Marsden, Christopher T. "Positive Discrimination and the ZettaFlood." *Net Neutrality: Towards a Co-regulatory Solution*. London: Bloomsbury Academic, 2010. 83–104. *Bloomsbury Collections*. Web. 26 Apr. 2023. <<http://dx.doi.org/10.5040/9781849662192.ch-003>>.

Downloaded from Bloomsbury Collections, www.bloomsburycollections.com, 26 April 2023, 09:00 UTC.

Copyright © Christopher T. Marsden 2010. You may share this work for non-commercial purposes only, provided you give attribution to the copyright holder and the publisher, and provide a link to the Creative Commons licence.

CHAPTER THREE

Positive Discrimination and the ZettaFlood

How should telcos help? First, create secure transaction environments online, ensuring that consumers' privacy concerns do not prevent m- and e-commerce. 'Walled gardens' ... secure credit and other online methods can achieve this. Second, don't become content providers. Telcos are terrible at providing media to consumers ... Third, do become content enablers ... in developing audio and video formats by participating in standards-building consortia.¹

Marsden (2001)

This chapter considers the case for 'positive' net neutrality, for charging for higher QoS, whether as in common carriage to all comers or for particular content partners in a 'walled garden'. As the opening quotation demonstrates, I have advocated variants of this approach before. The bottleneck preventing video traffic reaching the end-user may be a 'middle-mile' backhaul problem, as we will see. It is essential to understand that European citizens have supported a model in which there is one preferred content provider charged with provision of public service content (information, education as well as entertainment): the public service broadcaster (PSB). Some countries have more than one of these. The UK has four: the British Broadcasting Corporation (BBC) is publicly owned and publicly financed without advertising; Channel 4 is publicly owned but advertising financed; two, ITV and Channel 5, are privately owned and advertising financed.² These PSBs are accustomed to 'must carry' status on the terrestrial, cable and satellite networks of their own countries, and have leveraged that onto IPTV over the Internet and broadband. Thus, in Europe, it is not premium events such as sports broadcast that have led the way in the QoS debate: it is PSBs. The question thus becomes: is net neutrality only for PSBs (sic) or for all content? Obviously, 'must carry' for PSBs may squeeze other non-QoS content into a slow lane. Moreover, if citizens are accessing PSB content (already paid for by television regulatory arrangements), they will have little (or zero) propensity and incentive to pay more for the QoS to stream that in high definition (HD) over their Internet connections.

The debate will come to PSBs and commercial content providers, and you and me as bloggers and social network members, later in the chapter. First, I examine whether the Internet is about to drown in a sea of video-induced data, whether there will be a ‘ZettaFlood’.

The ZettaFlood

Ever since the Internet became commercial – and arguably long before that, even when the initial ARPANET was being built – engineers and analysts have worried that the increase in users, while essential in the astonishing growth that is summarized in Metcalfe’s Law, would eventually overwhelm the network, with traffic so dense that individual routers would drop too many packets for the reliability of the network. As the users reach over a billion, with several billion more using Internet-almost-ready mobile networks, and as their connections increase in speed a hundredfold from the old dial-up lines to broadband, there is a potential ‘meltdown’ brewing. This rapid – if not exponential – growth in traffic has led such doom-mongers to predict that ‘Something Must Be Done’. That a decade or two has passed since the problem and doom was first predicted has – if anything – increased the volume and intensity of the calls to slow down or monetize the growth.

To summarize: there are many more users than the network was built for, and they are using the network to a far greater extent than originally planned. Not only are there millions of university scientists and government users (the original user base) with fibre connections to extract enormous shared computing power, but there are a billion residential and corporate users with varying qualities of fixed or mobile connections. To increase the speed and volume of data transferred using the fibre links which large enterprises use is relatively trivial compared to the problems of those using mobile or old telephone and cable lines to access the data. I do not intend in this book to go into depth on the problems the aggregated data presents for the different ISPs in dealing with each other, but clearly an ISP with a million narrowband (dial-up) users’ data to send and receive is in a very different position to one with ten million broadband users, or a million fibre-enabled super-fast broadband users.

