Paradoxes of Traffic Engineering with Partially Optimal Routing

Daron Acemoglu, Ramesh Johari, Asuman Ozdaglar

***Abstract*— Most large-scale communication networks, such as the Internet, consist of interconnected administrative domains. While source (or selfish) routing, where transmission follows the least cost path for each source, may be reasonable across domains, within a subnetwork controlled by a single administrative domain, service providers often depart from selfish routing to improve operating performance. This observation suggests that the study of large-scale communication networks requires an analysis of *partially optimal routing*, where optimal routing within subnetworks is overlaid with selfish routing across domains. Such partially-optimal routing leads to a number of paradoxes, however. First, optimal routing within a subnetwork does not necessarily improve the performance of the overall network. In particular, we show that when partially optimal routing leads to worse overall network performance, it is because Braess' paradox can occur in the network. Second, in the absence of prices per unit of transmission, lower-layer traffic engineering within an administrative domain may prefer selfish routing to optimal routing because optimal routing typically increases the amount of flow into the subnetwork, thus increasing the delay costs.**

I. Introduction

Over the last decade, the Internet has undergone a dramatic transformation. Since the passage of the Telecommunications Act in 1996, ownership of the Internet has become increasingly decentralized. Today, thousands of network providers cooperate and compete to provide end-to-end network service to billions of users worldwide. There is an obvious tension that arises as a result; while individual network providers are led to optimize their own objectives, end users care about the performance observed across the entire network. The current Internet's architecture provides no guarantees that provider incentives will be aligned with end user objectives.

The emergence of *overlay routing* over the past five years has further highlighted the potentially conflicting objectives of the service provider and the end user. In overlay routing, end user software (such as peer-to-peer file-sharing software) makes route selection decisions on the basis of the best end-to-end performance available at any given time, but, at the same time, administrative domains control the routing of traffic within their own networks. They will typically wish to optimize performance in their network and may also react to the global routing decisions of overlay networks by changing the engineering of traffic patterns (e.g., [1]).

These considerations make it clear that the study of routing patterns and performance in large-scale communication networks requires an analysis of *partially optimal routing*, where end-to-end route selection is selfish (i.e., responds to aggregate route latency), but network providers redirect traffic within their own networks to achieve minimum intradomain total latency.

In this and a companion paper [2], we develop a model of partially optimal routing and investigate its implications for routing patterns and network performance. While recent research (e.g., [3], [4], [5]) has studied the interactions of overlay routing and traffic engineering, it has neither provided a model of partially optimal routing nor addressed the central question of whether partially optimal routing improves overall network performance.

We investigate these questions using a stylized model of traffic routing. All links in the network are characterized by a *latency function* describing the congestion level (e.g., delay) on the link as a function of the total flow passing through the link (e.g., [6]). Each source-destination pair in the network has a fixed amount of flow, and the standard assumption, captured in the notion of *Wardrop equilibrium*, is that each user chooses the minimum delay route among all the available routes. As is well known, because of the selfish routing behavior of users and the congestion externality thus created, the Wardrop equilibrium falls short of the optimal average latency of the network (e.g., [7], [8], [9]). We enrich this basic routing model by including service providers.

In our model, subsets of the links in the network are independently owned and operated by providers. To reflect the practice of traffic engineering by Internet service providers, we assume that each provider routes traffic within its subnetwork to minimize the average latency of that subnetwork. Source-destination pairs sending traffic across subnetworks perceive only the *effective latency* resulting from the traffic engineering (optimal routing) of the service provider. The resulting equilibrium, which we call a *partially optimal routing* (POR) equilibrium, is a Wardrop equilibrium according to the effective latencies seen by source-destination pairs.

We provide two basic sets of results on partially optimal routing. First, we investigate whether partially optimal routing (the presence of traffic engineering) improves overall network performance relative to pure selfish routing. We show that this issue is closely related to Braess' paradox: if partially optimal routing is less efficient than the Wardrop equilibrium, then Braess' paradox must occur in the network. Second,

This work was partially supported by research grants from the Okawa Foundation and the National Science Foundation.

Acemoglu: Department of Economics, Massachusetts Institute of Technology, Cambridge, MA, 02139 (e-mail: daron@mit.edu).

Johari: Department of Management Science and Engineering, Stanford University, Stanford, CA, 94305 (e-mail: ramesh.johari@stanford.edu).

Ozdaglar: Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 02139 (e-mail: asuman@mit.edu).

we investigate subnetwork performance under partially optimal routing, and determine conditions for a single network provider to prefer traffic engineering over selfish routing from the viewpoint of minimizing the total latency in its subnetwork. We show that, in the absence of prices per unit of transmission, providers may strategically choose selfish routing in order to reduce total flow into their subnetwork. For clarity of exposition, we focus on a network with single source and destination nodes, and subnetworks with unique entry and exit points. Our companion paper [2] extends our results to networks with multiple entry and exit points. It also provides bounds on the potential performance costs of partially optimal routing relative to optimal routing for the overall network and relative to pure selfish routing.

The remainder of the paper is organized as follows. Section II introduces the three basic routing paradigms: socially optimal routing, where total latency is minimized across the entire network; selfish routing, where end-to-end route selection is made based on minimum route latency; and partially optimal routing, where end-to-end route selection is still dependent on aggregate route latency, but providers also redirect traffic within their own networks to achieve minimum intradomain total latency. The latter model is an important contribution of our paper: it is the first extension of the classical traffic routing models to capture the phenomenon of traffic engineering.

