Congestion Resulting from Increased Capacity in Single-Server Queueing Networks

#252

Joel E. Cohen and Clark Jeffries, *Member, IEEE*

***Abstract*— In many networks for traffic and telecommunications, minimizing delays from entry to exit is a major concern of users. In user-optimal routing, each user chooses a path to minimize delay from entry to exit, given the existing paths chosen by all other users. Under user-optimal routing, at equilibrium all users experience the same delay. Many networks, especially data networks, are commonly modeled as networks of single-server queues. We report examples of single-server queueing networks with user-optimal routing in which adding servers or increasing the capacity of existing servers worsens the delay experienced by all users.**

I. Introduction

The optimal design and efficient use of traffic networks (including networks for voice communication and computer data) are important aspects of the infrastructure of economically complex societies [1]–[4]. Until recently, it was generally believed that increasing the capacity of links in traffic networks or adding links to networks could not worsen and would probably improve network performance. This plausible belief is now known to be false in several kinds of congested networks when traffic is routed by each individual participant to minimize its own delay, given the paths chosen by all other participants [3]–[10]. In such networks, additional capacity to process traffic, together with user-optimal routing, can—apparently paradoxically—worsen transit time through the network for all traffic. Similar apparent paradoxes can also arise in computer networks; for example, introducing a high-speed link in a serial chain of processors may increase the average transfer time from input to output [11].

Examples of such apparent paradoxes in networks consisting entirely of single-server queues are reported here. These examples were developed independently of a similar example recently reported in [12]. Such paradoxes are important practically because networks of single-server queues are widely used to model the network level of data networks [1]–[3]. Single-server queuing models at the network level focus on long-term average performance, summarizing the complexities of transient congestion through the service time distribution [1]–[2]. Transient congestion is also of practical importance at the physical and the data-link-control layers in data networks like the Internet [11].

Manuscript received January 17, 1996; revised July 14, 1996; approved by IEEE/ACM Transactions on Networking Editor U. Shankar. The work of J. E. Cohen was supported by the U.S. National Science Foundation under Grant BSR92-07293.

J. E. Cohen is with the Laboratory of Populations, Rockefeller University, New York, NY 10021-6399 USA, and the Earth Institute and Department of International and Public Affairs, Columbia University, New York, NY 10027 USA (e-mail: cohen@rockvax.rockefeller.edu.).

C. Jeffries is with the Laboratory of Populations, Rockefeller University, New York, NY 10021-6399 USA, on leave from the Department of Mathematical Sciences, Clemson University, Clemson SC 29634-1907 USA (e-mail: jclark@clemson.edu).

Publisher Item Identifier S 1063-6692(97)02775-1.

II. Single-Server Queues

An infinite-buffer single-server queue with Poisson message-packet arrivals takes an exponentially distributed amount of time (independently for all packets) to service each packet. Packets are processed in first-in, first-out order. (This idealization ignores packet loss and retries.) The capacity (number of packets processed per unit time) is the reciprocal of the average time the server requires to process one packet, and the flow is the average number of packets that arrive at the server per unit of time. The average delay experienced by a packet arriving as part of a Poisson stream at a single-server queue equals 1/(capacity − flow), always assuming capacity > flow. Thus, delay becomes arbitrarily great as flow in a link approaches capacity [1, pp. 385–390].

A network consists of nodes and links (directed edges, here modeled as single-server queues). Some pairs of distinct nodes are in–out pairs, meaning that traffic enters at the first node and leaves at the second at a constant flow rate F. For a given in–out pair, traffic can choose among paths $p_\alpha, \ldots, p_\beta$, so the flow F can be divided as $F = f_\alpha + \cdots + f_\beta$ where the flow in path p_α is some nonnegative number f_α. Link number i has a capacity C_i. When flows $f_\alpha, \ldots, f_\beta$ are all using link i, the delay d_i of each flow through that link is $d_i = 1/(C_i - f_\alpha - \cdots - f_\beta)$. Delay is defined and positive, even in the absence of flows, namely, $d_i = 1/C_i$. Under user-optimal routing, users seek the route that minimizes their average delay, given the routes already chosen by all other users. At equilibrium, all paths in use have the same delay, and that delay is less than that of any paths not in use. The total cost T of given flows in the network is defined to be the sum of all the link costs, where the cost in link i is $(f_\alpha + \cdots + f_\beta)d_i$.

