

Traffic Engineering with Elastic Traffic

Chiun Lin Lim and Ao Tang

School of Electrical and Computer Engineering, Cornell University
{c1377, at422}@cornell.edu

Abstract—In the network layer, an internet service provider controls the traffic across an autonomous system by load balancing via traffic engineering and by varying the offered traffic of the users via feedback signals. In the transport layer, users send traffic into the network using the TCP protocol, which adjusts offered traffic according to the received feedback. We investigate how feedback and current traffic engineering practice interact with congestion control under the network utility maximization framework. We show that the current interaction is stable, increases network utility, but does not necessarily improve the traffic engineering objective. The mismatch in outcome and incentives prompt us for a more holistic view using game theory. With suitable modification of the feedback, we show that the interaction converges to a socially optimal solution for users running either primal or dual algorithms. We further show that the results hold even when traffic engineering is performed at any irregular intervals. More generally, we show via heterogeneous feedback the same optimality result for a mix of users running primal and dual algorithms.

I. INTRODUCTION

To load balance traffic across an autonomous system (AS), an internet service provider (ISP) carries out a procedure called traffic engineering. First, the ISP measures the offered traffic into the AS. Next, using the data, the ISP ideally solves the well-known multi-commodity flow problem [1] for a given cost metric on link utilization. Once the solution is obtained, the desired traffic flow pattern is achieved by the setting of link weights used in the shortest-path computation of link-state or distance-vector routing algorithms. However, it turns out that setting the right link-weight is NP-hard for link-state algorithms [2], and many works in the traffic engineering literature have focused on solving this issue with a fixed offered traffic assumption, for instance [3], [4]. Our focus here is not on the link-weight setting problem but on *relaxing the fixed offered traffic assumption and investigating the interaction of current traffic engineering practices with elastic traffic*.

From the perspective of a layered network architecture, our work deals with the interaction of transport layer (congestion control) and network layer (routing). Though we view our problem by starting from the network layer, it turns out that the problem could equally well be interpreted from the view of transport layer under the context of optimal joint congestion control and routing. While the traffic engineering framework has been developed with a constant offered traffic assumption, the congestion control optimization framework, i.e. network utility maximization (NUM) has been developed assuming the traffic flow from each user follows a single static path. Relaxing the static assumption on either framework brings us to the joint congestion control and routing problem. Regardless, the tone and main perspective of this paper are of the network layer.

For congestion control, investigations on relaxation of the fixed routing assumption has been around for quite some time. The joint congestion control and multipath routing problem is already present in Kelly's seminal work [5]. There, rate variables are defined for all possible paths from source to destination, and each user has to determine all these rates to maximize utility. More recent works have shifted the routing decisions to the nodes. Each edge router either separately chooses a route to minimize congestion cost [6] or determines, for each destination, the load to place on all possible paths [7]; or all routers control the per-destination split ratios on outgoing links [8]. The paper that is closest in essence to ours is [9] though our stability and optimality results are not limited to the ring topology and our timescale separation model is more general. Many other related papers could be found in [10].

Across the aisle, the traffic engineering community has much fewer cross-layer works. The hardness of the link-weight setting problem and the lack of a solid analytical framework compared to NUM has limited the cross-layer works, for instance [11], to be simulation in nature. There is a vast literature on traffic engineering, however, if we do not limit ourselves to communication networks. Research on transportation traffic has been around for decades [12]. Beside the convex optimization tools [13] employed in NUM, another important theoretical tool here is game theory [14] where agents are modeled to behave selfishly. Under the game theoretical framework, each agent acts independently to choose a minimum-cost path, given the congestion costs caused by the actions of all other players. The problem is usually modeled as a non-cooperative congestion game, which is guaranteed to have a pure-strategy Nash equilibrium [15]. It is well-known that in general Nash equilibria are not necessarily socially optimal. How inefficient any Nash equilibrium can be has been characterized with the concept of Price of Anarchy in [16]. These works have been done with a constant offered traffic assumption and extensions to elastic traffic could be found in [17], [18].

With both frameworks in mind, we model the interaction of traffic engineering and NUM as a two-stage iterative process in Section II and explain how feedback, traffic engineering, and NUM interact with each other. In Section III, we find that the process converges, improves network utility, but does not guarantee improvement in the traffic engineering's objective. We thus relax the capacity constraint and propose a modification under a game theory framework, albeit *users are not necessarily selfish but instead act according to NUM*. In Section IV, we formulate a potential game perspective for the primal algorithm of NUM and show that the primal algorithm

converges to the socially optimal solution. We next relax our timescale separation assumption and show that the same convergence holds even when traffic engineering is performed at any irregular intervals. We extend to the dual algorithm for NUM in Section V and show that with proper modification of the feedback signal, the convergence results also hold here. In Section VI, we relax the homogeneous user assumption by allowing heterogeneous users running primal or dual algorithms and prove the same optimality result. We conclude the paper in Section VII.

