scalar. Following [27, Theorem 5.3.1] closely, we will utilize the so-called “ODE method” to show convergence. We rewrite (A.1) in the following form

$$\mathbf{x}(k+1) = \mathbf{v}(k) + \boldsymbol{\tau}(k) + \boldsymbol{\phi}(k) \quad (\text{A.2})$$

where

$$\begin{aligned} \mathbf{v}(k) &= \mathbf{x}(k) + \mathbf{A}(k)[- \nabla C(\mathbf{x}(k)) + \boldsymbol{\xi}(k) + \mathbf{b}(k)], \\ \boldsymbol{\xi}(k) &= \mathbf{E}[\hat{\mathbf{g}}(k)|\mathbf{x}(k)] - \hat{\mathbf{g}}(k), \\ \mathbf{b}(k) &= \nabla C(\mathbf{x}(k)) - \mathbf{E}[\hat{\mathbf{g}}(k)|\mathbf{x}(k)], \\ \boldsymbol{\tau}(k) &= \Pi_{\Theta}[\mathbf{v}^{\gamma}(k)] - \mathbf{v}^{\gamma}(k), \\ \mathbf{v}^{\gamma}(k) &= \mathbf{x}(k) + \mathbf{A}(k)[- \nabla C(\mathbf{x}(k)) + \boldsymbol{\xi}(k) + \mathbf{b}(k)]I(k), \\ \boldsymbol{\phi}(k) &= \mathbf{v}^{\gamma}(k) - \mathbf{v}(k) + (\Pi_{\Theta}[\mathbf{v}(k)] - \mathbf{x}(k))(1 - I(k)), \end{aligned}$$

$I(k)$ denotes the indicator function of the event $\|\mathbf{A}(k)\boldsymbol{\xi}(k)\| \leq \gamma(k)/2$, and $\gamma(k)$ is a sequence of positive real numbers such that (i) $\gamma(k) \rightarrow 0$ and (ii) $\|\mathbf{A}(k)\boldsymbol{\xi}(k)\| \leq \gamma(k)/2$ for all but a finite number of k w.p. 1. The following lemma guarantees the existence of such a sequence.

Lemma 1 *Under the assumptions A1-A4, for the SPSA gradient estimator $\hat{\mathbf{g}}(k)$ defined in (9), the bias and error terms given by $\mathbf{b}(k)$ and $\boldsymbol{\xi}(k)$, respectively, satisfy*

- (a) $\mathbf{b}(k) \rightarrow 0$ w.p. 1,
- (b) $\sum_{k=0}^{\infty} \mathbf{E}[\|\mathbf{A}(k)\boldsymbol{\xi}(k)\|^2] < \infty$ w.p. 1.

Proof See Appendix II.

Since $\sum_{i=k}^j \mathbf{A}(i)\boldsymbol{\xi}(i)$, $j \geq k$ is a martingale and $\sum_{i=k}^{\infty} \mathbf{E}[\|\mathbf{A}(i)\boldsymbol{\xi}(i)\|^2] < \infty$,

from Doob-Kolmogorov inequality ([28, p. 309]), for each $\epsilon > 0$, we have

$$P(\sup_{j \geq k} \|\sum_{i=k}^j \mathbf{A}(i)\boldsymbol{\xi}(i)\| \geq \epsilon) \leq \frac{1}{\epsilon^2} \sum_{i=k}^{\infty} \mathbf{E} [\|\mathbf{A}(i)\boldsymbol{\xi}(i)\|^2] . \quad (\text{A.3})$$

Because the right-hand side of (A.3) goes to zero as $k \rightarrow \infty$, it follows that

$$\lim_{k \rightarrow \infty} P(\sup_{j \geq k} \|\sum_{i=k}^j \mathbf{A}(i)\boldsymbol{\xi}(i)\| \geq \epsilon) = 0 . \quad (\text{A.4})$$

and the existence of the $\gamma(k)$ sequence is guaranteed.