III. Examples

Consider the networks shown in Fig. 1.

In the left network in Fig. 1, there are two flows from the input node to the output node. We label as f_1 the left flow using links with capacities C_1 and C_4, and as f_2 the right flow using links with capacities C_2 and C_5. Suppose $f_1 + f_2 = 1$, so

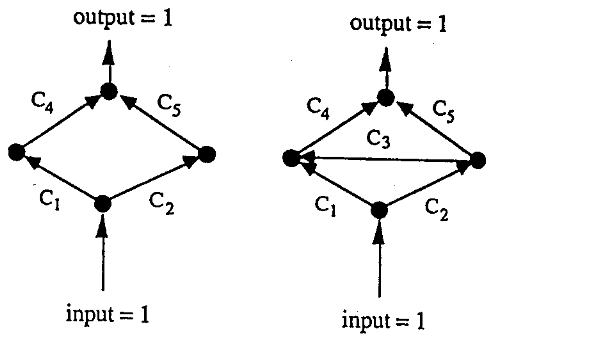

Fig. 1. Networks of single-server queues before and after the addition of a link with capacity C_3.

that the total cost equals the weighted total of all delays. With arbitrary positive capacities, delays are equal ($D_1 = D_2$) if

$$1/(C_1 - f_1) + 1/(C_4 - f_1) - 1/(C_2 - f_2) - 1/(C_5 - f_2) = 0$$

or, using $f_2 = 1 - f_1$,

$$1/(C_1 - f_1) + 1/(C_4 - f_1) - 1/(C_2 + f_1 - 1) - 1/(C_5 + f_1 - 1) = 0.$$

Note that this function of f_1 has a strictly increasing derivative, so any zero is the only zero.

The right network in Fig. 1 has a third flow, the zigzag flow f_3 using links with capacities C_2, C_3, and C_4. The right network satisfies the flow constraint $f_1 + f_2 + f_3 = 1$. The delays for the three flows are

$$D_1 = 1/(C_1 - f_1) + 1/(C_4 - f_1 - f_3) \quad \text{(left flow)}$$
$$D_2 = 1/(C_2 - f_2 - f_3) + 1/(C_5 - f_2) \quad \text{(right flow)}$$
$$D_3 = 1/(C_2 - f_2 - f_3) + 1/(C_3 - f_3) + 1/(C_4 - f_1 - f_3) \quad \text{(zigzag flow)}$$

and the total cost is

$$T = f_1/(C_1 - f_1) + (f_2 + f_3)/(C_2 - f_2 - f_3) + f_3/(C_3 - f_3) + (f_1 + f_3)/(C_4 - f_1 - f_3) + f_2/(C_5 - f_2).$$

Suppose the capacities are $C_1 = C_5 = 1, C_2 = C_4 = 3$, and C_3 varies from 0 (absence of the link) to infinity.

Equal delay traffic in the left network results in flows $f_1 = f_2 = 1/2$ and delays $D_1 = D_2 = 2.4$. Note that the total delay is $f_1 D_1 + f_2 D_2 = 2.4$. Even with no traffic in the zigzag path, the delay of that path in the right network is defined and is $1/(3-1/2)+1/(C_3-0)+1/(3-1/2) = 0.8+1/C_3$. Thus, no traffic will choose the zigzag path if $0.8+1/C_3 > 2.4$, that is, unless the capacity of the new link C_3 is at least 0.625. This is reflected in Fig. 2.

Suppose $C_3 = 1/2$, so no traffic would choose the zigzag path with user-optimal flows. Should flows be imposed as, say, $f_1 = f_2 = 0.49$ (and so $f_3 = 0.02$), then delays in the side paths are reduced to $D_1 = D_2 = 2.3624$. However, the zigzag delay would be 2.8865, contradicting the assumption of equal delays. This imposed flow would be an improvement in total delay because $f_1 D_1 + f_2 D_2 + f_3 D_3 = 2.3729 < 2.4$,