II. MODEL FORMULATION

We consider an autonomous system controlled by an ISP, whom we refer to as the traffic engineer henceforth. A set of N users, each representing a particular source-destination pair, send traffic into the AS. We assume users have infinite backlog and any user traffic could be split at any ratio across all possible paths from source to destination. An illustrative figure of the model is shown at Figure 1.

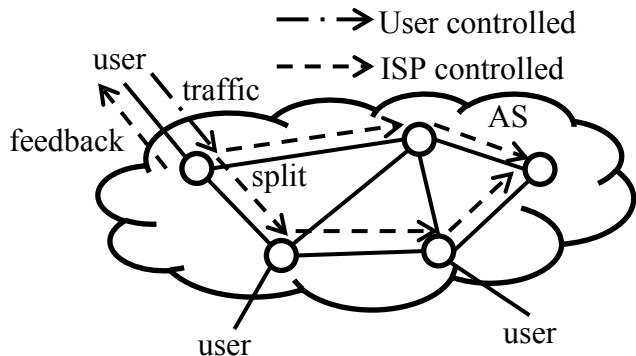


Figure 1. User controls input traffic while ISP owns the AS and controls feedback and traffic split at routers

Before proceeding further, for clarity, we explain the notational conventions adopted throughout the paper. Uppercase letters denote matrices, e.g. H , Q , R , or sets, e.g. L , N^1 , or constants, e.g. K , or utilities (or costs), e.g. U , Φ . Lowercase letter i represents user, j or k represents path, l represents link and other letters such as x , f , q denote vectors. Superscript denotes element of a vector, e.g. x^i or columns of a matrix, e.g. H^i associated with user i . Similarly, subscript on a vector or matrix, e.g. f_l , q_j^i denotes association with either link l or path j . A bar, e.g. \bar{x}^i , means the variable is an optimal solution. The bold number $\mathbf{1}$ is the vector of all ones while e_l or e_j^i is a unit vector for the corresponding element.

We model the AS as a graph with a set L of directed links. The links have finite capacities $c = \{c_l | l \in L\}$ and each link l is associated with a link-cost function $\Phi_l(f_l)$ that is convex and increasing in f_l , the traffic flow through link l . Examples

¹We will abuse notation by using L and N to denote both sets and their cardinality.

of $\Phi_l(\cdot)$ are link utilization f_l/c_l and link delay $f_l/(c_l - f_l)$ as given by the $M/M/1$ formula [19].

Each user $i \in N$ is associated with K^i acyclic paths, which are represented by a $L \times K^i$ 0-1 matrix H^i where

$$H_{lj}^i = \begin{cases} 1, & \text{if user } i \text{ uses link } l \text{ in path } j \\ 0, & \text{otherwise} \end{cases}$$

The matrix H^i does not necessarily contain all the possible paths from source to destination but instead could just be a subset of the paths. Let $K = \sum_i K^i$ and define the $L \times K$ matrix as

$$H = [H^1 \quad \dots \quad H^N]$$

A split ratio is specified over the K^i paths by the $K^i \times 1$ vector q^i where q_j^i represents the fraction of i 's flow on path j such that

$$q_j^i \geq 0, \quad \mathbf{1}^T q^i = 1$$

Collect the vectors q^i , $i = 1, \dots, N$ into a $K \times N$ block diagonal matrix $Q = \text{diag}(q^1, \dots, q^N)$. We represent the set of all possible traffic split matrices as

$$\mathcal{Q} = \{ Q | Q = \text{diag}(q^1, \dots, q^N) \in [0, 1]^{K \times N}, \mathbf{1}^T q^i = 1 \}$$

The matrix H defines the topology of the network while Q represents how the traffic of each user is split over the available paths. Their product is a $L \times N$ routing matrix $R = HQ$ with its R_{li} entry giving the fraction of i 's traffic at each link l .

The users and the traffic engineer interact in an iterative two-stage process. First, the traffic engineer, with knowledge of the offered traffic $x = \{x^i | i \in N\}$, solves a traffic engineering problem

$$(TE) \quad \min_{Q \in \mathcal{Q}} V(HQx) \\ \text{s.t.} \quad HQx = f \leq c$$

where $f = \{f_l | l \in L\}$ represents the load on all the links. Examples of the objective function V are maximum utilization, $\max_l f_l/c_l$, and for additive cost function, $\sum_l \Phi_l(e_l HQx) = \sum_l \Phi_l(f_l)$, in which case we could be minimizing total delay.