We interpolate $\mathbf{x}(k)$ into a continuous parameter process $\mathbf{X}^0(t)$, which is used to get the “mean” ODE, as follows:

$$\mathbf{X}^0(t) := \frac{t_{k+1} - t}{\hat{a}(k)} \mathbf{x}(k) + \frac{t - t_k}{\hat{a}(k)} \mathbf{x}(k+1) , \quad t \in (t_k, t_{k+1}) \quad (\text{A.5})$$

where $t_k = \sum_{i=0}^{k-1} \hat{a}(i)$.⁹

Let

$$\begin{aligned} \mathbf{B}^0(t_k) &:= \sum_{i=1}^{k-1} \mathbf{A}(i)\mathbf{b}(i), & \mathbf{M}^0(t_k) &:= \sum_{i=1}^{k-1} \mathbf{A}(i)\boldsymbol{\xi}(i), \\ \mathbf{T}^0(t_k) &:= \sum_{i=1}^{k-1} \boldsymbol{\tau}(i), & \boldsymbol{\Phi}^0(t_k) &:= \sum_{i=1}^{k-1} \boldsymbol{\phi}(i), \end{aligned}$$

and the corresponding piecewise linear interpolations are defined as

$$\begin{aligned} \mathbf{B}^0(t) &:= \frac{t_{k+1} - t}{\hat{a}(k)} \mathbf{B}^0(t_k) + \frac{t - t_k}{\hat{a}(k)} \mathbf{B}^0(t_{k+1}) \quad t \in (t_k, t_{k+1}), \\ \mathbf{M}^0(t) &:= \frac{t_{k+1} - t}{\hat{a}(k)} \mathbf{M}^0(t_k) + \frac{t - t_k}{\hat{a}(k)} \mathbf{M}^0(t_{k+1}) \quad t \in (t_k, t_{k+1}), \\ \mathbf{T}^0(t) &:= \frac{t_{k+1} - t}{\hat{a}(k)} \mathbf{T}^0(t_k) + \frac{t - t_k}{\hat{a}(k)} \mathbf{T}^0(t_{k+1}) \quad t \in (t_k, t_{k+1}), \\ \boldsymbol{\Phi}^0(t) &:= \frac{t_{k+1} - t}{\hat{a}(k)} \boldsymbol{\Phi}^0(t_k) + \frac{t - t_k}{\hat{a}(k)} \boldsymbol{\Phi}^0(t_{k+1}) \quad t \in (t_k, t_{k+1}). \end{aligned}$$

Using (A.2), $\mathbf{X}^0(t)$ can be written as

⁹ See Assumption A6 for the definition of $\hat{a}(k)$.

