Market Control of the Internet

— A Case Study in Engineering a Social Science

Bob Briscoe
<bob.briscoe@bt.com> <www.btexact.com/people/briscorj/>
BT Research, B54/130, Adastral Park, Martlesham Heath, Ipswich, IP5 3RE, England
Tel. +44 1473 645196

30 Mar 2002

Abstract

Our research to engineer control of the Internet indirectly through a fine-grained market is a classic case-study in the use of economic incentives to solve broken trust in a system with shared control. We outline why we believe solving this problem will 'save the Internet' from future stagnation. We highlight successes and pit-falls that are relevant to an audience on the boundary between economics and security.

1 Commercial Openness

We believe that the most successful way to offer an infrastructural service is to allow its customers total freedom to control it. A fine-grained market in the underlying resources used is all that is necessary to temper this freedom with responsibility. The service provider decides what the service will start as (design phase). Customers decide what it will be used for and how it will evolve (operational phase). We call this 'commercial openness', to capture the tension between freedom and responsibility.

In the early 1980s Saltzer *et al* articulated the end to end principle [5] behind the design of the early Internet. Viz. network infrastructure should be minimalist with all intelligent control on the end-systems using it. The motivation was to foster innovative evolution without compromising scalability. The rudimentary network capabilities were expected to be used to synthesise different capabilities from those originally envisaged. Providing different behaviours on each path through the Internet wouldn't cause design conflicts because the differentiating features were only at the very ends. The Internet could grow inexorably, as the growth in supply of resources on end systems naturally matched the growth in demand. Perversely, the dumber the network, the more valuable it could be.

Two decades later Clark & Blumenthal questioned the continued relevance of the end to end principle [3]. Arguing that the Internet had become a victim of its own success. There was now too much at stake commercially to be able to trust customers to co-operate in the control of the Internet. The inevitable consequence would be "a slow ossification of the form and function" of the Internet.

We agree that you can't have freedom without responsibility. But if some people are being irresponsible, you don't have to take away everyone's freedom. Instead, you can enforce responsibility. So we have designed, built and analysed the mechanisms of a fine-grained market between all Internet stakeholders. Below we report on the lessons we have learned from our efforts to 'save the Internet' from the stagnation prophesied by Clark & Blumenthal — a real case study of how to engineer a theoretical model of economic incentives to solve a hugely important problem of broken security.

2 Approach

Our approach is first and foremost not to promote a single approach. To function well, a market must be diverse, not just competitive. Therefore each network provider must be free to choose and to change its approach to the market, while still interoperating with others.

To be clear, our scope is far beyond the 'one size fits all' Internet of today. We even look beyond the market for differentiated classes of data delivery, that is emerging today. Currently this market is confined to islands of first-movers selling these capabilities as a market differentiator. We even look beyond the time when the economic advantages of the 'network effect' will outweigh the current first-mover advantages, so networks will harmonise their solutions to allow anyone to buy quality of service to anywhere else. We expect commercial openness to become important when the market in quality of service commoditises, perhaps in 5–7 years time. Only then will the benefit of giving customers fine-grained control of their own quality start to be realised. Unlike other approaches, we are preparing

for a future where customers will be able to run software that synthesises its own quality of service behaviours on a per-*packet* basis. Alongside this approach, we expect network operators to continue to embed capabilities *in* their networks to offer a restricted menu of qualities of service for human users on a per-*flow* basis. But our foresight will ensure that the industry doesn't embed technology in the network that will restrict it *only* to current preoccupations (e.g. streaming media). We foresee an Internet of the future where communities of interacting computers are vying for the last squeak of performance from the network, while others have far more elastic demands.

To this end, we have described an architecture [2] that encompasses the so-called 'Internet architectures' we see emerging today (we call them service plans), as well as novel service plans in the future and ones we and others have devised already. We have specified a principled way to allow network operators use policy control to switch from one service plan to the next as this market evolves, rather than re-embedding new technology in their network for each change. The whole internetwork of competing approaches will form a market managed multiservice Internet (M3I — the name of our collaborative project).

We have experimented with extensively with Kelly *et al*'s sample path shadow pricing (SPSP [4]) service plan. Briefly, it is now standard for a mark to be randomly applied to packets entering congested links. In SPSP a tiny fixed price is applied to these marks (congestion pricing). We have built a rate controller that reacts to this dynamic pricing, operating under the control of a 'quality buying policy'. Elastic policies cause it to back-off from congestion, while inelastic policies continue at full rate, or stop completely depending on the price (self-admission control). To populate these buying policies with data, we have conducted customer experiments to establish the bit rate that different people prefer for different tasks at different congestion prices. Congestion pricing also causes revenues to flow to the congested parts of the network.