Additionally, the traffic engineer sets a congestion price for all links. The price is represented by a price function $\phi = \{\phi_l(f_l) : l \in L\}$ that is related to the link cost function Φ . For each path used, the edge router of user i receives the sum of all link prices along the path as feedback. Thus, the price information received by the edge router is the vector $(H^i)^T \phi \in \mathbb{R}^{K^i}$. The edge router then calculates the expected congestion price $\phi^T H^i q^i = (HQe_i)^T \phi$ and passes it to the user as a feedback value. Users update their offered traffic according to the feedback value.

Solving (TE) gives a new traffic split matrix Q , and users update their offered traffic as the feedback values have changed. The update is governed by the NUM framework where users solve the following problem: Given Q ,

$$(NUM) \quad \max_{x \geq 0} \sum_{i \in N} U^i(x^i) \\ \text{s.t.} \quad Rx \leq c$$

where $U^i(x^i)$ is a nonnegative, increasing, and strictly concave utility function. The NUM framework comes with two flavors, primal and dual. Under the primal formulation, the capacity constraint is relaxed by adding a barrier function to the objective

$$(NUM-P) \quad \max_{x \geq 0} \sum_{i \in N} U^i(x^i) - \sum_{l \in L} B_l(e_l^T R x)$$

By proper choice of a convex barrier function, $B_l(\cdot)$, the optimal solution of (NUM) could be approximated by solving (NUM-P) [20]. For single path, i.e. when Q is restricted to be a 0 – 1 matrix, the primal algorithm has the users and price function update as follows [21]:

$$\dot{x}^i = \kappa^i(x^i) \left[(U^i)'(x^i) - (Re^i)^T \phi \right] \quad (1)$$

$$\dot{\phi}_l = B_l' \quad (2)$$

where $\kappa^i(\cdot)$ is non-negative, increasing and continuous. For the dual formulation, (NUM) is solved directly and the dual algorithm has the users and price function update as follows [22]:

$$x^i = (U^i)^{\prime-1} \left((Re^i)^T \phi \right) \quad (3)$$

$$\dot{\phi}_l = \begin{cases} h_l(f_l - c_l) & , \phi_l > 0 \\ h_l \max\{f_l - c_l, 0\} & , \phi_l = 0 \end{cases} \quad (4)$$

where h_l is a positive constant.

We assume that the update on Q is instantaneous. We also assume for now that the users converge to the optimal solution of (NUM) before the traffic engineer updates Q again. This timescale separation assumption will be relaxed in the later section to cover not only the 3 timescale separation models mentioned in [9], but also allowing Q to be updated at any irregular intervals. With this assumption, the iterative two-stage process has the traffic engineer and the users taking turns solving their respective optimization problems:

$$Q(t+1) = \arg \min_{Q \in \mathcal{Q}: HQx(t) \leq c} V(HQx(t)) \quad (5)$$

$$x(t+1) = \arg \max_{x \geq 0: HQ(t+1)x \leq c} \sum_{i \in N} U^i(x^i) \quad (6)$$

When there are more than one optimal solution to (5), we adopt the convention that if $Q(t)$ is an optimal solution, then we choose $Q(t+1) = Q(t)$. If not, an optimal solution is chosen at random. This convention of always choosing the current solution if optimal holds as well for (6) and in the later optimization problems.

We pause here to point out the critical differences between our work in the most general setting and prior literature. Most importantly, we have relaxed the fixed offered assumption in traffic engineering and consider an elastic traffic modeled by (NUM). We stress again that our main perspective lies with the traffic engineer, though our work has several implications for the congestion control framework as well. First, we adopt the view that network infrastructure is owned and controlled by the

traffic engineer. Thus, the traffic engineer has a complete view and control of the network topology while users could obtain information about the paths used via feedback but beyond that have limited to no knowledge of the network topology and possible paths to destination. This is in direct contrast with [5] where users could choose which path to use and with [6] where edge routers act independently of each other. Second, prior work with globally optimal convergence results, notably [7], [8], require synchronous updates by the user and traffic engineer. This is a timescale mismatch as congestion control typically converges to equilibrium in the order of round-trip time while routing is generally updated on a longer timescale. The paper [9] investigates various timescale separation models but its stability and optimality results are limited to the ring topology. Here, we complete the stability and globally optimal convergence proof for the general topology and for the general timescale model where the traffic engineer could update at any arbitrary interval.

III. CURRENT INTERACTION AND MODIFICATIONS

We first show that the iterative two-state process is stable in general.