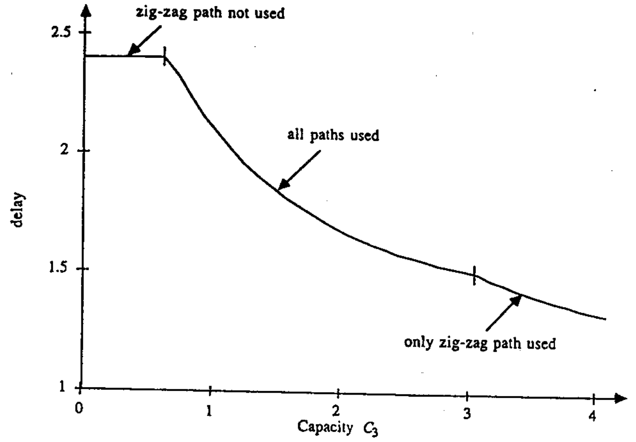

Fig. 2. Delay decreases as the capacity C_3 of the middle link in Fig. 1 increases, as intuition would suggest. Fixed capacities are $C_1 = C_5 = 1$ and $C_2 = C_4 = 3$. This graph (and Figs. 4 and 6) were generated from equilibrium values attained by a difference equation method described in Section VI.

but, again, an external factor would be needed to compel some traffic to use the zigzag path with longer delay.

More generally, for all single-server queuing networks connected as shown in Fig. 1 with $C_1 = C_5, C_2 = C_4$, the zigzag path never leads to increased delay, and in fact, reduces delay if used at all, as we now prove. Let $f_1 = f_2 = x, f_3 = 1 - 2x, 0 \le x \le 1/2$, so each flow and delay is a function of x. Traffic will shift to the new zigzag path precisely if $D_1(1/2) = D_2(1/2) > D_3(1/2)$ or, equivalently, if $1/(C_1-1/2) > 1/(C_2-1/2)+1/C_3$; this occurs if and only if $C_2 > C_1$ and C_3 is sufficiently large. To prove that such use of the zigzag path decreases delay for fixed $C_2 > C_1$, let us consider case (A) in which all three paths are used, and case (B) in which only the zigzag path is used. In case A, $D_1(x) = D_2(x) = D_3(x)$ for some $0 < x < 1/2$. Suppose, in addition, that $D_1(1/2) = D_2(1/2) < D_3(x)$ (by which the zigzag path would make delay worse). Thus, $D_1(x) = D_2(x) = D_3(x)$ and $D_1(1/2) = D_2(1/2) < D_3(x)$ imply

$$\frac{1}{C_1 - 1/2} + \frac{1}{C_2 - 1/2} < \frac{2}{C_2 - 1 + x} + \frac{1}{C_3 - 1 + 2x} = \frac{1}{C_1 - x} + \frac{1}{C_2 - 1 + x} \quad (0 < x < 1/2). \qquad \text{(A)}$$

Selecting terms from (A) yields

$$\frac{1}{C_1 - 1/2} + \frac{1}{C_2 - 1/2} < \frac{1}{C_1 - x} + \frac{1}{C_2 - 1 + x}$$

which can be rewritten as

$$\frac{1}{C_1 - 1/2} - \frac{1}{C_1 - x} < -\frac{1}{C_2 - 1/2} + \frac{1}{C_2 - 1 + x}.$$

Adding fractions then yields

$$\frac{1/2 - x}{(C_1 - 1/2)(C_1 - x)} < \frac{1/2 - x}{(C_2 - 1 + x)(C_2 - 1/2)}.$$

