from *SIAM News*, Volume 33, Number 2

Mathematical Modeling and Control of Internet Congestion

By Ramesh Johari

If your neighborhood has cable TV service, you may have been tempted by the idea of a cable modem. These ingenious devices connect you to the Internet over your cable line, 24 hours a day, 7 days a week. Unlike the standard modems in use since the 1980s, a cable modem leaves your phone line free.

Unfortunately, few technological advances come without a cost. In the case of cable modems, consumer enthusiasm has had a very undesirable side effect: congestion. Cable lines are typically shared within a neighborhood. As people have rushed to make use of the new technology, the amount of space on the line (the bandwidth) per user has proved to be far lower than anticipated.

The congestion experienced by cable modem users is one example of an Internet-wide problem. While available bandwidth is increasing, the demand for that bandwidth is growing even faster. Control of congestion on the Internet, therefore, is a very important engineering topic—the future stability of the network depends on it.

In fact, the Internet already implements a form of congestion control. The issue is significantly complicated by an economic problem, however. Consider the following situation: Three people are competing to use an Internet link connecting New York and London. Adam wants to use a Web phone service to call his girlfriend; Bonnie wants to download a 500-megabyte collection of digitized classical music; and Charlie wants to play a networked game against a friend in London. If the link is not large enough to support all three activities, congestion will result. But which of the three users should be forced off the network?

In economic theory, multiple demands for a scarce resource are mediated through a market. In that respect, limited bandwidth is no different from a scarce supply of shares on the stock market. The current Internet, however, does not differentiate among users or their demands: The bandwidth allocations made to Adam, Bonnie, and Charlie are, in an economic sense, arbitrary. It may be, for example, that Adam is willing to pay any amount to make his call; he would then naturally expect to receive use of the network. Unfortunately, the Internet today lacks a mechanism by which Adam can express his willingness to pay, and he is left at the mercy of the network.

If everyone used the Internet for the same purpose, this economic problem would not arise. As demonstrated in our simple example, however, all users are by no means the same: Some make calls, others transfer files, still others play interactive video games. When a resource is congested, some means of differentiating among users is needed. Market theory suggests that *differentiated service* can arise from a pricing scheme based on the level of congestion. As the resource becomes overloaded, only those who are willing to pay higher usage prices remain on the network. This proposal, *congestion-dependent pricing*, has spurred a large body of research on ways in which the Internet could be modified; for more details, see [3].

> *Market theory suggests that* differentiated service *can arise from a pricing scheme based on the level of congestion. As the resource becomes overloaded, only those who are willing to pay higher usage prices remain on the network.*

The Transmission Control Protocol

The operation of the Internet is probably mysterious to most end-users. What exactly happens when a hyperlink is clicked in Netscape, or when a file transfer starts? Let's take a brief look at TCP (Transmission Control Protocol), the primary protocol for controlling the transfer of data across the Internet.

At the finest level, data on the Internet are transferred in bundles called *packets*. All Internet traffic, from phone calls to file downloads, is broken down into packets. From the point of view of an individual link in the network, these packets are nearly identical; they carry information about the sender and receiver, but nothing about the type of traffic being sent. TCP controls the transfer of these packets across the network.

TCP is an *acknowledgement-based* system. In many countries, when you send a letter at the post office, you can request confirmation of delivery; once the letter has arrived safely, the post office will notify you with a postcard. TCP provides confirmation of delivery on a packet-by-packet basis: When a packet arrives safely at its destination, an acknowledgment (ACK) is sent to the sender.

The ACKs serve a dual purpose. First, if no ACK is received, the sender deduces that the packet has been lost. Second, the time between the sending of a packet and receipt of an acknowledgment gives the sender some idea of the communication delay, or round-trip time (RTT), between sender and receiver.

Our interest in TCP is mainly in the congestion avoidance algorithm, first proposed by Jacobson in [6]. This phase of TCP maintains a window of packets to be sent, of size **`cwnd`**. Each time an ACK is received, **`cwnd`** is incremented by 1/**`cwnd`**. Increasing **`cwnd`** only when ACKs arrive ensures that the algorithm is not sending data faster than the network can handle it—the *self-clocking* property of TCP. This structure has the practical effect of increasing **`cwnd`** by one per RTT (known as *additive increase*). The loss of a packet, an event the sending system interprets as an indication of congestion, has the effect of halving **`cwnd`** (known as *multiplicative decrease*). This basic dynamical system, and its variants, are used everywhere on the Internet today. Further details can be found in [1, 6, 7].