Theorem 1. *The iterative two-stage process (5), (6) converges to a fixed point.*

Proof: The sequence $\{\sum_i U^i(x^i(t)), t = 1, 2, \dots\}$ converges since it is bounded from above due to capacity constraint, and strictly increasing prior to convergence since at each time $t + 1$, $x(t)$ is a feasible solution of (6) and thus $\sum_i U^i(x^i(t+1)) \geq \sum_i U^i(x^i(t))$. If the inequality holds with equality, then the sequence has converged with $(Q(t+1), x(t))$ being a fixed point of equations (5) and (6). Note that the convention of always choosing the current solution if optimal prevents the possibility of oscillations. ■

Remark. While we have allowed the traffic engineer to split traffic arbitrarily on all possible paths, the proof still holds even if the traffic engineer is restricted to picking a fixed number of paths or even just one path for each user. In fact, in the single-path case, the process is guaranteed to converge in finite time since there is only a finite number of single-path routing configurations and the iterative process is guaranteed to pick a different configuration at each iteration. This is in contrast to results from [6] where the single-path case is shown to be unstable. As pointed out earlier, this is because the traffic engineer controls all routers instead of allowing each router to act independently of each other.

The proof shows that the two-stage iterative process improves network utility but the same guarantee does not necessarily hold for traffic engineering objective. Consider the following simple example.

Example 1. (No improvement for traffic engineer) Suppose the objective of (TE) is to minimize maximum utilization. The optimal solution for (NUM), however, will always have a bottleneck link, l where $\bar{f}_l = c_l$. This implies that the effort

of the traffic engineer is futile as at the end of each iteration, there will always be a link with maximum utilization.

Since the network utility $\sum_i U^i(x^i)$ strictly increases after every iteration prior to convergence, one could view the iterative process (5), (6) as trying to solve the joint optimization problem

$$\begin{aligned} \text{(JOINT)} \quad & \max_{x \geq 0, Q \in \mathcal{Q}} \sum_{i \in N} U^i(x^i) \\ & \text{s.t.} \quad HQx \leq c \end{aligned}$$

Equation (5) could now be interpreted as an update step that allows one to search \mathcal{Q} for better traffic split ratios. However, the iterative process is not guaranteed to reach the optimal solution of (JOINT) regardless of the cost function Φ as demonstrated in the following example.




Figure 2. A simple network used to illustrate why relaxing capacity constraint (Example 2) and assuming positive traffic (Example 3) are required

Example 2. (Capacity constraint prevents update of Q) Consider the network shown in Figure 2. The cost function for each link l is $\Phi_l(f_l) = f_l^2/2$. There are two users with the same strictly concave utility function $U(\cdot)$. Initially, the traffic engineer sets the 1 unit of offered traffic from user 1 to travel along $A \rightarrow D$ and splits the 2 units of traffic from user 2: 1 along $B \rightarrow D$ and 1 along $B \rightarrow C \rightarrow D$. Due to capacity constraints, regardless of the traffic engineering objective V , the optimal solution of (TE) maintains the current traffic split ratios and thus the iterative process has converged. However, the optimal solution of (JOINT) is $2U(1.5) > U(1) + U(2)$ by concavity of U . The iterative process fails to improve the solution further as the current solution is trapped due to capacity constraints.

Our goal now is to modify the iterative process such that the incentives of both the users and the traffic engineer are taken into account. Since we have two groups with incentives that are not necessarily aligned, the framework of game theory lends itself naturally. We stress that though we borrow the terminologies and tools of game theory, and we may refer to users as trying to maximize a certain utility, we do not assume users are rational and selfish but instead they act according to the governing algorithm.

If we view the group of users as a single player, then we can view the iterative process (5), (6) as a two-player game between the player and the traffic engineer. The player and the traffic engineer take turns to make a best response to the

opponent's strategy and we have shown in Theorem 1 that this best-response dynamics converges to a Nash equilibrium. It is natural to define the aggregate surplus or the social value of an offered traffic and traffic split combination (x, Q) as

$$S(x, Q) = \sum_{i \in N} U^i(x^i) - V(HQx)$$

Refining our goal further, we now want to modify the iterative process such that it has a fixed point that is socially optimal. We know from Example 2 that a solution may get trapped due to capacity constraints. This observation leads us to consider a relaxation of the capacity constraint. To do that, first we restrict the objective of (TE) to doubly differentiable and additive cost functions, $\sum_l \Phi_l(f_l)$. Next we assume that for all l , Φ_l is large when it is near or over capacity to penalize capacity overshooting. As with the primal formulation (NUM-P), for the users, the capacity constraint is replaced with Φ_l as the barrier function. We thus arrive at the following Gauss-Seidel system [23]:

$$\begin{aligned} Q(t+1) &= \arg \min_{Q \in \mathcal{Q}} \sum_{l \in L} \Phi_l(e_l^T HQx(t)) \quad (7) \\ x(t+1) &= \arg \max_{x \geq 0} \sum_{i \in N} U^i(x^i) \\ &\quad - \sum_{l \in L} \Phi_l(e_l^T HQ(t+1)x) \quad (8) \end{aligned}$$

But we are not quite done with the assumptions yet since a solution may also get trapped at the other extreme where $x^i = 0$.