Since $1/2 > x$, it follows that

$$(C_1 - 1/2)(C_1 - x) > (C_2 - 1 + x)(C_2 - 1/2).$$

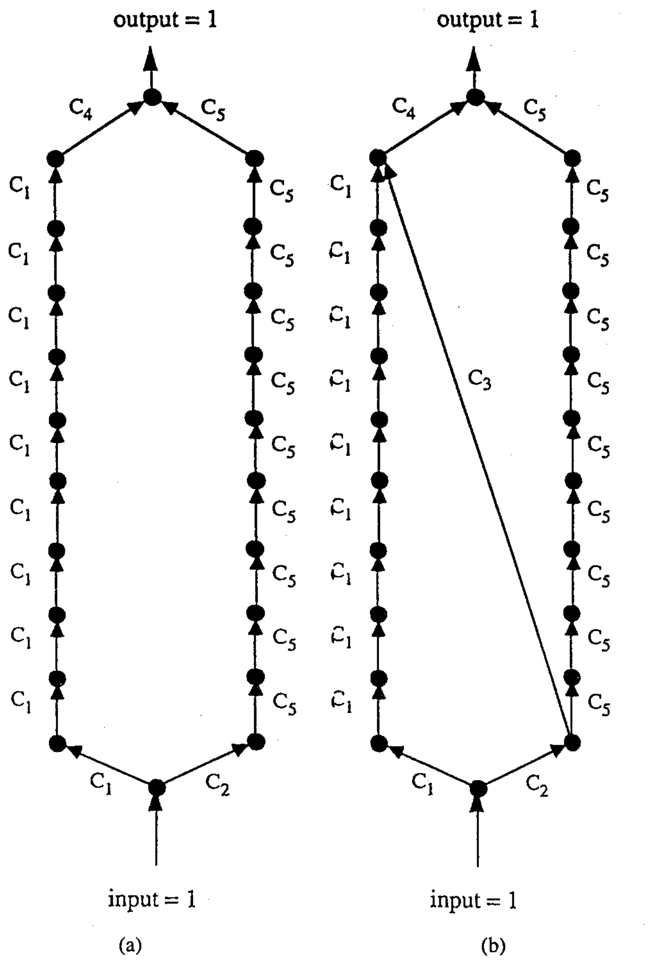

Fig. 3. For certain flow capacity values, the right network with the shortcut has longer delay than the left network.

Selecting terms in (A) also yields

$$\frac{1}{C_1 - x} > \frac{1}{C_2 - 1 + x}.$$

Thus $C_2 - 1 + x > C_1 - x$. The inequalities $(C_1 - 1/2)(C_1 - x) > (C_2 - 1 + x)(C_2 - 1/2)$ and $C_2 - 1 + x > C_1 - x$ together imply

$$(C_1 - 1/2)(C_1 - x) > (C_1 - x)(C_2 - 1/2)$$

Thus $C_1 - 1/2 > C_2 - 1/2$, but this contradicts $C_2 > C_1$ from the assumption that traffic initially shifts to the new link.

For case B, $D_1(0) = D_2(0) > D_3(0)$. Suppose, in addition, that $D_1(1/2) = D_2(1/2) < D_3(0)$ (by which the zigzag path would make delay worse). Substituting $x = 0$ and otherwise using the same line of reasoning leads again to $C_1 > C_2$, a contradiction.

We do not know whether the symmetries $C_1 = C_5, C_2 = C_4$ are crucial for decreasing delays in response to increasing C_3. However, in certain networks with the same topology as Fig. 1, but with links not responding as single-server queues, delay does not always decrease as C_3 increases [5], [9].

IV. Paradoxical Flows

Contrary to some intuition, including ours, user-optimal delay need not decline with increasing capacity in larger networks of single-server queues. For example, consider the networks in Fig. 3.

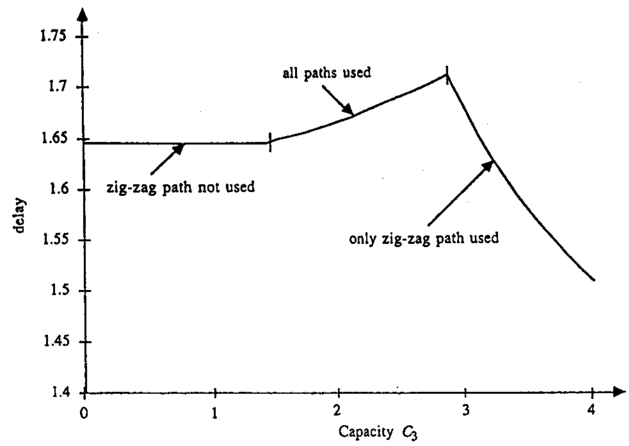

Fig. 4. The effect of C_3 values on delay in the network in Fig. 3.

Suppose the capacities in Fig. 3 are $C_1 = C_5 = 8.9$ and $C_2 = C_4 = 2.7$, while C_3 varies from 0 to infinity. As in the previous section, there are three flows: left (f_1), right (f_2), and zigzag (f_3). If, initially, the link with capacity C_3 is absent (the left network), then the flow with $f_1 = f_2 = 1/2$ is both the flow with equal delays $D_1 = D_2 = 1.65$ (decimal values are approximate) and the flow with optimal cost $T = 1.65$. Under user-optimal routing, the addition of the link with capacity C_3 results in delay as shown in Fig. 4.