Below we highlight the major pitfalls and successes of relevance to an audience interested in the interplay between economics and security.

2.1 Perils

Economics doesn't replace security: Although a system aligned with major economic incentives removes some need for traditional security mechanisms, they are still needed for the accounting system that polices the market; for initial customer access to the market-controlled system (triggered by creditworthiness) and for protection against fraud within a complex economic entity like a network provider.

Also on a philosophical level, economic incentives are trumped by more irrational forces. A popular Internet hackers' tee-shirt [citation unknown] shows an eleven layer reference model for system interconnection (unlike the traditional seven). The four additional layers are commercial, legal, political and at the very top, religious. Nothing could be closer to the truth. The commercial layer is about grey-scale policies, whereas higher layers tend to be black and white, and inconsistently related to willingness to pay.

Dynamics: A congestion price is delayed for one round trip time. Risk insensitive policies can overflow the control system in this time, moving it temporarily into an uncontrolled loss regime.

Imperfect competition: A per-packet market mechanism can encourage but not ensure perfect competition for *every* link in the global network. This includes the notorious 'last mile' links where competition tends to raise rather than lower costs.

A utility function for every occasion: To engineer a economic control system, a utility function is needed for every task, every person and every environment. Of course, we assume utility functions for many scenarios are approximately equal, or at least the same shape, but scaled. However, this problem is still of unknown size.

2.2 Triumphs

Structural security: When data is *added* to a system in order to implement a charging system, a fraud incentive is created if it can be removed without affecting the service. Instead we only apply prices to inherent features of the service itself. Just as increasing the money-supply devalues money, creating information in order to charge for it devalues the added information.

Market diversity: Inter-operation between providers each choosing different approaches to the market is eased because money is a globally understood standard that intermediates between each approach [1].

Synthesise old from new: We (the M3I consortium) have built gateways between the dynamic world we predict and the current static world. Wholesale networks can be controlled with congestion pricing while the retail fringes

offer static quality guarantees sold at fixed prices. Our test-bed experiments show that our 'risk broker' controlled gateway successfully synthesises hard admission controlled guarantees for flows at the edge, with absolutely no flow knowledge in or between core networks.

A market in 'pre-congestion': Instead of charging more for worse service, we can sense when the underlying load is *about to* cause congestion of any resource and raise the price to *avoid* worse service.

Cost of charging: Itemisation is the root of most market mechanism costs. Charging for bulk service can still allow fine-grained market control, as the customer has the incentive (and software) to itemise their own charges. With increasing use of encryption this will become the only practical model anyway.

3 Conclusion

Despite the perils listed earlier, we now have a complete solution for buying and selling Internet quality that is in sympathy with the Internet Architecture.

The original end to end principle was generalised to any capital intensive artefact that would outlive requirements captured only at its inception. Of course, its inherent flaw of shared control with divergent incentives was generalised with it. We have recently started new cross-disciplinary research to generalise 'commercial openness' to similar infrastructures.

Acknowledgements

The contributors to the M3I consortium and my colleagues at BT.

Author's biography

Bob Briscoe, MA, heads BT's Edge Research Lab, which specialises in improving the growth prospects of the communications industry by moving intelligence out of networks to their edge and solving the resulting division of control problem. He specialises in innovative commerce and security mechanisms for global scale distributed systems, in particular for multicast applications and previously for Web services. He is technical director of the Market Managed Multi-service Internet (M3I) consortium. He won BT's medal for innovation in 2001. He is also studying part-time for a PhD at University College London. <http://www.btexact.com/people/briscorj/>

References

[1] Bob Briscoe. The direction of value flow in connectionless networks. In *Proc. 1st International COST264 Workshop on Networked Group Communication (NGC'99)*, volume 1736, URL: http://www.btexact.com/projects/mware/, November 1999. Springer LNCS. (Invited paper).

[2] Bob Briscoe. M3I Architecture PtI: Principles. Deliverable 2 PtI, M3I Eu Vth Framework Project IST-1999-11429, URL: http://www.m3i.org/, February 2002. (To appear).

[3] David Clark and Marjory Blumenthal. Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world. In *Proc. Telecommunications Policy Research Conference (TPRC'00)*, URL: http://www.tprc.org/abstracts00/rethinking.pdf, September 2000.

[4] Frank P. Kelly, Aman K. Maulloo, and David K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49, 1998.

[5] Jerome H. Saltzer, David P. Reed, and David D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, November 1984. An earlier version appeared in the Second International Conference on Distributed Computing Systems (April, 1981) pages 509–512.