Example 3. (Phantom user) We reuse the network shown in Figure 2. There are two users with the same strictly concave utility function $U^i(x^i) = -(x^i)^2 + 2x^i$, for $x^i \leq 1$. The first user sends traffic along $A \rightarrow D$ and $A \rightarrow C \rightarrow D$ while the second user sends traffic along $C \rightarrow D$. The link-cost function is linear $\Phi_l(f_l) = 3f_l/2$. Initially, user 1 is configured to send all its traffic along the second path, i.e. $q_{A \rightarrow C \rightarrow D}^1 = 1$. With this setup, solving (8) gives $x^1 = 0$, $x^2 = 1/4$. On the next iteration, the traffic split ratio remains the same by convention. The Gauss-Seidel system thus converges with a social value of $1/16$. However, the optimal social value is $1/8$ with $\bar{q}_{A \rightarrow D}^1 = 1$ and $\bar{x}^1 = \bar{x}^2 = 1/4$.

We therefore assume that a solution of (8) always gives $\bar{x} > 0$. The assumption makes sense in practice because the traffic engineer would not be able to know the existence of a user if it is not offering any traffic. This assumption is satisfied with, for instance, utility functions U satisfying $\lim_{x \rightarrow 0^+} dU/dx = \infty$, e.g. α -fair utility functions [24].

We next prove the optimality of the Gauss-Seidel system.

Theorem 2. *The Gauss-Seidel system (7), (8) converges to a socially optimal fixed point, (\bar{x}, \bar{Q}) .*

Proof: The social value $S(x, Q)$ is increasing after each update and is bounded and thus the Gauss-Seidel system converges to a fixed point (\bar{x}, \bar{Q}) . To show that the fixed point

is socially optimal, we rely on Karush-Kuhn-Tucker (KKT) conditions [13]. Both (7) and (8) are convex optimization problems that satisfy regularity conditions since the constraints are linear and thus the KKT conditions are necessary and sufficient. The KKT conditions for (7) state that \bar{Q} is optimal if and only if there exists Lagrange multiplier λ such that

$$(He_j^i)^T \Phi'(H\bar{Q}x) \begin{cases} = \frac{\lambda^i}{x^i}, & \text{if } \bar{q}_j^i > 0 \\ \geq \frac{\lambda^i}{x^i}, & \text{if } \bar{q}_j^i = 0 \end{cases} \quad (9)$$

where $\Phi'(H\bar{Q}x) = \{\Phi_l'(e_l^T H\bar{Q}x), l \in L\}$.

The KKT conditions for (8) state that $\bar{x} > 0$ is optimal if and only if

$$(U^i)'(\bar{x}^i) = (HQe^i)^T \Phi'(HQ\bar{x}) \quad (10)$$

However, the optimization problem

$$\max_{x \geq 0, Q \in \mathcal{Q}} S(x, Q) \quad (11)$$

is not necessarily convex. To apply the KKT conditions, we transform it into a convex problem with a change of variables

$$w_j^i = q_j^i x^i \text{ or } w = Qx \quad (12)$$

where w_j^i is now the traffic flow on path j for user i .

$$\max_{w \geq 0} \sum_{i \in N} U^i \left(\sum_j w_j^i \right) - \sum_{l \in L} \Phi_l(e_l^T Hw) \quad (13)$$

The optimization problem is now convex with linear constraints and thus the KKT conditions are necessary and sufficient. We denote its Lagrangian as $L_S(w; \theta)$. At $\bar{w}_j^i = \bar{q}_j^i \bar{x}^i$, since $\bar{x}^i > 0$, we start with the inverse relationship

$$x^i = \sum_j w_j^i, \quad q_j^i = \frac{w_j^i}{\sum_k w_k^i} \quad (14)$$

For $\bar{w}_j^i > 0$, choose $\theta_j^i = 0$ and apply chain rule

$$\frac{\partial L_S(\bar{w}, \theta)}{\partial w_j^i} \quad (15)$$

$$= \frac{\partial x^i}{\partial w_j^i} \frac{\partial L_S(\bar{Q}\bar{x}, \theta)}{\partial x^i} + \sum_k \frac{\partial q_k^i}{\partial w_j^i} \frac{\partial L_S(\bar{Q}\bar{x}, \theta)}{\partial q_k^i} \quad (16)$$

$$= (U^i)'(\bar{x}^i) - (HQe^i)^T \Phi'(H\bar{Q}\bar{x}) \quad (17)$$

$$+ (1 - \bar{q}_j^i) \left[-(He_j^i)^T \Phi'(H\bar{Q}\bar{x}) \right] \quad (18)$$

$$- \sum_{k \neq j} \bar{q}_k^i \left[-(He_k^i)^T \Phi'(H\bar{Q}\bar{x}) \right] \quad (19)$$