No traffic uses the zigzag path for $0 \leq C_3 \leq 1.36$. For $1.36 \leq C_3 \leq 2.87$, all three paths are used. For $C_3 > 2.87$, all traffic uses the zigzag path, and the delay is $D_3 = 1.18 + 1/(C_3 - 1)$. For C_3 between 1.36 and 3.14, the delay is, paradoxically, worse than it would be without the zigzag path. The delay diminishes to 1.18 as C_3 goes to infinity.

For any fixed $C_3 > 1.36$, the network with user-optimal routing will reach equilibrium in a curious manner. With the opening of the zigzag route, the delay D_3 of the zigzag path p_3 remains lower than D_1 and D_2, so all traffic eventually shifts to the shortcut. The delay D_3 remains lower than D_1 and D_2 because of congestion on the bottleneck links with capacities C_2 and C_4. However, when equilibrium is reached, the delay can be worse than before the addition of the new link.

The above apparent paradox arises because an individual's choice of least delay path does not consider delays imposed by congestion on other individuals. If flow in a path is reduced, then in itself that reduction causes a reduction in delays in that path. However, if the flow is displaced to a second path which shares a link with and traffic from the first path, the increase in delay in the common link can more than negate the decrease in delays in other links. That is, delay in the first path may increase because of the increased delay in the shared link. In Fig. 3, as the zigzag link with C_3 between $1.36 \leq C_3 \leq 3.14$ is introduced, nonequilibrium traffic abandons the left and right paths for shorter delay in the zigzag path. Initially, traffic in the zigzag path does enjoy shorter delay, but at the expense of increasing delay in the left and right paths because of congestion in the shared links. As traffic approaches equilibrium, delay in the zigzag path increases past the original

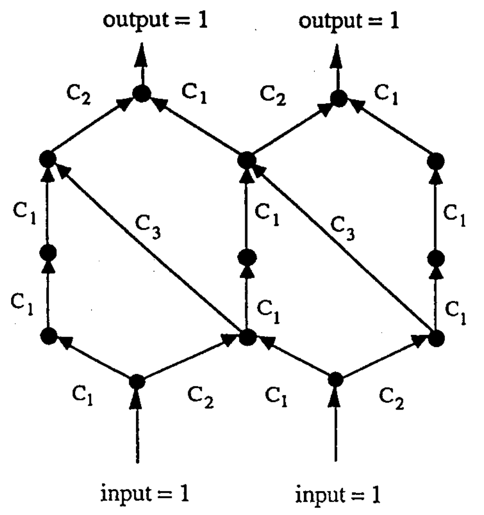

Fig. 5. A network with two in–out pairs.

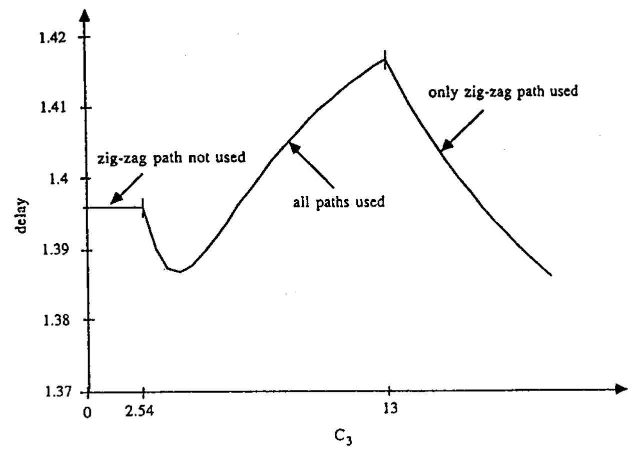

Fig. 6. Delay for the flows in Fig. 5 as a function of the capacity C_3 of the middle links in the zigzag paths.

delay, but delay in the right and left paths increases even more. Thus, traffic continues to choose the zigzag path to the extent that, when equilibrium is reached, delays have increased along the old paths.