$$\begin{aligned}
\mathbf{X}^0(t) &= \frac{t_{k+1}-t}{\hat{a}(k)}\mathbf{x}(k) + \frac{t-t_k}{\hat{a}(k)}\left(\mathbf{x}(k) + \mathbf{A}(k)\left[-\nabla C(\mathbf{x}(k)) + \boldsymbol{\xi}(k) + \mathbf{b}(k)\right] \right. \\
&\quad \left. + \boldsymbol{\tau}(k) + \boldsymbol{\phi}(k)\right) \\
&= \mathbf{x}(k) + \frac{t-t_k}{\hat{a}(k)}\left(\mathbf{A}(k)\left[-\nabla C(\mathbf{x}(k)) + \boldsymbol{\xi}(k) + \mathbf{b}(k)\right] + \boldsymbol{\tau}(k) + \boldsymbol{\phi}(k)\right) \\
&= \mathbf{x}(1) + \sum_{i=1}^{k-1}\left(\mathbf{A}(i)\left[-\nabla C(\mathbf{x}(i)) + \mathbf{b}(i) + \boldsymbol{\xi}(i)\right] + \boldsymbol{\tau}(i) + \boldsymbol{\phi}(i)\right) \\
&\quad + \frac{t-t_k}{\hat{a}(k)}\left(\mathbf{A}(k)\left[-\nabla C(\mathbf{x}(k)) + \boldsymbol{\xi}(k) + \mathbf{b}(k)\right] + \boldsymbol{\tau}(k) + \boldsymbol{\phi}(k)\right) \\
&= \mathbf{x}(1) + \mathbf{B}^0(t) + \mathbf{M}^0(t) - \sum_{i=1}^{k-1}\hat{a}(i)\nabla C(\mathbf{x}(i)) \\
&\quad - \frac{t-t_k}{\hat{a}(k)}\hat{a}(k)\nabla C(\mathbf{x}(k)) + \mathbf{Z}^0(t) + \mathbf{T}^0(t) + \boldsymbol{\Phi}^0(t) \\
&= \mathbf{x}(1) + \mathbf{B}^0(t) + \mathbf{M}^0(t) + \mathbf{H}^0(t) + \mathbf{Z}^0(t) + \mathbf{T}^0(t) + \boldsymbol{\Phi}^0(t), \quad (\text{A.6})
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{Z}^0(t) &= \sum_{i=0}^{k-1}(\hat{a}(i) - \mathbf{A}(i))\nabla C(\mathbf{x}(i)) + \frac{t-t_k}{\hat{a}(k)}(\hat{a}(k) - \mathbf{A}(k))\nabla C(\mathbf{x}(k)), \\
\mathbf{H}^0(t) &= -\int_0^t \nabla C(\bar{\mathbf{x}}(s))ds, \\
\bar{\mathbf{x}}(s) &= \mathbf{x}(k), \quad s \in [t_k, t_{k+1}).
\end{aligned}$$

Since we are interested in the tail properties of the interpolated process, let us define the left shifted and centered process $\mathbf{X}^k(t)$:

$$\mathbf{X}^k(t) := \mathbf{x}(k) + \mathbf{B}^k(t) + \mathbf{M}^k(t) + \mathbf{H}^k(t) + \mathbf{Z}^k(t) + \mathbf{T}^k(t) + \boldsymbol{\Phi}^k(t), \quad (\text{A.7})$$

where

$$\begin{aligned}
\mathbf{B}^k(t) &= \mathbf{B}^0(t_k + t) - \mathbf{B}^0(t_k), & \mathbf{M}^k(t) &= \mathbf{M}^0(t_k + t) - \mathbf{M}^0(t_k), \\
\mathbf{T}^k(t) &= \mathbf{T}^0(t_k + t) - \mathbf{T}^0(t_k), & \boldsymbol{\Phi}^k(t) &= \boldsymbol{\Phi}^0(t_k + t) - \boldsymbol{\Phi}^0(t_k), \\
\mathbf{Z}^k(t) &= \mathbf{Z}^0(t_k + t) - \mathbf{Z}^0(t_k), & \mathbf{H}^k(t) &= -\int_0^t \nabla C(\bar{\mathbf{x}}(t_k + s))ds.
\end{aligned}$$

Similarly as in [27], we work with a fixed $w \notin \Omega_0$ where Ω_0 denotes the set for which $\|\mathbf{A}(k)\boldsymbol{\xi}(k)\| \geq \gamma(k)/2$ infinitely often or the conditions given in Lemma 1 do not hold. In this case $\mathbf{M}^k(t)$ and $\mathbf{B}^k(t)$ converge to zero uniformly on

finite intervals as $k \rightarrow \infty$. In addition, $\Phi^k(t) \rightarrow 0$ uniformly on finite intervals as $k \rightarrow \infty$ because we have

$$\begin{aligned} \|\phi(k)\| \leq & \| -\mathbf{A}(k)\nabla C(\mathbf{x}(k) + \mathbf{A}(k)\boldsymbol{\xi}(k) + \mathbf{A}(k)\mathbf{b}(k)) \| (1 - I(k)) \\ & + \|\Pi_\Theta[\mathbf{v}(k)] - \mathbf{x}(k)\| (1 - I(k)). \end{aligned}$$

and only a finite number of terms of $(1 - I(k))$ are nonzero.