Equation (10) implies line (17) equals zero, while equation (9) implies lines (18) and (19) cancel off each other. One could easily verify that the other KKT conditions are satisfied with proper choices of θ and hence \bar{w} is an optimal solution of (13) implying that (\bar{x}, \bar{Q}) is the socially optimal solution. ■

With Theorem 2 in mind, we know we could achieve social optimality if the offered traffic of the users converges to a solution of

$$\max_{x \geq 0} \sum_{i \in N} U^i(x^i) - \sum_{l \in L} \Phi_l(e_l^T HQx) \quad (20)$$

So far, our results and examples have not touched on the dynamics of how users vary their offered traffic. We introduce in the modeling section how the expected congestion prices are passed to the users as feedback values and how users react to the feedback values according to either primal or dual algorithm. For the next two sections, we assume all users are running either the primal or the dual algorithm. Without altering the existing algorithms, our goal in these sections is to propose congestion prices such that the offered traffic of the users converges to a solution of (20).

IV. PRIMAL ALGORITHM AS A POTENTIAL GAME

We know the primal algorithm (1), (2) is globally asymptotically stable and converges to the solution of (20) for the single path case [21]. The crux of the proof is to show that the objective function of (20) acts as a Lyapunov function and is, in essence, similar to the proof technique of Theorem 1. The proof for the multipath case could easily be generalized using the same Lyapunov function argument but we will take a slightly different path by viewing the proof from the perspective of a potential game [25]. The potential game interpretation is useful when we relax the timescale assumptions and when we extrapolate the ideas here to the dual algorithm.

Consider the following game where each user i receives utility $U^i(x^i)$ for an offered traffic of x^i but is charged the sum of link costs of all links used i.e. $\mathbf{1}^T (H^i)^T \Phi(HQx)$ where $\Phi(HQx) = \{\Phi_l(e_l^T HQx), l \in L\}$. Thus, each user i tries to maximize the total utility:

$$\max_{x^i \geq 0} \Gamma^i(x) \triangleq U^i(x^i) - \mathbf{1}^T (H^i)^T \Phi(HQx)$$

Lemma 3. For fixed Q , the users are playing an exact potential game with exact potential function

$$P(x) = \sum_{i \in N} U^i(x^i) - \sum_{l \in L} \Phi_l(e_l^T HQx)$$

Proof: We check by definition that, for all $i \in N$

$$\frac{\partial \Gamma^i}{\partial x^i} = (U^i)'(x^i) - (HQe^i)^T \Phi'(HQx) = \frac{\partial P}{\partial x^i} \quad \blacksquare$$

An exact potential game has the property that any user i that unilaterally switches from rate x^i to y^i increases (or decreases) his utility by the same amount as the potential:

$$\Gamma^i(y^i, x^{-i}) - \Gamma^i(x^i, x^{-i}) = P(y^i, x^{-i}) - P(x^i, x^{-i}) \quad (21)$$

This property implies that it is possible to arrive at a Nash equilibrium if the users take turns performing its best response to the current strategy.

Lemma 4. For fixed Q , the unique Nash equilibrium, x_{NE} satisfies

$$x_{NE} = \arg \max_{x \geq 0} \sum_{i \in N} U^i(x^i) - \sum_{l \in L} \Phi_l(e_l^T H Q x) \quad (22)$$

Proof: At Nash equilibrium,

$$x_{NE}^i = \arg \max_{x_i \geq 0} U^i(x^i) - \mathbf{1}^T (H^i)^T \Phi(H Q x) \quad (23)$$

The optimization problem on the RHS of (23) is convex, has linear constraints and the KKT conditions state that for each user i

$$(U^i)'(x_{NE}^i) = (H Q e^i)^T \Phi'(H Q x_{NE})$$

The N conditions together constitute the KKT conditions of (22). Uniqueness arises from the fact that (20) is strictly concave. ■

We now show that the primal algorithm (1) generalizes to the multipath case.

Lemma 5. If users update their strategies according to gradient ascent and congestion prices are set to the first derivatives of link cost:

$$\dot{x} = \kappa^i(x^i) \left[(U^i)'(x^i) - (H Q e^i)^T \phi \right] \quad (24)$$

$$\dot{\phi} = \Phi'(H Q x) \quad (25)$$

then the strategies converge to an optimal solution of (20).