In a recent paper by Korilis *et al.* [12], a network like that in Fig. 3, but with even more repeated links (54 on each side) exhibits increased delay (compared to $C_3 = 0$) for all C_3 above a threshold, even infinite C_3. To avoid the paradox in the special cases where all users have identical demands or where users route all flows through links (or paths) of minimal delay, the authors advocate adding additional capacity either to all links in proportion to existing capacity or directly from the source node to the destination node. However, recipes for avoiding the paradox with more general classes of users do not appear to be known.

Similar apparent paradoxes arise in single-server queueing networks with more than one in–out pair. The network in Fig. 5 has two in–out pairs, each with unit total flow. All flow from the left input is assumed to go to the left output, and all flow from the right input is assumed to go to the right output. If capacities are $C_1 = 4$ and $C_2 = 2.5$, then as C_3 varies from 0 to infinity, the two zigzag paths are not used, used with all other paths, or used exclusively. For $6.04 < C_3 < 17.0$, the delay is worse than in the absence of the C_3 link (Fig. 6).

V. Congestion Taxes

Congestion taxes have been proposed as a means of improving the efficiency of other congested networks [3], [4], and the principle applies here as well. To illustrate this, we refer again to the networks in Fig. 3. Various combinations of taxes could be imposed on the short-cut link with capacity C_3 or on the the congested links with capacities C_2 and C_4 so that user-optimal routing with taxation would minimize the total cost T. For example, if, as before, $C_1 = C_5 = 8.9$ and $C_2 = C_4 = 2.7$, the capacity $C_3 = 2.5$ results in user-optimal delays $D_i = T = 1.6907$ in all three paths with flows $f_1 = f_2 = 0.1066$ and $f_3 = 0.7868$. By contrast, the optimal total cost is $T = 1.5874$ when $f_1 = f_2 = 0.3193$ and $f_3 = 0.3614$. Then $D_1 = D_2 = 1.6606$ and $D_3 = 1.4580$, so the delay of traffic along the zigzag path is less than that along the left and right paths by $1.6606 - 1.4580 = 0.2026$. If a monetary tax equivalent to an additional time delay of 0.2026 is imposed on traffic through the link with capacity C_3, then traffic seeking the user-optimal routing with taxation will reach an equilibrium with the above optimal total cost $T = 1.5874$. The tax could be imposed by agreement among all participants to remove any apparent unfairness in the flow that minimizes total cost. The proceeds of the taxation could be distributed equally to all participants, possibly as services or infrastructure. Under user-optimal routing with optimal taxation, the delay along every path would be lower than the delay under user-optimal routing without taxation.

VI. Time-Dependent Flows

The delay values in the above graphs can be calculated by allowing certain systems of difference equations to converge to equilibrium states. Time is discrete with values $t = 0, \Delta t, 2\Delta t, \dots$, and for every in–out pair of nodes, there is a given sequence of total flows $F(0), F(\Delta t), F(2\Delta t), \dots$ (each F is constantly 1 in our examples). As explained in Section II, the flow total $F(t)$ of a given in–out node pair can be partitioned as $F(t) = f_\alpha(t) + \cdots + f_\beta(t)$ where the flow in path p_α is some nonnegative function $f_\alpha(t)$. Thus, the set $\{f_\alpha(t)\}$ of all flows along every possible path associated with every in–out node pair comprises a set of microscopic dynamic variables. Since, at equilibrium, delays along all paths with positive flows are equal, we devised heuristic equations that increase flows along paths with low delays and decrease flows along paths with high delays, until delays along all paths are equal.

For example, in Fig. 5, there are two in–out node pairs. The flow of the first in–out node pair goes in at bottom left and out at top left, and $F_1(t) = 1$ for all time. Three paths are possible, and so there are three flows with $f_1(t) + f_2(t) + f_3(t) = 1$. Likewise, the flow of the second in–out node pair goes in at bottom right and out at top right, and $F_2(t) = 1$ for all time. Again, three paths are possible, and so there are three more flows with $f_4(t) + f_5(t) + f_6(t) = 1$. Flow dynamics as C_3 changes can be modeled by a six-dimensional difference equation system with system variables $\{f_1(t), f_2(t), \ldots, f_6(t)\}$. (The constancy of $F_1(t)$ and $F_2(t)$ here implies that some of these variables are redundant, but this harms nothing in obtaining the equilibrium flows.)