Section III analyzes the performance of partially optimal routing. We show that there may exist situations where optimization within a subnetwork can actually lead to lower global network performance. We prove that if this situation occurs, then Braess' paradox must occur in the network. Section IV considers the choice of routing policy by a single service provider. We show that a provider may prefer *not* to optimize within its domain because this reduces total flow and therefore delay in its subnetwork, and we provide conditions under which traffic engineering (optimal routing) is optimal for a provider in parallel link topologies.

II. Preliminaries: Different Routing Paradigms

We consider a network $G = (V, A)$, with distinguished source and destination nodes $s, t \in V$, respectively. Let P denote the set of paths available from s to t using the edges in A; we view each path $p \in P$ as a subset of A, $p \subset A$. Each link $j \in A$ has a strictly increasing, nonnegative latency function $l_j(x_j)$ as a function of the flow on link j. We assume that X units of flow are to be routed from s to t. We call the tuple $R = (V, A, P, s, t, X, \mathbf{l})$ a *routing instance*.

A. Socially Optimal Routing

Given a routing instance $R = (V, A, P, s, t, X, \mathbf{l})$, we define the social optimum, denoted by $\mathbf{x}^{SO}(R)$, as the optimal solution of the following optimization problem:

$$\begin{aligned} \text{minimize} \quad & \sum_{j \in A} x_j l_j(x_j) \\ \text{subject to} \quad & \sum_{p \in P : j \in p} y_p = x_j, \quad j \in A; \\ & \sum_{p \in P} y_p = X; \\ & y_p \geq 0, \quad p \in P. \end{aligned}$$

Under our assumption that each latency function is strictly increasing, it is possible to show that the objective function above is strictly convex, so that $\mathbf{x}^{SO}(R)$ is uniquely defined. The total latency cost at the social optimum is given by:

$$C(\mathbf{x}^{SO}(R)) = \sum_{j \in A} x_j^{SO}(R) l_j(x_j^{SO}(R)).$$

B. Selfish Routing

When traffic routes "selfishly," all paths with nonzero flow must have the same total delay. A flow configuration with this property is called a *Wardrop equilibrium*. The Wardrop equilibrium flow for a given routing instance R, denoted $\mathbf{x}^{WE}(R)$, is the unique solution to the following optimization problem:

$$\begin{aligned} \text{minimize} \quad & \sum_{j \in A} \int_0^{x_j} l_j(z)\, dz \qquad (1) \\ \text{subject to} \quad & \sum_{p \in P : j \in p} y_p = x_j, \quad j \in A; \\ & \sum_{p \in P} y_p = X; \\ & y_p \geq 0, \quad p \in P. \end{aligned}$$

The total latency cost at the Wardrop equilibrium is given by

$$C(\mathbf{x}^{WE}(R)) = \sum_{j \in A} x_j^{WE}(R) l_j(x_j^{WE}(R)). \qquad (2)$$

It is well-known (e.g. Dafermos and Sparrow [10], Smith [11]) that a feasible solution $\mathbf{x}^{WE}$ of Problem (1) is a Wardrop equilibrium if and only if it satisfies

$$\sum_{j \in A} l_j(x_j^{WE})(x_j^{WE} - x_j) \leq 0, \qquad (3)$$

for all feasible solutions $\mathbf{x}$ of Problem (1).

C. Partially Optimal Routing

Next, we consider partially optimal routing in a subnetwork with unique entry and exit points. We assume there is a distinguished network inside of G, denoted $G_0 = (V_0, A_0)$, within which a network operator optimally routes all incoming traffic. Let $s_0 \in V_0$ denote the unique entry point to G_0, and let $t_0 \in V_0$ denote the unique exit point from G_0. Let P_0 denote the available paths from s_0 to t_0 using the edges in A_0. We make the assumption that every path in P passing through any node in V_0 must contain a path in P_0 from s_0 to t_0; this is consistent with our assumption that G_0 is an independent autonomous system, with a unique entry and exit point. We call $R_0 = (V_0, A_0, P_0, s_0, t_0)$ a *subnetwork* of G, and by abusing the notation, we say that $R_0 \subset R$.

Given an incoming amount of flow X_0, the network operator chooses a routing of flow to solve the following optimization problem:

$$\begin{aligned} \text{minimize} \quad & \sum_{j \in A_0} x_j l_j(x_j) \qquad (4) \\ \text{subject to} \quad & \sum_{p \in P_0 : j \in p} y_p = x_j, \quad j \in A_0; \\ & \sum_{p \in P_0} y_p = X_0; \\ & y_p \geq 0, \quad p \in P_0. \end{aligned}$$

Let $L(X_0)$ denote the optimal value of the preceding optimization problem; we define $l_0(X_0) = L(X_0)/X_0$ as the *effective latency* of partially optimal routing in the subnetwork R_0, with flow $X_0 > 0$. If traffic in the entire network G routes selfishly, while traffic is optimally routed within G_0, then replacing G_0 by a single link with latency l_0 will leave the Wardrop equilibrium flow unchanged elsewhere in G.

We have the following simple lemma.