Proof: We show that the potential is always increasing

$$\begin{aligned} \dot{P}(x) &= \sum_i \frac{\partial P(x)}{\partial x^i} \dot{x}^i \\ &= \sum_i \kappa^i(x^i) \left[(U^i)'(x^i) - (H Q e^i)^T \Phi'(H Q x) \right]^2 \end{aligned}$$

which is always greater than zero when x is not the Nash equilibrium. ■

We have shown two ways for the potential game to converge to the Nash equilibrium, one via best response dynamics and one via gradient ascent. For best response, only one player is modifying the strategy while for gradient ascent, all players are modifying the strategies. From the proofs, it should be clear that these are not the only two ways where the potential game could converge to the Nash equilibrium. We simply require dynamics where the potential is always increasing and where all players get a chance to update their strategies.

Now, we relax the timescale separation assumption and allow the traffic engineer to perform updates at any irregular intervals. We still assume that the update happens instantaneously.

Consider the potential game earlier with N players. We add in the traffic engineer as an additional player who tries to maximize its utility:

$$\max_{Q \in \mathcal{Q}} V(Q, x) \triangleq \sum_{l \in L} \Phi_l(e_l^T H Q x)$$

Lemma 6. The game with $N + 1$ players is an exact potential game with exact potential function

$$P(Q, x) = \sum_{i \in N} U^i(x^i) - \sum_{l \in L} \Phi_l(e_l^T H Q x)$$

Proof: We check by definition that

$$\frac{\partial V}{\partial q_j^i} = x^i (H e_j^i)^T \Phi'(H Q x) = \frac{\partial P}{\partial q_j^i}$$

■

Theorem 7. Suppose the N users update their strategies according to (24) while the traffic engineer sets congestion prices according to (25) and performs best response to update Q at any irregular intervals, then the potential game converges to the socially optimal solution.

Proof: We have shown that the potential is always increasing when the N users update according to gradient ascent. When the traffic engineer performs a best response update to Q , the potential increases instantaneously according to (21). Thus the potential is always increasing when the game is not at the Nash equilibrium. The Nash equilibrium is the fixed point of the Gauss-Seidel system (7), (8) and we have shown that it is socially optimal. ■

Remark. In fact, the traffic engineer only needs to find a new traffic split matrix Q that strictly increases $V(Q, x)$.

V. MODIFYING DUAL ALGORITHM

The dual algorithm shown in equations (3) and (4) have a capacity-dependent congestion price. Since we have relaxed the capacity constraint, we have to modify the price function accordingly. A direct attempt by taking the dual of the primal formulation does not work out but instead the modification will arise naturally from a game theory perspective by investigating the symmetry of equations (1), (2), and (3) of the primal and dual algorithms.

In the primal formulation, users are treated as players while the links are non-player entities providing the congestion price for link usage. The roles and dynamics are swapped for the dual formulation. Users are now non-player entities providing the offering traffic according to

$$x^i = (U^i)'^{-1} \left((H Q e^i)^T \phi \right) \quad (26)$$

Note that this is the best response for the users in the potential game. Each link l is now a player maximizing the utility

$$\max_{\phi_l \geq 0} \beta_l(\phi_l) \triangleq -\frac{1}{2} (\Phi'_l - \phi_l)^2$$

Note that the links are doing best response in the potential game.

We are now ready to prove an analogous result of Lemma 5 and Theorem 7 for the dual algorithm.

Lemma 8. If users set offered traffic according to (26), and links update their strategies according to gradient ascent:

$$\dot{\phi}_l = \Phi'_l - \phi_l \quad (27)$$

then the strategies converge to an optimal solution of (20).

Proof: From (26), we obtain the time derivative of

$$(U^i)''(x^i) \dot{x}^i = (HQe^i)^T \dot{\phi} \quad (28)$$

Using (26), (27) and (28), one can show that the potential P is always increasing as

$$\dot{P}(x) = \sum_i - \frac{[(U^i)'(x^i) - (HQe^i)^T \Phi']^2}{(U^i)''(x^i)} \geq 0$$

Theorem 9. *Suppose the N users update the offered traffic according to (26), and the traffic engineer sets the L links to update according to (27) and performs best response update to Q at any irregular intervals, then the process converges to the socially optimal solution.*

Proof: Analogous to Theorem 7. ■

We have assumed that all the users adopt the same algorithm in the last two sections. Our final contribution is to consider the case when there is a mix of users running primal and dual algorithms.

VI. HETEROGENEOUS USERS AND FEEDBACK

The result in this section is a straightforward combination of the results from the prior two sections. Suppose the users are divided into two non-intersecting subsets: a set N^P of users running primal algorithm and a set N^D of users running dual algorithm. The traffic engineer is assumed to know the type of algorithm each user is running. A different congestion price and thus feedback is calculated for each algorithm. For the primal algorithm, the congestion prices ϕ^P update according to (25) while for the dual, ϕ^D update according to (27).