We recall from Section II that link number i has capacity C_i, and that when flows $f_\alpha, \ldots, f_\beta$ are all using link i, the (always positive) delay $d_i(t)$ of each flow through that link is $d_i(t) = 1/(C_i - f_\alpha(t) - \cdots - f_\beta(t))$. The total delay $D_\alpha(t)$ of a flow $f_\alpha(t)$ in a path for an in–out node pair is the current sum of all delays of links in that path. For a given in–out node pair and nonequilibrium flows $\{f_\alpha(t)\}$, we first calculate the variables

$$\begin{aligned} &\sim f_\alpha(t + \Delta t) \\ &\quad = f_\alpha(t) + \sum_{\beta \in \text{ all possible paths for the given in-out pair}} \\ &\qquad f_\alpha(t)[D_\beta(t) - D_\alpha(t)]. \end{aligned}$$

Thus, for example, if delays for all paths other than path α are consistently less than $D_\alpha(t)$, then variable $\sim f_\alpha(t)$ is less than the flow $f_\alpha(t)$. Let $\sim F(t + \Delta t)$ be the sum of $\{\sim f_\alpha(t + \Delta t)\}$ for the given in–out node pair. The new flow values $\{f_\alpha(t + \Delta t)\}$ are computed by

$$f_\alpha(t + \Delta t) = \frac{F(t + \Delta t)}{\sim F(t + \Delta t)} \sim f_\alpha(t + \Delta t).$$

Thus, for each t, flows $\{f_\alpha(t)\}$ and delays $\{D_\alpha(t)\}$ are used to define terms $\{\sim f_\alpha(t+\Delta t)\}$ and $\sim F(t+\Delta t)$, and then new flows $\{f_\alpha(t+\Delta t)\}$ with the correct given total flow $F(t+\Delta t)$. The equilibrium states of the corresponding six-dimensional system were used to produce the delay curves in Fig. 6.

A variety of iterative solution techniques [1, p. 388] could be used here. For example, Nash equilibria in noncooperative product-form networks are discovered by linear programming in [13]. Our difference equation scheme is motivated by the nature of noncooperative network dynamics.

With small perturbations, a path with initially zero flow and smaller delay than other used paths will acquire traffic. Likewise, flow in an initially used path with high delay can asymptotically approach zero. Provided the rates of change of total flows are sufficiently small, user-optimal traffic dynamics are captured by this method.

Time-dependent $M/M/1$ networks in which traffic greedily switches to least loaded paths can exhibit routing instabilities such as oscillations and loops. For example, the original ARPANET routing scheme, using a link metric in which packet delay was averaged over a 10 s interval, performed effectively under light to moderate loads. However, under heavy loads, routing was potentially unstable and far from optimal (see Khanna and Zinky [14]). A new, heuristic metric installed in July 1987 caused an immediate and dramatic decrease in the frame drop rate. The new metric incorporates averaging with previous utilization estimates, as well as upper and lower limits on both metric values and the rates of change of metric values, all with the goal of damping oscillations within routing overhead constraints.

Admittedly, our paradoxical examples deal only with constant equilibrium flows. However, we anticipate that networks to which such paradoxically counterproductive links are added and in which users choose among currently *or recently* optimal paths would be especially good candidates for unstable time-dependent flows.

VII. Concluding Remarks

Practicing designers are well aware that adding links might not improve the performance of some systems. This awareness stems from experience. However, our point is to present a mathematically rigorous set of examples based solely on the simplest single-server queues, bringing the phenomenon out of designer lore and into a scientific framework.

The above apparent paradoxes arise because an individual's choice of least delay path takes no account of the delays imposed by congestion on other individuals. If two paths share a link, and traffic from the first path and other paths changes to follow the second path, the delay in the first path may increase because of the increased delay in the shared link. For example, in Fig. 3, as the zigzag link with C_3 between $1.36 \leq C_3 \leq 3.14$ is introduced, nonequilibrium traffic abandons the left and right paths for shorter delay in the zigzag path. Immediately, however, the delays in the left and right paths increase because of congestion in the shared links with capacities C_2 and C_4. Even though delay in the zigzag path increases, delay in the right and left paths increases even more. Thus, traffic continues to choose the zigzag path to the extent that, even when equilibrium is reached, delays have increased along the old paths.