Lemma 1

(a) The effective latency $l_0(X_0)$ is a strictly increasing function of $X_0 > 0$.
(b) The total cost $L(X_0)$ is a convex function of X_0.

Proof.

(a) Fix X_0, and $\hat{X}_0 < X_0$ with $\hat{X}_0 > 0$. Let $(\mathbf{x}, \mathbf{y})$ be an optimal solution to problem (4) with total flow X_0. Define $\hat{y}_p = y_p \hat{X}_0 / X_0$, for all $p \in P_0$. Then this is a feasible solution to problem (4) with total flow $\hat{X}_0$. Furthermore, the total latency at this solution is easily seen to be:

$$\sum_{p \in P_0} \hat{y}_p \sum_{j \in p} l_j \left(\sum_{q \in P_0 : j \in q} \hat{y}_q \right) = \sum_{j \in A_0} \frac{x_j \hat{X}_0}{X_0} l_j \left(\frac{x_j \hat{X}_0}{X_0} \right).$$

Now we observe that:

$$\begin{aligned} l_0(\hat{X}_0) &= \sum_{j \in A_0} \frac{x_j}{X_0} l_j \left(\frac{x_j \hat{X}_0}{X_0} \right) \\ &< \frac{1}{X_0} \sum_{j \in A_0} x_j l_j(x_j) = l_0(X_0), \end{aligned}$$

where the last inequality follows because l_j is strictly increasing and $\hat{X}_0 < X_0$.

(b) Note that $L(X_0)$ is the primal (or perturbation) function of optimization problem (4). Since the objective function is convex and the constraints are linear, the result follows using standard arguments from convex analysis (see [12]). ■

In light of the preceding lemma, we can extend the definition of l_0 so that $l_0(0) = \lim_{x_0 \downarrow 0} l_0(x_0)$; the preceding limit is well defined because l_0 is strictly increasing.

We define the overall network performance under partially optimal routing as follows. Given a routing instance $R = (V, A, P, s, t, X, \mathbf{l})$, and a subnetwork $R_0 = (V_0, A_0, P_0, s_0, t_0)$ defined as above, we define a new routing instance $R' = (V', A', P', s, t, X, \mathbf{l}')$ as follows:

$$\begin{aligned} V' &= (V \setminus V_0) \bigcup \{s_0, t_0\}; \\ A' &= (A \setminus A_0) \bigcup \{(s_0, t_0)\}; \end{aligned}$$

P' corresponds to all paths in P, where any subpath in P_0 is replaced by the link (s_0, t_0); and $\mathbf{l}'$ consists of latency functions l_j for all edges in $A \setminus A_0$, and latency l_0 for the edge (s_0, t_0). Thus R' is the routing instance R with the subgraph G_0 replaced by a single link with latency l_0; we call R' the *equivalent POR instance* for R with respect to R_0. The overall network flow in R with partially optimal routing in R_0, $\mathbf{x}^{POR}(R, R_0)$, is defined to be the Wardrop equilibrium flow in the routing instance R':

$$\mathbf{x}^{POR}(R, R_0) = \mathbf{x}^{WE}(R').$$

(Note that this leaves undefined the exact flow in the subnetwork R_0; this is to be expected, since problem (4) may not have a unique solution.) This corresponds to a model where traffic is routed selfishly, given the effective latency l_0 of the subnetwork R_0.

The total latency cost of the equivalent POR instance for R with respect to R_0 is given by

$$C(\mathbf{x}^{POR}(R, R_0)) = \sum_{j \in A'} x_j^{POR}(R, R_0) l_j(x_j^{POR}(R, R_0)).$$

III. Performance of Partially Optimal Routing

We first consider the effect of optimal routing within subnetworks on the performance of the network. Intuitively one would expect that doing optimal routing within some or all of the subnetworks would improve the overall performance. The following example shows that this need not always be the case.

Example 1 Consider the network $G = (V, A)$ with source and destination nodes $s, t \in V$ illustrated in Figure 1(a). Let $R = (V, A, P, s, t, 1, \mathbf{l})$ be the corresponding routing instance, i.e., one unit of flow is to be routed over this network. The subnetwork G_0 consists of the two parallel links in the middle, links 5 and 6, with latency functions

$$l_5(x_5) = 0.31, \qquad l_6(x_6) = 0.4\ x_6.$$

The latency functions for the remaining links in the network are given by

$$\begin{aligned} l_1(x_1) = x_1, &\qquad l_2(x_2) = 3.25, \\ l_3(x_3) = 1.25, &\qquad l_4(x_4) = 3x_4. \end{aligned}$$