Theorem 10. *If users update their offered traffic and links update the offered traffic and congestion prices according to*

$$\begin{aligned} \dot{x}^i &= \kappa^i(x^i) \left[(U^i)'(x^i) - (HQe^i)^T \phi^P \right], \quad i \in N^P \\ \dot{x}^i &= (U^i)^{\prime-1} \left((HQe^i)^T \phi^D \right), \quad i \in N^D \\ \dot{\phi}_l^P &= \Phi'_l, \quad \forall l \in L \\ \dot{\phi}_l^D &= \Phi'_l - \phi_l^D, \quad \forall l \in L \end{aligned}$$

while the traffic engineer performs best response update to Q at any irregular intervals, then the process converges to the socially optimal solution.

Proof: We show that the potential P is always increasing

$$\dot{P}(x) = \sum_{i \in N^P} \frac{\partial P(x)}{\partial x^i} \dot{x}^i + \sum_{i \in N^D} \frac{\partial P(x)}{\partial x^i} \dot{x}^i$$

As shown in the proofs of Lemma 5 and 8, the sum of the two terms are always positive except at the optimal solution of (20). The rest of the proof is analogous to Theorem 7. ■

VII. CONCLUSION

In this paper, we have shown that it is possible to induce social optimality by modifying the feedback value while maintaining all other aspects of the current interaction between traffic engineering and elastic traffic governed by (NUM). The optimality result holds even when the traffic engineer updates at any irregular intervals and the users are heterogeneous. The key to all the results is a simple idea: induce a user behavior that always increases the social value. It will be interesting to see if this idea is applicable for other cross-layer interactions in the layered network architecture.

REFERENCES

- [1] R. Gallager, "A minimum delay routing algorithm using distributed computation," *IEEE Trans. Commun.*, vol. 25, no. 1, pp. 73 – 85, Jan 1977.
- [2] B. Fortz and M. Thorup, "Increasing internet capacity using local search," *Computational Optimization and Applications*, vol. 29, pp. 13–48, 2004.
- [3] —, "Internet traffic engineering by optimizing OSPF weights," in *INFOCOM*, 2000.
- [4] N. Michael, A. Tang, and D. Xu, "Optimal link-state hop-by-hop routing," in *ICNP*, 2013.
- [5] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, no. 3, pp. pp. 237–252, 1998.
- [6] J. Wang, L. Li, S. Low, and J. Doyle, "Cross-layer optimization in TCP/IP networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 582 – 595, June 2005.
- [7] J. He, M. Bresler, M. Chiang, and J. Rexford, "Towards robust multi-layer traffic engineering: Optimization of congestion control and routing," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 5, pp. 868–880, June 2007.
- [8] F. Paganini and E. Mallada, "A unified approach to congestion control and node-based multipath routing," *IEEE/ACM Trans. Netw.*, vol. 17, no. 5, pp. 1413 –1426, Oct. 2009.
- [9] J. He, M. Chiang, and J. Rexford, "TCP/IP interaction based on congestion price: Stability and optimality," in *ICC*, June 2006.
- [10] M. Chiang, S. Low, A. Calderbank, and J. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proc. IEEE*, vol. 95, no. 1, pp. 255 –312, Jan. 2007.
- [11] B. Fortz, J. Rexford, and M. Thorup, "Traffic engineering with traditional ip routing protocols," *IEEE Commun. Mag.*, vol. 40, no. 10, pp. 118 – 124, Oct 2002.
- [12] C. B. W. Martin J. Beckmann, C. B. McGuire, *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [13] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.
- [14] D. Fudenberg and J. Tirole, *Game theory*. MIT Press, 1991.
- [15] R. W. Rosenthal, "A class of games possessing pure-strategy nash equilibria," *Int. J. Game Theory*, vol. 2, pp. 65–67, 1973.
- [16] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, Mar. 2002.
- [17] R. Cole, Y. Dodis, and T. Roughgarden, "Bottleneck links, variable demand and the tragedy of the commons," in *SODA*, 2006.
- [18] R. Johari and J. Tsitsiklis, "A scalable network resource allocation mechanism with bounded efficiency loss," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 5, pp. 992 – 999, May 2006.
- [19] J. Walrand, *An introduction to queueing networks*. Prentice Hall, 1988.
- [20] D. Bertsekas, *Nonlinear programming*. Athena Scientific, 1999.
- [21] R. Srikant, *The mathematics of Internet congestion control*. Birkhäuser, 2004.
- [22] S. Low and L. D.E., "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Trans. Netw.*, vol. 7, no. 6, pp. 861–874, 1999.
- [23] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.
- [24] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 556–567, Oct 2000.
- [25] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124 – 143, 1996.