Under assumption A6, using similar arguments one can easily verify that $\mathbf{Z}^k(t) \rightarrow 0$ as $k \rightarrow \infty$. Therefore, all the time varying terms on the right side of (A.7) other than $\mathbf{H}^k(t) + \mathbf{T}^k(t)$ go to zero as $k \rightarrow \infty$, and Assumption A5 is sufficient to guarantee the equicontinuity of $\mathbf{H}^k(t) + \mathbf{T}^k(t)$ as shown in [27, Theorem 5.3.1]. Hence, as a consequence of Arzela-Ascoli Theorem ([11, p. 72]), there exists a convergent subsequence $\mathbf{X}(t) \rightarrow \mathbf{x}^*$, which in fact results in $\mathbf{x}(k) \rightarrow \mathbf{x}^*$, and as shown in [27, Theorem 5.3.1], \mathbf{x}^* is determined by the solution of the following ODE:

$$\dot{\mathbf{x}} = \bar{\Pi}_\Theta[-\nabla C(\mathbf{x})], \quad (\text{A.8})$$

where

$$\bar{\Pi}_\Theta[-\nabla C(\mathbf{x})] := \lim_{\delta \downarrow 0} \frac{\Pi_\Theta[\mathbf{x} - \delta \nabla C(\mathbf{x})] - \mathbf{x}}{\delta}.$$

Note that the stationary points of the ODE above is the set of Kuhn-Tucker points

$$KT = \{\mathbf{x} : \exists \lambda^i \geq 0 \text{ such that } \nabla C(\mathbf{x}) + \sum_{i: \mathbf{q}_i(\mathbf{x})=0} \lambda^i \mathbf{q}_i(\mathbf{x}) = 0, \} \quad (\text{A.9})$$

where $\mathbf{q}_i(\cdot)$ terms are the constraints of the problem. Since the cost function $C(\mathbf{x})$ is assumed to be convex, any KT point belongs also to \mathbf{x}^* which concludes the proof.

B Appendix II

Proof of Lemma 1

Given any $N_s \times 1$ vector $\boldsymbol{\vartheta}$, the solution of the minimization problem

$$\begin{aligned} \min_{\boldsymbol{\eta}} \quad & \|\boldsymbol{\vartheta} - \boldsymbol{\eta}\|^2 \\ \text{s. t.} \quad & \boldsymbol{\eta}^T \mathbf{u} = r_s \end{aligned} \quad (\text{B.1})$$

is given by

$$\eta_i = \vartheta_i + \frac{r_s - \sum_{j=1}^{N_s} \vartheta_j}{N_s}, \quad (\text{B.2})$$

where $\mathbf{u} = [1, 1, \dots, 1]^T$. Obviously, if $\eta_i \geq 0$ for all i , this solution is equivalent to the L_2 projection. Here for the purpose of temporary perturbation we replace (5) with a non-negativity constraint. Thus, the projection of $\mathbf{x}_s(k) + c_s(k)\mathbf{\Delta}_s(k)$ can be calculated using (B.2) if

$$x_{s,i}(k) + c_s(k) \left(\Delta_{s,i}(k) - \frac{\sum_{j=1}^{N_s} (\Delta_{s,j}(k))}{N_s} \right) \geq 0. \quad (\text{B.3})$$

Recall that $\Delta_{s,i}(k)$ is bounded by α from Assumption A2. Hence, (B.3) holds if

$$c_s(k) \leq \frac{\min_j \{x_{s,j}(k)\}}{2\alpha}. \quad (\text{B.4})$$

From (5) we know $\frac{\min_j \{x_{s,j}(k)\}}{2\alpha} \geq \frac{\epsilon}{2\alpha}$. Since, $c_s(k) \rightarrow 0$, there exists finite K_1 such that $c_s(k) \leq \frac{\epsilon}{2\alpha}$ for all $k > K_1$. Therefore, (B.2) can be used to compute the projection of $\mathbf{x}_s(k) + c_s(k)\mathbf{\Delta}_s(k)$ for sufficiently large $k > K_1$.