Assume first that the flow through the subnetwork G_0 is routed selfishly, i.e., according to Wardrop equilibrium. Given a total flow X_0 through the subnetwork G_0, the effective Wardrop latency can be defined as

$$\tilde{l}_0(X_0) = \frac{1}{X_0}\ C(x^{WE}(R_0)), \qquad (5)$$

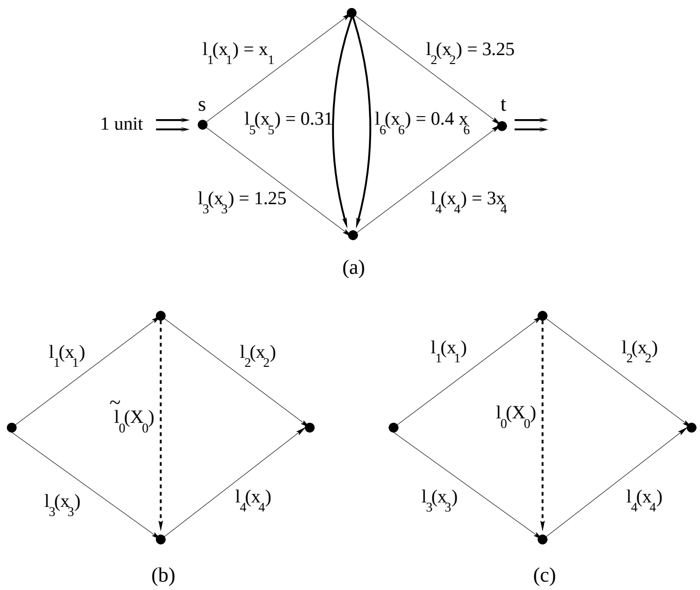

Fig. 1. A network for which POR leads to a worse performance relative to selfish routing. Figures (b) and (c) illustrate representing the subnetwork with a single link with Wardrop effective latency $\tilde{l}_0(X_0)$ and optimal effective latency $l_0(X_0)$, respectively.

[cf. Eq. (2)], where R_0 is the routing instance corresponding to the subnetwork G_0 with total flow X_0. The effective Wardrop latency for this example is given by

$$\tilde{l}_0(X_0) = \min\{0.31, 0.4X_0\}.$$

Substituting the subnetwork with a single link with latency function $\tilde{l}_0$ yields the network in Figure 1(b). It can be seen that selfish routing over the network of Figure 1(b) leads to the link flows $x_1^{WE} = 0.94$ and $X_0^{WE} = 0.92$, with a total cost of $C(\mathbf{x}^{WE}(R)) = 4.19$.

Assume next that the flow through the subnetwork G_0 is routed optimally, i.e., as the optimal solution of problem (4) for the routing instance corresponding to G_0. Given a total flow X_0 through the subnetwork G_0, the effective latency of optimal routing within the subnetwork G_0 can be defined as

$$l_0(X_0) = \frac{L_0(X_0)}{X_0},$$

where $L_0(X_0)$ is the optimal value of problem (4). The effective optimal routing latency for this example is given by

$$l_0(X_0) = \begin{cases} 0.4X_0, & \text{if } 0 \leq X_0 \leq 0.3875; \\ 0.31 - \frac{0.0961}{1.6X_0}, & \text{if } X_0 \geq 0.3875. \end{cases}$$

Substituting the subnetwork with a single link with latency function l_0 yields the network in Figure 1(c). Note that selfish routing over this network leads to the partially optimal routing (POR) equilibrium. It can be seen that at the POR equilibrium, the link flows are given by $x_1^{POR} = 1$ and $X_0^{POR} = 1$, with a total cost of $C(\mathbf{x}^{POR}(R)) = 4.25$, which is strictly greater than $C(\mathbf{x}^{WE}(R))$.

We next investigate the relationship between the performance of partially optimal routing and Braess' paradox. Informally, Braess' paradox occurs in a network if reducing the link latency functions actually degrades overall network performance, as measured by the total latency. We have the following definition.

Definition 1 (Braess' paradox) Consider a routing instance $R = (V, A, P, s, t, X, \mathbf{l})$. We say that *Braess' paradox* occurs in R if there exists another routing instance $R_m = (V, A, P, s, t, X, \mathbf{m})$, with a vector of strictly increasing, nonnegative latency functions, $\mathbf{m} = (m_j, j \in A)$, such that $m_j(x_j) \leq l_j(x_j)$ for all $x_j \geq 0$ and

$$C(\mathbf{x}^{WE}(R_m)) > C(\mathbf{x}^{WE}(R)).$$

The routing instance R' differs from R only by a reduction of the latency functions on some (or all) links. Nevertheless, in a network topology where Braess' paradox occurs, this can yield a higher total latency.

We also have the following alternative definition of a Braess' paradox.

Definition 2 (Braess' paradox centered at a subnetwork) Consider a routing instance $R = (V, A, P, s, t, X, \mathbf{l})$ and a subnetwork $R_0 = (V_0, A_0, P_0, s_0, t_0) \subset R$. We say that *Braess' paradox* occurs in R centered at R_0 if there exists another routing instance $R_m = (V, A, P, s, t, X, \mathbf{m})$, with a vector of strictly increasing, nonnegative latency functions, $\mathbf{m} = (m_j, j \in A)$, such that for all $x_j \geq 0$,

$$m_j(x_j) \leq l_j(x_j),\ \forall\ j \in A_0, \quad m_j(x_j) = l_j(x_j),\ \forall\ j \notin A_0,$$

and

$$C(\mathbf{x}^{WE}(R_m)) > C(\mathbf{x}^{WE}(R)).$$

The following definition captures the counterintuitive phenomenon that traffic engineering within some subnetwork, i.e., partially optimal routing, leads to a degradation in the overall performance compared to pure selfish routing.