Let us briefly explain what happens. ‘Freetards’ – and this certainly is a pejorative term³ – is the term employed most infamously by technology commentator Orlowski to describe those who continue to believe that ‘information wants to be free’ and should be, accusing them of free-riding

on the back of the average user. Leaving aside the ‘information communism’ implications of the philosophy, if it can loosely be called that, the debate has focused around alternatives to current pricing and distribution models on the Internet, including the far from retarded ideas of Wikipedia, Creative Commons and P2P software distribution. The claim is that ‘freetards’ are using a dangerously disproportionate share of consumer bandwidth in their almost unlimited use of domestic connection as a P2P file-sharing upload and download facility. It can be illustrated as follows: a consumer with a 10 Mbps download/1 Mbps upload speed on his/her domestic connection uses about 3 GB of data per month (this is an average value, but it is certainly in the 1–10 GB range). By contrast, using the maximum theoretical speed at all times, it is possible to use a total of 3.6 TB (3,600 GB approximately).⁴ Maths isn’t my strongest suit but this is approximate, theoretical, therefore highly implausible (as connections are not that reliable 24/7/365). That one mythical ‘freetard’ would be using a thousand times the monthly usage of the average user, or the entire usage of a small town. Could attacking this one user be beneficial to the other 999 people in that town using that telephone exchange? Users are summarily terminated or suspended. This can be conducted by any ISP and may well be justified, but could be made more transparent.⁵ ISPs choose to filter P2P traffic, typically popular file-sharing programs, as in a best effort environment without congestion charging⁶ that content has insufficient disincentives to prevent its flourishing. ISPs can choose to filter P2P traffic of various kinds; typically it is unencrypted relatively crude versions of popular file-sharing programmes, such as BitTorrent, which is used to provide upgrades to the most popular multiplayer online game ‘World of Warcraft’. Many security assertions are made about the implications of certain types of traffic, but regulators currently have no basis for deciding if such assertions represent real problems.⁷

The virtual lynch mob in question is of course an illustration, and in any case this is not the main problem. P2P networks are a problem because they seed many concurrent ‘streams’ of data in order to ‘max out’ the available bandwidth – that is also why they are efficient compared to single ‘streams’ of data. It is as if every second person in the queue at the ticket office is actually working for the same person sitting outside, drinking coffee, while the rush-hour throng mills past. She will of course receive her ticket more quickly, but at the cost to everyone else in the queue, many of whom will be ‘timed out’

of buying a ticket and be forced to buy on the train or find alternatives. As Felten rather more technically, accurately puts the dilemma:⁸

a single router (in the ‘middle’ of the network) ... has several incoming links on which packets arrive, and several outgoing links on which it can send packets ... if the outgoing link is busy transmitting another packet, the newly arrived packet will have to wait—it will be ‘buffered’ in the router’s memory, waiting its turn until the outgoing link is free. Buffering lets the router deal with temporary surges in traffic. But if packets keep showing up faster than they can be sent out on some outgoing link, the number of buffered packets will grow and grow, and eventually the router will run out of buffer memory. At that point, if one more packet shows up, the router has no choice but to discard a packet.

This is where the buffer has set rules for that critical decision to discard packets:

When a router is forced to discard a packet, it can discard any packet it likes. One possibility is to assign priorities to the packets, and always discard the packet with lowest priority. This mechanism defines one type of network discrimination, which prioritizes packets and discards low-priority packets first, but only discards packets when that is absolutely necessary. I’ll call it minimal discrimination, because it only discriminates when it can’t serve everybody. With minimal discrimination, if the network is not crowded, lots of low-priority packets can get through. Only when there is an unavoidable conflict with high-priority packets is a low-priority packet inconvenienced.

This minimal discrimination is what has always happened on the network. P2P services are designed for this orderly non-prioritized packet queuing, and flood the queue to make sure they get through. That makes the job of the router more difficult and is what leads to positive net neutrality problems – when packets are deliberately put into a slow lane. As Felten puts it, this is ‘another, more drastic form of discrimination, in which routers discard some low-priority packets even when it is possible to forward or deliver every packet.’ He terms that ‘non-minimal discrimination’ and its key characteristic is the decision to put a high-priority toll-lane in place which necessarily limits the speed of all other traffic.