Let $\bar{\mathbf{\Delta}}_s(k)$ be an $N \times 1$ vector, where values of entries corresponding to those of SD pair s are $\Delta_{s,i}(k)$ and zero otherwise. Hence, $\sum_{s \in S} \bar{\mathbf{\Delta}}_s(k) = (\Delta_{s,i}, s \in S, i \in P_s)$. Similarly, \mathbf{u}_s is an $N \times 1$ vector, where the values of entries corresponding to those of SD pair s are one and the remaining entries are zero. Following the proof in [10] and using Taylor's theorem, for $k > K_1$ and $s \in S$ we have

$$\begin{aligned} \mathbf{E}[\hat{g}_{s,i}(k)|\mathbf{x}(k)] &= \frac{N_s}{N_s - 1} \mathbf{E} \left[\frac{C_s^+(k) - C_s^-(k) + \mu_s^+(k) - \mu_s^-(k)}{c_s(k)\Delta_{s,i}(k)} \middle| \mathbf{x}(k) \right] \\ &= \frac{N_s}{N_s - 1} \mathbf{E} \left[\frac{\mathbf{E}[C_s^+(k) - C_s^-(k)|\mathbf{\Delta}(k)]}{c_s(k)\Delta_{s,i}(k)} \middle| \mathbf{x}(k) \right] \\ &= \frac{N_s}{N_s - 1} \left(\mathbf{E} \left[\frac{\nabla \Lambda_s^T(\mathbf{x}(k)) \sum_{s' \in S} c_{s'}(k) \bar{\mathbf{\Delta}}_{s'}(k)}{c_s(k)\Delta_{s,i}(k)} \middle| \mathbf{x}(k) \right] \right. \\ &\quad \left. - \mathbf{E} \left[\frac{\nabla \Lambda_s^T(\mathbf{x}(k)) \sum_{s' \in S} \frac{c_{s'}(k) \sum_{j=1}^{N_{s'}} \Delta_{s,j}(k)}{N_{s'}} \mathbf{u}_{s'}}{c_s(k)\Delta_{s,i}(k)} \middle| \mathbf{x}(k) \right] \right. \\ &\quad \left. + \mathbf{E} \left[\sum_{s' \in S} \frac{O(c_{s'}^2(k) \mathbf{\Delta}_{s'}^2(k))}{c_s(k)\Delta_{s,i}(k)} \right] \right) \\ &= \frac{N_s}{N_s - 1} \left(\frac{N_s - 1}{N_s} \nabla \Lambda_{s,i}(\mathbf{x}(k)) + O(c_s(k)) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{l \in L} \frac{\partial C_l(x^l(k))}{\partial x_{s,i}} + O(c_s(k)) \\
&= \frac{\partial C(\mathbf{x}(k))}{\partial x_{s,i}} + O(c_s(k))
\end{aligned} \tag{B.5}$$

where $C_s^-(k) = \Lambda_s(\mathbf{x}(k)) (= \sum_{l \in L^s} C_l(x^l(k)))$, $C_s^+(k) = \Lambda_s(\mathbf{x}(k) + \sum_{s' \in S} c_{s'}(k) \tilde{\Delta}_{s'}(k))$, $\nabla \Lambda_s(\mathbf{x}(k)) := (\partial \Lambda_s(\mathbf{x}(k)) / \partial x_{sp}; s \in S, p \in P_s)$, and

$$\tilde{\Delta}_{s'}(k) = \bar{\Delta}_{s'}(k) - \frac{\sum_{j=1}^{N_{s'}} \Delta_{s',j}(k)}{N_{s'}} \mathbf{u}_{s'}.$$

Therefore, one can see that $\mathbf{b}(\mathbf{k}) \rightarrow 0$ with probability one.