Definition 3 (POR paradox) Consider a routing instance $R = (V, A, P, s, t, X, \mathbf{l})$, and a subnetwork $R_0 = (V_0, A_0, P_0, s_0, t_0)$. We say that the *POR paradox* (partially optimal routing paradox) occurs in R with respect to R_0 if

$$C(\mathbf{x}^{POR}(R, R_0)) > C(\mathbf{x}^{WE}(R)).$$

Intuitively, the POR paradox appears to be a form of "generalized Braess' paradox", in the following sense. Given a total flow X_0 routed through the subnetwork G_0, we define the effective Wardrop latency $\tilde{l}_0$, as follows:

$$\tilde{l}_0(X_0) = \frac{1}{X_0} \sum_{j \in A_0} x_j^{WE}(R') l_j(x_j^{WE}(R')) = \frac{C(\mathbf{x}^{WE}(R'))}{X_0}, \tag{6}$$

where $R' = (V_0, A_0, P_0, s_0, t_0, X_0, \mathbf{l})$ is a routing instance corresponding to the subnetwork R_0 with total flow X_0 [cf. Eq. (5)]. As in Lemma 1, it is straightforward to show that $\tilde{l}_0$ is strictly increasing. Furthermore, it is clear that $\tilde{l}_0(X_0) \geq l_0(X_0)$ for all $X_0 \geq 0$, since $\mathbf{x}^{WE}(R')$ is a feasible solution to problem (4). Thus when we contrast $\mathbf{x}^{POR}(R)$ and $\mathbf{x}^{WE}(R)$, it is *as if* we are lowering the effective latency of the subnetwork R_0. If this increases the total latency, then we are observing a form of Braess' paradox.

In fact, it is possible to show a stronger result: whenever the POR paradox occurs in R with respect to some $R_0 \subset R$, then Braess' paradox occurs in R centered at R_0.

Proposition 1 Consider a routing instance $R = (V, A, P, s, t, X, \mathbf{l})$ and a subnetwork $R_0 = (V_0, A_0, P_0, s_0, t_0) \subset R$. Assume that the POR paradox occurs in R with respect to R_0. Then Braess' paradox occurs in R centered at R_0.

Proof. Our approach will be to uniformly lower the latency functions in the subnetwork R_0, such that we exactly ensure at a Wardrop equilibrium the effective latency of R_0 is given by l_0, the effective latency of optimal routing within R_0. This will allow selfish routing to "replicate" the partially optimal routing of flow, and imply Braess' paradox.

Let $\mathbf{x}^{WE}(R)$ be the Wardrop equilibrium flow for the routing instance R, with corresponding path flows $\mathbf{y}^{WE}(R)$. Similarly, let $\mathbf{x}^{POR}(R, R_0)$ be the flow with partially optimal routing in R_0, with corresponding path flows $\mathbf{y}^{POR}(R, R_0)$. Let $X_0 = x_{s_0 t_0}^{POR}(R, R_0)$ represent the flow routed through the subnetwork R_0 under partially optimal routing. Note that $X_0 > 0$ since by assumption POR paradox occurs in R with respect to R_0. Let l_0 denote the effective latency of R_0 under partially optimal routing, and $\tilde{l}_0$ denote the effective latency of R_0 under selfish routing [cf. Eq. (6)].

Define a routing instance $R_0' = (V_0, A_0, P_0, s_0, t_0, X_0, \mathbf{l})$ and let $\mathbf{x}^{WE}(R_0')$ be the Wardrop equilibrium flow for the routing instance R_0'.

We define a new collection of latency functions as follows. For all $j \notin A_0$, define $m_j = l_j$. For $j \in A_0$, we choose a new strictly increasing, nonnegative latency function m_j with $m_j(x_j) \leq l_j(x_j)$ for all $x_j \geq 0$, such that

$$m_j(x_j^{WE}(R_0')) = \frac{l_0(X_0)}{\tilde{l}_0(X_0)} l_j(x_j^{WE}(R_0')).$$

Observe that such a choice is possible, since $l_0(X_0) \leq \tilde{l}_0(X_0)$ implies that $m_j(x_j^{WE}(R_0')) \leq l_j(x_j^{WE}(R_0'))$.

Let $T_0 = (V_0, A_0, P_0, s_0, t_0, X_0, \mathbf{m})$; i.e., T_0 is the routing instance R_0' with latencies replaced by $\mathbf{m}$. We claim that $\mathbf{x}^{WE}(T_0) = \mathbf{x}^{WE}(R_0')$. This follows from the definition of $\mathbf{m}$: all values $m_j(x_j^{WE}(R_0'))$ are proportional to $l_j(x_j^{WE}(R_0'))$, with common constant of proportionality $l_0(X_0)/\tilde{l}_0(X_0)$. Thus if $\mathbf{x}^{WE}(R_0')$ is the Wardrop equilibrium flow with latencies $\mathbf{l}$, it must remain so with latencies $\mathbf{m}$. Furthermore, observe that for any path p with positive flow, we have

$$\sum_{j \in p} m_j(x_j^{WE}(T_0)) = \frac{l_0(X_0)}{\tilde{l}_0(X_0)} \sum_{j \in p} l_j(x_j^{WE}(R_0')) = l_0(X_0),$$

because the second summation above is equal to $\tilde{l}_0(X_0)$. Thus we conclude

$$C(\mathbf{x}^{WE}(T_0)) = \sum_{j \in A_0} x_j^{WE}(T_0) m_j(x_j^{WE}(T_0)) = X_0 l_0(X_0). \tag{7}$$