So the lazy consumer can pump out 3.6 TB theoretically – multiply him by a million (or find a billion ‘normal’ lightweight users) and you have not petabytes (PB) but exabytes (EB). That’s a lot of data. Cisco Systems, which makes many of the routers used for Internet traffic and therefore has an incentive to suggest rapid traffic growth, has estimated that we are growing traffic towards several hundred EB a year, and Gilder and Swanson⁹ have talked of the next one up, the zettabyte (ZB): ‘We estimate that by 2015, U.S. IP traffic could reach an annual total of one zettabyte, or one million million billion bytes. We began using the term “exaflood” in 2001.’ This flood of data, which they believe will be largely video data, could overwhelm the network: ‘Today’s networks are not remotely prepared to handle this exaflood. Wall Street will finance new telco and cable fibre optic projects, but only with some reasonable hope of a profit. And that is what net neutrality could squelch.’ They claim, moreover, that net neutrality which prevents networks making special arrangements with high-volume video service providers, and thus perhaps investments by networks in fibre, will prevent the next stage of less limited fibre-employed bandwidth at which point net neutrality could really be employed.

One element of this is almost certainly correct: as fibre is built out towards the consumer from the local telephone exchange, the costs of that upgrade (80 percent of it is ‘civil works’, digging holes in the road and trenches in the pavement) will be prohibitive for smaller players, and only possible by monopoly or duopoly networks. As Banerjee and Sirbu summarize¹⁰ in a conclusion backed by Van Der Berg, facilities-based competition among fibre network providers is unlikely due to the economies of scale in installing fibre for an entire street or neighbourhood. They therefore hold that regulators should be wary in removing open access requirements in return for investment in fibre, as it could lead to remonopolization, or perhaps duopoly competition. This is of course what has already happened in the United States under the Bush-appointed head of the FCC, and the recent activity in 2007–9 over net neutrality is the belated response of the regulators to alleged abuses. If fibre means that this duopoly – or monopoly – takes control of the customer access line, then it appears vital to ensure competition for services on that line, whether by Skype, Google or whomever else. As Sirbu further holds, if service competition is limited only to ISPs which own facilities, this would lead to greatly reduced service-level competition, and that reduced service-level competition raises the net neutrality issue.

So we will almost certainly need net neutrality rules at the next stage, the fibre-based consumer Internet. Can we not afford them now, as Gilder and

Swanson suggest? Is it a luxury, so that copper-based networks will not be able to carry HD video? And if ISPs cannot afford the maximum use of their networks now, what should they do about it? That is the first set of questions addressed in this section, before a short journey around what happens to that data as it is aggregated and passed around the Internet (where fortunately we already have fibre), and then the \$64 billion question of what policymakers are currently doing about it, and what they could and should in future.

‘Walled Gardens’ or Preferred Partners

The obvious answer to the data flood is to constrain the customer to locally hosted and/or cached bandwidth. That means a ‘walled garden’, as my quotation opening the chapter suggests. Carriers can offer exclusive, preferential treatment to one application provider over others, creating a type of ‘walled garden’ of preferred suppliers. This is less distorting than blocking, depending on the type of walled garden and the ‘height of the walls’. I differentiate ‘walled gardens’ from an open/interoperable access ‘commons’.¹¹ A ‘walled garden’ is a type of IP content service offered without access to the wider Internet. For example, most mobile telephone networks provided walled gardens to their subscribers. This has wider regulatory implications, involving the development of gatekeepers rather than open access models. Continuing the analogy with commons and walled gardens, one can imagine that a walled garden can be protected and entry or exit charges imposed. By contrast, a commons is open access, with no controls. The walled garden gatekeeper is likely to be the owner of the garden, the operator. ‘Walled gardens’ have historically described content or services bundled by an access provider as a package with fixed or mobile Internet access. The content is usually supplied under contract by content/services providers and presented to the end-user by the access provider as a branded portal. The content or service can be acquired from a third party in exchange for a direct payment. An agreement to share advertising revenue is an increasingly common model. The service offered by an access provider may restrict users to content only in that walled garden. In this case, the access provider is a gatekeeper (like Vodafone Live! when it first launched). Alternatively, the access provider may give users the freedom to access the wider Internet and consume other content and services, including those which may compete with those in the portal. In this case, he is not a gatekeeper to the Internet (like the fixed ISP model).