From the assumption that $E[\mu_s^+(k) - \mu_s^-(k) | \mathcal{F}_k] = 0$ and using the independence of $\mu_s^\pm(k)$ and $\Delta_s(k)$, we can bound the second moment of $\hat{\mathbf{g}}_s(k)$ as follows:

$$\begin{aligned}
\mathbf{E} [(\hat{g}_{s,i}(k))^2] &= \mathbf{E} \left[\left(\frac{C_s^+(k) - C_s^-(k) + \mu_s^+(k) - \mu_s^-(k)}{c_s(k) \Delta_{s,i}(k)} \right)^2 \right] \\
&= \mathbf{E} \left[\left(\frac{C_s^+(k) - C_s^-(k)}{c_s(k) \Delta_{s,i}(k)} \right)^2 + \left(\frac{\mu_s^+(k) - \mu_s^-(k)}{c_s(k) \Delta_{s,i}(k)} \right)^2 \right]
\end{aligned} \tag{B.6}$$

Following a similar argument used above one can show that the first term in (B.6) is $O(1)$ and using the bounds on $\mathbf{E}[(\Delta_s(k))^2]$, $\mathbf{E}[(\Delta_s(k))^{-2}]$, and $\mathbf{E}[(\mu_s^\pm(k))^2]$ the second term is $O(c_s(k)^{-2})$. Therefore, $\sum_{i=k}^\infty \mathbf{E}[\|\mathbf{A}(k)\boldsymbol{\xi}(k)\|^2] < \infty$ w. p. 1.

C Appendix III

Proof of Proposition 2

For the proof of the constant step size case, since the step size is fixed and the same for all SD pairs, we can directly apply the result of Theorem 8.2.1 of [27, p. 219]. Let $a_s(k) := a \leq 1$ and $c_s(k) := c = \frac{a \cdot \epsilon}{2\alpha}$. Then, as in the decreasing step size case with $k > K_1$, we can replace the projection of $\mathbf{x}_s(k) + c_s(k) \Delta_s(k)$ with its linear approximation. Hence $\mathbf{E}[\hat{g}_{s,i}(k) | \mathbf{x}(k)] = \frac{\partial C(\mathbf{x}(k))}{\partial x_{si}} + O(c)$ and we have

$$\lim_{m,k,a} \frac{1}{m} \sum_{i=k}^{k+m-1} \mathbf{E}[O(c) | \mathcal{F}_k] = \lim_{m,k,a} \frac{1}{m} \sum_{i=k}^{k+m-1} \mathbf{E} \left[O \left(\frac{a \cdot \epsilon}{2\alpha} \right) | \mathcal{F}_k \right] = 0 \text{ in mean.}$$

where the $\lim_{m,k,a}$ means that the limit is taken as $m \rightarrow \infty$, $k \rightarrow \infty$, and $a \rightarrow 0$ simultaneously in any way at all. Then, $\hat{\mathbf{g}}(k) = \{\hat{\mathbf{g}}_s(k), s \in S\}$ satisfies all the conditions given in Theorem 8.2.1 of [27, p. 219], establishing the convergence.