Let $T = (V, A, P, s, t, X, \mathbf{m})$. Define a feasible flow $\mathbf{x} = [x_j]_j \in A$ as follows:

$$x_j = \begin{cases} x_j^{POR}(R, R_0), & \text{if } j \notin A_0; \\ x_j^{WE}(R_0'), & \text{if } j \in A_0. \end{cases}$$

We claim that $\mathbf{x}^{WE}(T) = \mathbf{x}$. This claim follows easily since we have already established that $\mathbf{x}^{WE}(T_0) = \mathbf{x}^{WE}(R_0')$, and (7) holds. In the flow $\mathbf{x}$ for the routing instance T, the effective latency perceived by any flow crossing the subnetwork R_0 is exactly equal to the partially optimal routing effective latency $l_0(X_0)$ (by (7)). But then since all routing outside the subnetwork R_0 is performed according to $\mathbf{x}^{POR}(R, R_0)$, we conclude that in fact $\mathbf{x}^{WE}(T) = \mathbf{x}$, as required.

Combining the preceding, we obtain

$$\begin{aligned}
&\sum_{j \in A} x_j^{WE}(T) m_j(x_j^{WE}(T)) \\
&\quad = \sum_{j \notin A_0} x_j^{POR}(R, R_0) l_j(x_j^{POR}(R, R_0)) \\
&\qquad + \sum_{j \in A_0} x_j^{WE}(R_0') m_j(x_j^{WE}(R_0')) \\
&\quad = \sum_{j \notin A_0} x_j^{POR}(R, R_0) l_j(x_j^{POR}(R, R_0)) + X_0 l_0(X_0) \\
&\quad = \sum_{j \in A'} x_j^{POR}(R, R_0) l_j(x_j^{POR}(R, R_0)) \\
&\quad = C(\mathbf{x}^{POR}(R, R_0)).
\end{aligned}$$

Since we assumed that the POR paradox occurs in R with respect to R_0, we obtain from the preceding that

$$C(\mathbf{x}^{WE}(T)) = C(\mathbf{x}^{POR}(R, R_0)) > C(\mathbf{x}^{WE}(R)),$$

implying that Braess' paradox occurs in R centered at R_0. ■

An immediate corollary of the preceding proposition is the following: Given a routing instance R, if Braess' paradox does not occur in R, then partially optimal routing with respect to any subnetwork always improves the network performance. Milchtaich has shown that Braess' paradox does not occur in directed graphs where the underlying undirected graph has a *series-parallel* structure [13]. This implies that as long as the network under consideration has a series-parallel structure (for example, a network of parallel links), partially optimal routing always improves the overall network performance.

IV. Subnetwork Performance

In this section, we consider a model where a subnetwork can choose a routing policy to achieve the minimum (total) latency within its subnetwork (implicitly assuming there are no prices per unit of transmission and the subnetwork ignores revenues from transmission). While optimal routing seems like the natural means to achieve this goal, end-to-end route selection may counteract any expected performance gains from this type of intradomain traffic engineering. As a result, the provider may prefer to allow traffic to route selfishly in order to reduce flow to and total delay in its subnetwork. The following example illustrates this scenario.

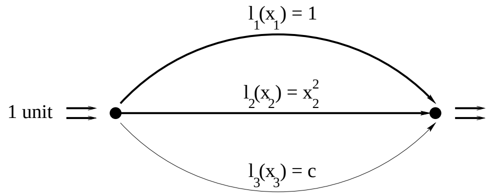

Fig. 2. A parallel link network. Links 1 and 2 form a subnetwork that is controlled by an independent administrator.

Example 2 Consider the parallel-link network illustrated in Figure 2. The latency functions are given by

$$l_1(x_1) = 1, \qquad l_2(x_2) = x_2^2, \qquad l_3(x_3) = c,$$

for some constant $c > 0$. Assume that links 1 and 2 form a subnetwork, denoted by G_0, which is controlled by an independent administrator. Assume that one unit of flow is to be routed over this network.

Assume first that the flow through the subnetwork G_0 is routed selfishly, i.e., according to Wardrop equilibrium. It can be seen in this case that $\sqrt{c}$ units of traffic is routed through the subnetwork, leading to a total cost of $C(\mathbf{x}^{WE}) = c$, and a subnetwork cost of $C_{G_0}(\mathbf{x}^{WE}) = c\sqrt{c}$.

Assume next that the flow through the subnetwork G_0 is routed optimally, i.e., the flow is routed through the overall network according to POR equilibrium. Assume that the constant $c \in \left[1 - \frac{2}{3\sqrt{3}}, 1\right]$. It can be seen in this case that the entire traffic is routed through the subnetwork, leading to a total and subnetwork cost of $C(\mathbf{x}^{POR}) = C_{G_0}(\mathbf{x}^{POR}) = 1 - \frac{2}{3\sqrt{3}}$. Note that for $c\sqrt{c} < 1 - \frac{2}{3\sqrt{3}}$, we have

$$C_{G_0}(\mathbf{x}^{POR}) > C_{G_0}(\mathbf{x}^{WE}).$$

As the preceding example demonstrates, lower-layer traffic engineering may prefer selfish to optimal routing. It is equally easy to construct examples where optimal routing will be preferred. The simplest example is a situation in which the total traffic entering the subnetwork is constant, regardless of whether selfish or optimal routing is used. This will be the case in the example above when $c > 1$, and a similar analysis immediately implies that optimal routing will be preferred within the subnetwork in this case.

