

Service Differentiation: Congestion Pricing, Brokers and Bandwidth Futures

Peter Key

Microsoft Research,
St George House, 1 Guildhall Street,
Cambridge, CB2 3NH, UK
peterkey@microsoft.com

1 Background

In most Networks or Systems, users compete for a set of limited resources. Resources might be point-to-point bandwidth, buffer space, memory, CPU cycles etc. and users can be humans, applications or processes. Connection-oriented Telecoms services use either reservation, where resources are booked in advanced or Admission Control — think of telephony or ATM for example. Packet-based services and Operating Systems rely mainly on either priority schemes, for instance using simple priority queues, or flow-control feedback mechanisms such as TCP. The current Internet uses flow-control with an admit-all policy that relies on secondary measures, such as users' willingness to tolerate current service, to throttle excess load. Within the IETF, IntServ is attempting to define connection-oriented admission control schemes whilst DiffServ is based on connectionless ideas, attempting to use priorities and complex queuing disciplines to provide QoS classes.

But if we step outside the Computer Science or Telecommunication mindset, we see the same problem occurring in other areas, but solved using resource pricing as a fundamental ingredient. For example, think of resources as airline seats, units of (electrical) power or economic wealth. Prices or taxes are used to control the load, or to maximise the return by offering differential services at different prices — think for instance of power companies who use price incentives and penalties, or airlines.

We claim not only that it is impossible to separate service differentiation from pricing, but also that some use of resource pricing is the only sensible way to proceed. The first point is clear: if there are two qualities which both cost the same, any rational person will chose the higher quality! As an example, most current DiffServ proposals are interesting but fundamentally flawed, caused by attempting to define technologies independently of economics or implementation. They attempt to define local priority schemes and scheduling policies, whereas an application/user is interested in

end-to-end performance. Also, the cost and complexity of scheduling policies is not addressed. Even if the hardware cost is not significant, there are real management overheads. A salutary example is provided by the ABR service of ATM: this only really makes sense if ABR operates from end-system to end-system, which now looks highly unlikely. Moreover, despite being more complex to implement, it would have to be charged at a *cheaper* tariff than say a CBR service operating at the MCR, despite using the same bandwidth. (Otherwise a smart user with delay insensitive traffic could just use CBR and overflow excess traffic to a best-effort traffic class that has be to an order of magnitude cheaper).

The Explicit Congestion Notification (ECN) [1] proposal within DiffServ looks attractive, and we believe can be extended to give as much diversity as is required. Our contribution is to break away from shackles of mandatory TCP-like behaviour inherent in the current ECN RFC. We can even incorporate call admission control at the flow or micro-flow level, an anathema to the current DiffServ!

2 Architecture and Framework

The idea of Congestion Pricing has been explored in a number of papers [2,3], where the basic idea is to signal back derivative information to the user (the Lagrange multiplier of the congestion costs, representing the marginal incremental costs). Prices can be simply calculated by each resource; if a resource is lightly loaded, the price is zero, and rises as the load increases. Feedback signals are proportional to flows, enabling multiplicative decrease. There are links with the Vickrey-auction proposal of MacKie-Mason and Varian [4], though we use bounded costs, and no explicit signalling. Marking only lost packets is *not* a sensible pricing mechanism. With this framework, complex queuing scheduling is not needed: FIFO can be used, with a suitable marking scheme using counters or adapting RED ideas.

These signals have to be communicated to the user. For IP based networks we advocate marking at the IP level, which then allows control to be exercised

at any higher layer, *or the IP layer itself*. This also allows unresponsive applications or domains to co-exist with responsive ones — unresponsive flows will in general incur higher charges. Other approaches could be used elsewhere, such as a virtual signalling channel.

Then the fundamental idea is to allow end-systems to react as they please, so users who are prepared to “pay” more receive more. The only assumption required is that they must feel some pain associated with the feedback signals / marks. Two possible scenarios are

1. Users have complete freedom to behave as they wish, in this case the network will run below congestion provided the aggregate demand is less than the capacity (where we need to take expectations over time, and demand is some maximal flow vector). This will happen if the sum of the users’ “willingness to pay” is less than the maximum charge rate for the system. The system provider can raise prices to achieve this goal, which is appropriate for an operating system, whilst in a competitive network there is an incentive for the provider to increase capacity until equilibrium is reached.
2. The user may have restricted freedom, and a mandated control algorithm that allocates capacity based on user parameters. For example, an ISP/Intranet might give a user a particular flow control algorithm, which has some user or process defined parameter, that reacts to feedback signals mediated by the ISP. Think, for instance, of a TCP-like algorithm where increase and decrease rates are parameterised.

If users’ preferences can be modelled by utility functions, and users seek to maximise their net utility, then the system will behave as a distributed optimisation and converge to the global optimum [5], provided the prices are right, matched to aggregate user demand. In the second case, a global will only be reached if each user is able to use the set of parameters as a proxy for true preferences. These statements are saying that correct taxation leads to the social optimum! The second scenario above hints at congestion pricing for aggregates. For example, a corporate client or network access point to an ISP might receive pricing signals that relate to an aggregate of channels/connections, and decide to redistribute these signals amongst connections/users according to some (corporate) policy. In this case the optimisation is acting at an aggregate level, but even here can reach an end-user optimum depending on how the marks/signals are redistributed – in effect using a gradient projection

for the optimisation. A simple example is an aggregate of voice and video channels: assuming that the per-packet price is the same, voice marks can be handed-off to video channels, causing the video channel to react and back-off rather than the voice channel. Another example arises from asymmetry: suppose a user is retrieving a file from a remote server where the server has some mandated flow control algorithm, then the user will receive marks for the flow from the server, and can decide how to pass them back to change server behaviour.