References

- [1] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja and X. Xiao, Overview and principles of internet traffic engineering, *RFC 3272* (2002).
- [2] B. Fortz and M. Thorup, Internet traffic engineering by optimizing OSPF weights, *Proceedings of the Conference on Computer Communications (IEEE Infocom)* (Tel-Aviv, Israel, 2000).
- [3] M. A. Rodrigues and K. G. Ramakrishnan, Optimal routing in shortest-path networks, *International Telecommunications Symposium (IEEE ITS)* (Rio de Janeiro, Brazil, 1994).
- [4] A. Elwalid, C. Jin, S. Low and I. Widjaja, MATE: MPLS adaptive traffic engineering, *Proceedings of the Conference on Computer Communications (IEEE Infocom)* (Anchorage, Alaska, 2001).
- [5] C. Kommareddy, T. Güven, B. Bhattacharjee, R. J. La and M. A. Shayman, Overlay routing for path multiplicity, *Tech. Rep. UMIACS-TR# 2003-70* <<http://www.cs.umd.edu/Library/TRs/CS-TR-4500/CS-TR-4501.pdf>>.
- [6] A. Elwalid, C. Jin, S. Low and I. Widjaja, MATE: Multipath adaptive traffic engineering, *Computer Networks-The International Journal of Computer and Telecommunications Networking* **40** (2002) 695–709.
- [7] J. C. Spall, Stochastic optimization, stochastic approximation and simulated annealing, *Encyclopedia of Electrical and Electronics Engineering (J.G. Webster, ed.)*, Wiley, New York **20** (1999) 529–542.
- [8] J. C. Spall, Stochastic optimization and the simultaneous perturbation method, *Proceedings of the Winter Simulation Conference* (Phoenix, AZ, 1999).
- [9] J. C. Spall, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, *IEEE Trans. Automat. Contr.* **37** (1992) 332–347.
- [10] M. Fu and S. D. Hill, Optimization of discrete event systems via simultaneous perturbation stochastic approximation, *Transactions of the Institute of Industrial Engineers* **29** (1997) 223–243.
- [11] H.J. Kushner and G.G. Yin, Stochastic approximation algorithms and applications, *Springer-Verlag* (1997).
- [12] J. N. Tsitsiklis and D. P. Bertsekas, Distributed asynchronous optimal routing in data networks, *IEEE Trans. Automat. Contr.* **31** (1986) 325–332.

- [13] D. Bertsekas and R. Gallager, Data networks, *Prentice-Hall Inc.*, 2nd edition (1992).
- [14] J. Kiefer and J. Wolfowitz, Stochastic estimation of a regression function, *Ann. Math. Stat.* **23** (1952) 462–466.
- [15] J. R. Blum, Multidimensional stochastic approximation methods, *Ann. Math. Stat.* **25** (1954) 737–744.
- [16] P. Sadegh, Constraint optimization via stochastic approximation with a simultaneous perturbation gradient approximation, *Automatica* **33** (1997) 889–892.
- [17] A. Collins, The detour framework for packet rerouting, *PhD thesis*, (University of Washington, 1998).
- [18] D. Andersen, H. Balakrishnan, F. Kaashoek and R. Morris, Resilient overlay networks, *Proc. 18th ACM Symp. on Operating Systems Principles (SOSP)* (Banff, Canada, 2001).
- [19] D. P. Heyman and M. J. Sobel, Stochastic models in operations research, *McGraw-Hill* (1982).
- [20] K. Claffy, G. Miller and K. Thompson, The nature of the beast: recent traffic measurements from an internet backbone, *Internet Society's Networking Conference (INET)* (Geneva, Switzerland, 1998).
- [21] E. Gelenbe, On approximate computer system models, *Journal ACM* **22** (1975) 261–269.
- [22] J.M. Smith and F.R.B. Cruz, The buffer allocation problem for general finite buffer queueing networks, <<http://www.ecs.umass.edu/mie/faculty/smith/>> *Unpublished*.
- [23] K. Sinha and S. Patek, OpIATE: optimization integrated adaptive traffic engineering, *Tech. Rep.* <<http://www.sys.virginia.edu/techreps/2002/sie-020001.pdf>> (2001).
- [24] S. Nelakuditi and Z. L. Zhang, A localized adaptive proportioning approach to QoS routing, *IEEE Commun. Mag.* **40** (2002) 66–71.
- [25] G. Apostolopoulos, R. Guerin, S. Kamat and S. Tripathi, Quality of service based routing: a performance perspective, *ACM SIGCOMM* (Vancouver, 1998).
- [26] Q. Ma and P. Steenkiste, On path selection for traffic with bandwidth guarantees, *IEEE International Conference on Network Protocols* (1997).
- [27] H.J. Kushner and D.S. Clark, Stochastic approximation methods for constrained and unconstrained systems, *Springer-Verlag*, (1978).
- [28] G.R. Grimmett and D.R. Stirzaker, "Probability and random processes, *Oxford Science Publications*, 2nd edition (1998).