The “Smart” Market

With a basic idea of how TCP works, we can consider possible alternatives that incorporate congestion-dependent pricing. One interesting and conceptually simple proposal is the “smart” market of Mackie-Mason and Varian [9]. The smart market aims to replace TCP’s congestion avoidance mechanism with a market-based congestion control mechanism.

Terminology aside, the smart market proposal is just basic economics at work. Suppose that a single link in the Internet receives a flood of data all at once, with the packets numbered $1, \ldots, n$. If the link can accept only m of these packets for processing, which m should be chosen? Mackie-Mason and Varian propose that each packet i carry with it a “bid” w_i, a price the sender is willing to pay to have that packet sent safely. The network then chooses the m highest bids and processes those packets.

Easy enough, but where does pricing come in? After all, if the bids do not correlate to hard money, the congestion control can’t be enforced. It’s here that the market gets “smart.” Suppose, for convenience, that the packets are already in decreasing bid order: $w_1 \geq w_2 \geq \ldots \geq w_n$. Packets $1, \ldots, m$ will be accepted, and Mackie-Mason and Varian identify w_{m-1} as the price to be charged to packets $1, \ldots, m$. In economic terms, this is just the “marginal cost,” the cost of sending one additional packet. When price equals marginal cost, a market is in equilibrium—precisely the desired behavior.

While conceptually appealing, the smart market proposal suffers from several im-plementation problems. First, it is unreasonable to expect users to bid on a packet-by-packet basis, especially on today’s fast-moving Internet. With 64 bytes per packet, a user can easily download 1000 packets while viewing a single Web page!

Another problem with the smart market proposal is the substantial investment in new router hardware that would be necessary. This is a general problem with schemes that suggest modifications at Internet links: Router hardware is expensive to implement and is thus upgraded very slowly. For this reason, upgrades of user software are considered a more feasible means of modifying TCP.

Finally, because the smart market proposal involves individual links, we have no guarantees of network-wide stability. The network should not fluctuate wildly, or collapse altogether, under any congestion control scheme. How is this stable behavior to be ensured?

Stability and Network-wide Congestion Control

The smart market concept makes the point that simple economics can lead to congestion control. Kelly et al. [8] have applied this idea to develop a network protocol that combines a market mechanism with the additive increase/multiplicative decrease properties of TCP.

Three fundamental quantities are of interest: rates, bids, and charges. The sending rate is the number of packets transferred per unit time. The bid is the amount a user is willing to pay per unit time. Finally, the charge to users of congested links is based on the level of congestion.

The dynamical system proceeds as follows: Users begin sending packets. When the total flow through a given resource is moving at a rate of y packets per unit time, the charge per unit flow is $p(y)$. Thus, a user sending at rate x will be charged $x p(y)$. More generally, the charge to a user of multiple links will be proportional to the sending rate. By adjusting the sending rate through an additive increase/multiplicative decrease scheme, the user attempts to bring the network’s charge into equilibrium with his bid. (For details, see [8].)

“Equilibrium” is a key word here. The link to the smart market comes from the fact that the equilibrium is achieved when demand and supply are balanced. Moreover, this equilibrium is “globally stable”: The system always converges to it [8].

But what about implementation? Remember, changing router technology is expensive—network managers don’t like to do it. The system, however, has two important features: First, it is the end-users who im-plement the congestion control; they determine what they are being charged, and adjust their rates until the system reaches equilibrium. This part of the algo-rithm can be implemented by simple modifications at the software level [4]. Second, the links are asked to determine a charge based only on the total flow passing through them.

Enter ECN—Explicit Congestion Notification [2]. This (relatively new) feature of TCP controls a single bit in each packet. At the discretion of a link in the network, a packet can be “marked”—i.e., the ECN bit can be switched on—to signal the sender that congestion is occurring. The number of marks a sender receives can serve as a packet-level implementation of the charging scheme described earlier [4]. The only problem, of course, is that links must use marks as a vehicle to convey prices.

Again, we can apply basic economics. In a fair market, the price is just the marginal cost. What is the cost to a link of accepting an additional packet? If the link “overflows” while the packet is passing through it, i.e., if capacity is exceeded, then the cost is one lost packet. If the link does not overflow during the lifetime of the pa-cket, then no cost is incurred. But this “0–1” pricing can be conveyed perfectly through the ECN bit! A packet is marked 1 if the link overflows during the packet’s lifetime, and 0 otherwise. The number of marks received will be proportional (on average) to the sending rate; hence, this scheme can be used to implement the congestion control algorithm of [8].