These examples illustrate, for single-server queueing networks, a general phenomenon in noncooperative games, often illustrated by the Prisoner's Dilemma: user-defined optima are not system optimal [15], [16]. More technically, under broad assumptions, Nash equilibria are Pareto-inefficient [17], [18].

If traffic chooses paths with minimal delay, given the choices of other traffic, the addition of a link can paradoxically degrade system performance. However, when traffic is routed to minimize total cost, the addition of links never worsens the total cost [1]. In general, enforcing flows with minimal total cost implies that some traffic uses paths with nonminimal delay. Incentives, priorities, or systems of compensation are sometimes required in congested networks to induce users to choose routes that minimize the total cost [4].

Acknowledgment

The first author, J. E. Cohen, acknowledges the hospitality of W. T. Golden and Mrs. S. R. Golden. The authors thank the referees for valuable comments.

References

[1] J. Spragins, J. Hammond, and K. Pawlikowski, *Telecommunications Protocols and Designs*. Reading MA: Addison-Wesley, 1991.
[2] D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1987.
[3] F. P. Kelly, *Phil. Trans. Roy. Soc. (London)*, ser. A, vol. 337, pp. 343–367, 1991.
[4] R. Arnott and K. Small, "The economics of traffic congestion," *Amer. Scientist*, vol. 82, pp. 446–455, 1994.
[5] D. Braess, "Über ein Paradoxon der Verkehrsplanung," *Unternehmensforschung*, vol. 12, pp. 258–268, 1968.

[6] R. Steinberg and W. Zangwill, "The prevalence of Braess's paradox," *Transportation Sci.*, vol. 17, pp. 301–318, 1983.
[7] J. E. Cohen and F. P. Kelly, "A paradox of congestion in a queuing network," *J. Appl. Prob.*, vol. 27, pp. 730–734, 1990.
[8] B. Calvert and G. Keady, "Braess's paradox and power-law nonlinearities in networks," *J. Australian Math. Soc. B*, vol. 35, pp. 1–22, 1993.
[9] J. E. Cohen and P. Horowitz, "Paradoxical behavior of mechanical and electrical networks," *Nature*, vol. 352, pp. 699–701, 1991.
[10] J. E. Cohen, "The counterintuitive in conflict and cooperation," *Amer. Scientist*, vol. 76, pp. 576–583, 1988.
[11] R. Jain, "Congestion control in computer networks: Issues and trends," *IEEE Network Mag.*, vol. 4, pp. 24–30, May 1990.
[12] Y. Korilis, A. Lazar, and A. Orda, "Architecting noncooperative networks," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1241–1251, 1995.
[13] Y. Korilis and A. Lazar, "On the existence of equilibria in noncooperative optimal flow control," *J. Assoc. Comput. Mach.*, vol. 42, pp. 584–613, 1995.
[14] A. Khanna and J. Zinky, "The revised ARPANET routing metric," presented at *SIGCOMM'89 Symp.*, also, *Comput. Commun. Rev.*, vol. 19, pp. 45–56, 1989.
[15] G. Hardin, "The tragedy of the commons," *Science*, vol. 162, pp. 1243–1248, 1968.
[16] R. Axelrod, *The Evolution of Cooperation.* New York: Basic Books, 1984.
[17] S. Dafermos and A. Nagurney, "On some traffic equilibrium theory paradoxes," *Transportation Res. Bull.*, vol. 18, pp. 101–110, 1984.
[18] P. Dubey, "Inefficiency of Nash equilibria," *Math. Oper. Res.*, vol. 11, pp. 1–8, 1986.

Joel E. Cohen received the A.B., A.M., Ph.D. degrees (applied mathematics), M.P.H. and Dr.P.H. degrees (population sciences and tropical public health) from Harvard University.

He is the Abby Rockefeller Mauze Professor of Populations at Rockefeller University and Professor of Populations at Columbia University, both in New York City. His most recent book is *How Many People Can the Earth Support?* (New York: W. W. Norton, 1995).

Clark Jeffries (M'92), received the Ph.D. (mathematics) from the University of Toronto, Ont., Canada.

He is Professor of mathematics at Clemson University, Clemson SC, and presently a visiting senior engineer with the Networking Hardware Division of IBM, Research Triangle Park, NC. He is the author of *Code Recognition and Set Selection with Neural Networks* (New York: Birkhauser, 1991) and co-inventor of three neural network patents.