To gain more insights, let us next consider a "partial equilibrium" analysis of routing within a subnetwork, taking the strategies of all other subnetworks as given. To illustrate the main issues, we consider a network consisting of parallel links between a single origin-destination pair with d units of total traffic. Suppose that there are $N+1$ providers and each network provider owns a subset of the links in the network. We represent network provider i, for $i = 1, \ldots, N$, by a single link with effective latency l_i (corresponding to the intradomain routing policy chosen by provider i, whether optimal routing or not). We assume all these latency functions l_i are continuous and strictly increasing.

As in the preceding discussion, we assume that if provider 0 pursues an optimal intradomain routing policy, then the effective latency is given by l_0, and if provider 0 allows purely selfish routing within his network (corresponding to a Wardrop equilibrium), the effective latency is $\tilde{l}_0$. To simplify the discussion here, let us also assume that l_0 and $\tilde{l}_0$ are both continuous and strictly increasing (this will be the case, for example, when all latency functions of the links in the subnetwork are continuous and strictly increasing). As before, recall that $\tilde{l}_0(x) \geq l_0(x)$ for all $x \geq 0$. Moreover, for simplicity, let us assume that $\tilde{l}_0(x) > l_0(x)$ if $x > 0$ (though the arguments can be generalized to the case without this assumption).

We assume that the subnetwork owner can randomize between the two policies, so any convex combination of optimal and selfish routing can be achieved. In other words, the subnetwork owner chooses a $\delta \in [0,1]$ corresponding to an effective latency given by:

$$m_0(x,\delta) = (1-\delta)\, l_0(x) + \delta \tilde{l}_0(x),$$

where $\delta = 0$ corresponds to optimal routing, while $\delta = 1$ corresponds to selfish routing.

We continue to use $\mathbf{x}^{POR}$ to denote a Wardrop equilibrium with respect to the latency functions $m_0, l_1, \ldots, l_N$, so that, $\mathbf{x}^{POR}$ satisfies:

$$\begin{aligned}
m_0(x_0^{POR}) &\geq \lambda; \\
l_i\left(x_i^{POR}\right) &\geq \lambda \text{ for } i = 1,...,N; \\
\sum_{i=0}^{N} x_i^{POR} &= d\ ; \\
x_i^{POR} &\geq 0 \text{ for } i = 0,...,N; \\
\lambda &= \min\left\{m_0\left(x_0^{POR}\right), l_1\left(x_1^{POR}\right), \ldots, l_N(x_N^{POR})\right\}.
\end{aligned}$$

First consider the routing of flow through the links $1, \ldots, N$. If a total flow x is routed through links $1, \ldots, N$, then the resulting flow allocation must satisfy:

$$l_i(x_i) = \min\{l_1(x_1), \ldots, l_N(x_N)\} \text{ if } x_i > 0; \tag{8}$$

$$\sum_{i=1}^{N} x_i = x; \tag{9}$$

$$x_i \geq 0, \ i = 1, \ldots, N. \tag{10}$$

In view of the assumption that $l_1, \ldots, l_N$ are strictly increasing, the preceding equations have a unique solution. We define $l_R(x)$ as the latency at this solution, i.e.,

$$l_R(x) = \min\{l_1(x_1), \ldots, l_N(x_N)\},$$

where $(x_1, \ldots, x_N)$ is the unique solution to (8)-(10). Since each l_i is strictly increasing and continuous, the function l_R is also strictly increasing and continuous.

Next consider the traffic engineering problem faced by subnetwork 0. The network provider will choose a value of δ that minimizes the total latency inside the subnetwork, given that traffic will follow the Wardrop equilibrium pattern for the resulting effective latencies. Formally, the optimization problem of subnetwork 0 is the following:

$$\min_{0 \leq x_0 \leq d, \delta \in [0,1]} \left((1-\delta)\, l_0(x_0) + \delta \tilde{l}_0(x_0)\right) x_0 \tag{11}$$

subject to

$$\begin{aligned}(1-\delta)\, l_0(0) + \delta \tilde{l}_0(0) &\geq l_R(d), \quad \text{if } x_0 = 0;\\ (1-\delta)\, l_0(d) + \delta \tilde{l}_0(d) &\leq l_R(0), \quad \text{if } x_0 = d;\\ (1-\delta)\, l_0(x_0) + \delta \tilde{l}_0(x_0) &= l_R(d - x_0), \quad \text{if } 0 < x_0 < d.\end{aligned}$$

Since $\tilde{l}_0(x_0) \geq l_0(x_0)$ for all $x_0 \geq 0$, and l_R is strictly increasing, as δ increases from $\delta = 0$ (purely optimal routing) to $\delta = 1$ (purely selfish routing), the flow routed through subnetwork 0 at the POR equilibrium must be nonincreasing.