Unfortunately, the difficulties that arise with this scheme are much like those encountered by Internet stock investors. In retrospect, many Internet stocks are seen to be huge winners, but choosing the winners in advance would be an exercise in black magic. Similarly, the network suffers from a lack of *previsibility*: Although the pricing scheme described here works, a link has no way of knowing in advance which packets will arrive during an overflow period. Much interesting research, therefore, concerns marking strategies that can give the fair market behavior described here. For details, see [11].

Only the Beginning

Having discussed possible congestion-dependent pricing mechanisms, we might reasonably ask which of them is the "right" solution. But part of the excitement in this area stems from the fact that, currently, no one knows what the best solution will be.

Recall, for example, that the smart market has been criticized for requiring the user to make constant bids. While the model proposed by Kelly et al. in [8] may not require this constant updating, it still allows a great degree of freedom in the choice of rates and payment. In contrast to these schemes, some advocate the "Paris Metro" pricing scheme of Odlyzko [10] (named for the Paris subway system). This scheme establishes several separate logical networks, which differ only in price—the idea being that a more expensive network will effectively provide higher bandwidth, because fewer people will be willing to pay the higher price. While it is not an efficient market, advocates of Paris Metro pricing believe its advantage to lie in its simplicity: With just a few service classes to choose from, the confusion and difficulty of a more complex pricing scheme are avoided.

Other questions remain concerning these strategies in a competitive environment. The Internet comprises many smaller networks, each separately owned and operated. How can we ensure that a new congestion control scheme based on pricing will survive? A recent paper demonstrates, for example, that Paris Metro pricing may not persist in a competitive environment [5].

All these issues have made network modeling and control a dynamic field at the start of the new millennium. Economists, applied mathematicians, and engineers are joining research efforts to help develop the next generation of networks. No one can say what path networks will take, but one thing is certain: As networks grow on a global scale, and telephone, data, and even television networks are combined, network models will grow in importance. It no longer suffices to build a network and hope it works; models such as those presented here aim to provide a solid theoretical basis for future design.

References

[1] J. Crowcroft and P. Oechslin, *Differentiated end-to-end Internet services using a weighted proportionally fair sharing TCP*, ACM Comp. Commun. Rev., 28 (1998), 53–67. Available from http://www.cs.ucl.ac.uk/staff/jon/hipparch/pricing.html.

[2] S. Floyd, *TCP and explicit congestion notification*, ACM Comp. Commun. Rev., 24 (1994), 10–23. Available from http://www-nrg.ee.lbl.gov/floyd/ecn.html.

[3] R.J. Gibbens, *Control and pricing for communication networks*, to appear in Phil. Trans. R. Soc. Lond. A (2000). Available from http://www.statslab.cam.ac.uk/cgi-bin/resreps.pl?term=1999-10.

[4] R.J. Gibbens and F.P. Kelly, *Resource pricing and the evolution of congestion control*, Automatica, 35 (1999), 1969–1985. Available from http://www.statslab.cam.ac.uk/~frank/evol.html.

[5] R.J. Gibbens, R. Mason, and R. Stein-berg, *Internet service classes under competition* (1999). Available from http://www.soton.cam.ac.uk/~ram2/papers/pmp10.pdf.

[6] V. Jacobson, *Congestion avoidance and control*, Proceedings of ACM SIG-COMM '88, 1988, 314–329. Available from ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z.

[7] F.P. Kelly, *Mathematical modelling of the Internet*, to appear in Proceedings of the Fourth International Congress on Industrial and Applied Mathematics, 1999. Available from http://www.statslab.cam.ac.uk/~frank/mmi.html.

[8] F.P. Kelly, A.K. Maulloo, and D.H.K. Tan, *Rate control in communication networks: Shadow prices, proportional fairness, and stability*, J. Oper. Res. Soc., 49 (1998), 237–252. Available from http://www.statslab.cam.ac.uk/~frank/rate.html.

[9] J.K. Mackie-Mason and H. Varian, *Pricing the Internet*, in B. Kahin and J. Keller, eds., *Public Access to the Internet*, MIT Press, Cambridge, MA, 1995. Available from http://www-personal.umich.edu/~jmm/papers.html#pi.

[10] A. Odlyzko, *A modest proposal for preventing Internet congestion*, 1997. Available from http://www.research.att.com/~amo/doc/networks.html.

[11] D. Wischik, *How to mark fairly*, Workshop on Internet Service Quality Economics, MIT. Available from http://www.statslab.cam.ac.uk/~djw1005/Stats/Research/marking.html.

Ramesh Johari (r.johari@statslab.cam.ac.uk) is a graduate student of Frank Kelly at the Statistical Laboratory, University of Cambridge, Cambridge, U.K.