3 Adaptive Applications and Real Time Services

Feedback systems are particularly appropriate for elastic applications, such as file transfer or Web browsing. If user is prepared to pay a fixed amount per unit time, a simple “willingness-to-pay” algorithm [2,3] adjusts the transmit rate according the difference between this amount and the actual congestion price/feedback signals. It has proved straightforward to adapt TCP to such a scheme: all that is needed is to change the way the congestion window increases and reacts to feedback. Retransmission behaviour is unaffected. If a user desires to transfer a file as fast as possible, a “willingness to pay” rate can be inferred, which has the effect that when the prices are high certain users will stop sending, preferring to wait for the price to drop.

The jury is still out on whether VBR will be used for real-time transport. But there are certainly adaptive coding algorithms, e.g. transcoding or layered coding, and these can be made to react to feedback signals. However, they still need some minimum bandwidth to be of use for the user, which has to last for the life of the connection/flow, which for Video can be extremely long compared to congestion timescales.

It is often assumed that reservation or CAC is necessary for guaranteed services. Not so! CAC relies on the system determining if there is enough capacity, and allowing on flows/processes or rejecting them. But if the network can set the correct price, then connections will decide whether to admit themselves or not, which obviously relies on “price matching” of users to the network. It is possible to show rigorously that threshold control strategies, known to be optimal for certain connection-oriented problems, are exactly equivalent to pricing schemes where a congestion tax is payable by the user. There are some research issues here: a connection is interested in a price over a long time interval, whereas congestion prices change on shorter time-scales, but in another context this is solved with fixed-rate mortgages! Recent suggestions for sending probing packets can

be incorporated in our scheme. As a rule of thumb, sensible congestion pricing will mark at least an order of magnitude more packets than lost packets, so sending several probes will produce better estimation for acceptance decisions.

4 Options, Futures and Bandwidth

An alternative suggestion is to use an agent (e.g. an access node) to act as a broker, taking the risk by appearing to the end-user as a call admission control, whilst paying congestion prices to the network. This effectively uses aggregation to improve statistical estimation efficiency over a single-user probe.

A more extreme possibility is to view bandwidth as a commodity. The seeds of this have already been sown with bandwidth brokerage, and commercial concerns like RateXchange, which acts as a trading exchange for wholesalers. But why not extend this to a much lower granularity? Congestion prices within a network will fluctuate on a short timescale due to the inherent critical time-scale of resources, and on a longer time scale, e.g. diurnal, due to user demand. The critical time-scale should be shortened by having small buffers within the network; this avoids the rush-hour effect, means delay is dominated by transmission delay and enables point-to-point bandwidth to be treated as a commodity, since jitter is minimal a one-dimensional descriptor suffices. A network with fluctuating prices and underlying commodities allows the tools of financial securities to be used: there are direct applications of Futures Contracts, Options and Swaps. For example a fixed video-conference could be “booked” by taking long position in a Futures contract or by a (European) call Option, while a moveable broadcast could use an American style option, with a broker taking the corresponding short position or trading it. On-demand users have to pay the prevailing spot prices. It is natural to see these operating first between providers.

5 Experimental Framework and Preliminary Results

The picture we have painted suggests that the users are in effect playing a game with the network, so why not test the framework by playing a multi-user distributed game? This is exactly what we have done: we have built a real-time distributed event simulator that allows arbitrary network topologies and user behaviours which also allows users (clients) to remotely connect to the simulator and “play”. Our work is inspired by Axelrod’s work [6] on the repeated prisoner’s dilemma. Initial results have validated the concepts: congestion pricing does work, gives more to those who are prepared to pay more, and is simple to implement

at the TCP layer. We are currently implementing controls and users strategies in Windows 2000 to try out IP level controls within a network.

6 Summary

The fundamental position is the following: keep the core of the network very simple and push complexity (if needed) to the end-systems. This is reminiscent of the ATM vision before it was obscured by baroque signalling stacks and traffic management complexity! All that is needed within the core is for switches/ routers to give feedback signals based on congestion prices. With marking or feedback signals at the IP level, legacy end-systems can co-exist with adaptive ones: they will literally pay the price for having sub-optimal reaction to network signals, giving a clear economic reason for upgrade. With an adequately priced network, buffers within the core of network can be small. This not only ensures that queuing delay is negligible, but also facilitates the use of bandwidth as a commodity. Once pricing mechanisms implicitly controls networks, it is natural for derivative markets to flourish. Brokers and arbitrage agents can act as exchanges or risk takers, selling products which mirror those in the financial derivatives sector, or selling services that are repackaged to appear to the end-user as a traditional service.

Selected References:

- [1] K. Ramakrishnan and S. Floyd. A proposal to add explicit congestion notification (ECN) to IP. RFC 2481, IETF, January 1999. <ftp://ftp.isi.edu/in-notes/rfc2481.txt>.
- [2] R.J. Gibbens and F.P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, <http://www.statslab.cam.ac.uk/~frank/PAPERS/evol.html>, 1999.
- [3] P. B. Key and D. R. McAuley. Differential QoS and pricing in networks: where flow control meets game theory. *IEE Proceedings Software*, 146(2), March 1999.
- [4] J.K. MacKie-Mason and H.R. Varian. Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. Prentice-Hall, New Jersey, 1994.
- [5] F.P. Kelly, A.K. Maulloo, and D.K.H Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [6] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, NY, 1984.