Next note that when $\tilde{l}_0(0) \geq l_R(d)$, the subnetwork can achieve the minimum total latency of zero by choosing $\delta = 1$ (since the POR equilibrium will route no traffic across subnetwork 0). Similarly, if $l_0(d) \leq l_R(0)$, then regardless of provider 0's policy, all the flow will be routed across subnetwork 0. As a result, in this scenario the optimal strategy is $\delta = 0$ (optimal routing), as this minimizes the total latency. For the remainder of this section, we assume that $\tilde{l}_0(0) < l_R(d)$ and $l_0(d) > l_R(0)$. Since $\tilde{l}_0 \geq l_0$, this also implies

$$l_R(0) < l_0(d) \leq \tilde{l}_0(d); \qquad l_0(0) \leq \tilde{l}_0(0) < l_R(d). \tag{12}$$

The preceding conditions, together with the fact that l_0, $\tilde{l}_0$, and l_R are strictly increasing and continuous, ensure that the following two equations have unique solutions, x_0^{MIN} and x_0^{MAX}:

$$\tilde{l}_0(x_0^{MIN}) = l_R(d - x_0^{MIN}); \qquad l_0(x_0^{MAX}) = l_R(d - x_0^{MAX}).$$

Moreover, given our assumptions, x_0^{MIN} is the minimum flow that can go through subnetwork 0 (achieved exactly when $\delta = 1$, i.e., at purely selfish routing), and x_0^{MAX} is the maximum flow for subnetwork 0 (achieved exactly when $\delta = 0$, i.e., at optimal routing). Furthermore, any flow $x_0 \in [x_0^{MIN}, x_0^{MAX}]$ is achievable, by choosing δ such that:

$$\delta = \frac{l_R(d - x_0) - l_0(x_0)}{\tilde{l}_0(x_0) - l_0(x_0)},$$

where $0 \leq \delta \leq 1$ since conditions (12) hold and $\tilde{l}_0(x_0) > l_0(x_0)$ for all $x_0 > 0$. Finally, observe that for all $x_0 \in [x_0^{MIN}, x_0^{MAX}]$, if x_0 results as the partially optimal routing flow through subnetwork 0, we must have the relation $m_0(x_0) = l_R(d - x_0)$ (since $0 < x_0 < d$).

As a result, the optimization problem for the owner of subnetwork 0 becomes:

$$\min_{x_0 \in [x_0^{MIN}, x_0^{MAX}]} x_0 l_R(d - x_0), \tag{13}$$

the solution of which determines the delay-minimizing routing policy of the subnetwork. Since the game between service providers is one of complete information, all the latency functions are common knowledge and the owner of subnetwork can compute x_0^{MIN}, x_0^{MAX}, and l_R, and hence the optimal flow through the subnetwork. Clearly, if $x_0 l_R(d - x_0)$ increases as x_0 increases in the neighborhood of x_0^{MIN}, the provider will (locally) prefer selfish routing. Similarly, if $x_0 l_R(d - x_0)$ decreases as x_0 decreases in the neighborhood of x_0^{MAX}, the provider prefers selfish routing.

This analysis shows that with complete information and a parallel-link network, the delay-minimizing policy of the network is straightforward to characterize. We leave several issues for future work: the analysis of networks with more general topologies; general equilibrium structures where all subnetworks optimize; situations in which latency functions of other providers are unknown; and scenarios in which the objective of subnetworks may be profit maximization rather than delay minimization.

References

[1] J. Rexford, *Route optimization in IP networks*. Springer Science and Business Media, 2006, pp. 679–700.

[2] D. Acemoglu, R. Johari, and A. Ozdaglar, "Partially optimal routing," 2006, under preparation.

[3] L. Qiu, Y. Yang, Y. Zhang, and S. Shenker, "On selfish routing in internet-like environments," in *Proceedings of ACM SIGCOMM*, 2003, pp. 151–162.

[4] H. Zhang, Y. Liu, W. Gong, and D. Towsley, "Understanding the interaction between overlay routing and traffic engineering," University of Massachusetts CMPSCI, Technical Report 04-63, 2004.

[5] ——, "On the interaction between overlay routing and underlay routing," in *Proceedings of IEEE Infocom*, 2005, pp. 2543–2553.

[6] M. Beckmann, C. B. McGuire, and C. B. Winsten, *Studies in the Economics of Transportation*. New Haven, Connecticut: Yale University Press, 1956.

[7] A. Pigou, *The Economics of Welfare*. London, U.K.: Macmillan, 1920.

[8] T. Roughgarden and É.. Tardos, "How bad is selfish routing?" *Journal of the ACM*, vol. 49, no. 2, pp. 236–259, 2002.

[9] J. Correa, A. Schulz, and N. Stier-Moses, "Selfish routing in capacitated networks," *Mathematics of Operations Research*, vol. 29, no. 4, pp. 961–976, 2002.

[10] S. Dafermos and F. Sparrow, "The traffic assignment problem for a general network," *Journal of Research of the National Bureau of Standards-B. Mathematical Sciences*, vol. 73, no. 2, pp. 91–118, 1969.

[11] M. Smith, "The existence, uniqueness and stability of traffic equilibria," *Transportation Research*, vol. 13B, pp. 295–304, 1979.

[12] D. Bertsekas, A. Nedic, and A. Ozdaglar, *Convex Analysis and Optimization*. Cambridge, Massachusetts: Athena Scientific, 2003.

[13] I. Milchtaich, "Network topology and the efficiency of equilibrium," Department of Economics, Bar-Ilan University, Working Paper 12-01, 2005.