An example is mobile users, who inhabit a more personal and pervasive environment than fixed Internet users, with additional constraints on full openness. The mobile industry has developed hitherto on the basis that operators control the use of their networks and the devices that connect to them. For that reason, the initial content offerings of mobile providers have tended to be provided in a walled garden in which the customer experience is guaranteed by the operator and discriminatory pricing can be imposed on third-party content providers through their contracts with the mobile operator. It is an open question whether net neutrality measures should be imposed on the practices of such mobile operators, a topic explored in Chapter 7.¹²

In a ‘walled garden’, the number of content providers, and the innovation they are permitted to undertake, is effectively regulated by their relationship with an ISP. ‘Walled gardens’ may evolve so that access providers (both mobile and fixed) are likely to continue to offer content and services to their customers, bundled with broadband access. These services are often provided with guaranteed QoS (e.g. IPTV services). These services are not necessarily anti-competitive *if* the end-user can access the wider Internet and choose to consume other content. The critical competition issues are:

- Access providers who provide bundled services could be motivated to degrade content services or applications which compete with their own portal services
- If they do this, they will not be incentivized to tell their customers that the QoS for these services is inferior
- They could use this to leverage payment from content and applications providers
- Access providers may agree preferential arrangements with some content or applications providers but not make the same terms available to others

It is clear that discrimination and other forms of quality control are exercisable in a manner which does not fully support open access to content. If networks and commercial content providers cannot monetize their respective parts of the value chain, network effects can reverse into a vicious circle, in which neither content nor network can secure investment to provide service.¹³ Instead, the inflexion points at which investment in the lagging element is needed to prime the next phase of disruptive growth can become crisis points. At this point, investments may be constrained and a ‘virtuous circle’ of

investment may be replaced by a vicious circle of under-investment. As users currently display relatively little apparent motivation to price discrimination in order to gain greater bandwidth, knowing that the extra bandwidth is only in the 'last mile' and does not necessarily result in higher speeds for their favourite service, the incentives for end-users to signal willingness to pay for greater service may be weak. Again, there is an information problem, with ISPs unwilling to demonstrate clearly the practical advantages of advertised speeds of for instance 'up to 8/16/24Mbps'. This investment conundrum is claimed by some ISPs as a justification for traffic management and price discrimination.

Problems in the 'Middle Mile'

The Internet used to claim a cost-free transport model – effectively free peering, in that it was assumed that ISPs sent the same amount of information over a link as they received. However, as ISPs specialized in consumer or business customers through the 1990s, it became obvious that their traffic flows were diverging, and that many consumer ISPs, for instance, received much more information than they sent. Consumers are assumed to act in a mainly passive mode, receiving web pages and email in response to short (in packet terms) requests, which is why asymmetric connections are provided. Your bandwidth is predicated on receiving 10 times more packets than you send, and your speed of connection is based on that number. That is one of the reasons why consumers who act as 'broadcasters', by sending more information than they receive, are difficult for network engineering: they are not 'supposed' to act like that. The reasons for sending much more information vary, but uploading a lot of files, acting as a P2P exchange, or hosting a popular website on your local connection, can all cause more traffic to flow out than in. When it does, that Internet connection behaves in a manner which is more like a business connection than one for consumers. A real business dedicated 'leased line' (or 'partial circuit' depending on the terminology) is a vastly more expensive link than a consumer broadband connection, perhaps 50 times more expensive. It is not difficult to see why ISPs prefer consumers to act like consumers when they use their network.

With many thousands of ISPs, and the demand for peering much greater than many of the larger (so-called 'Tier 1') ISPs could meet, smaller ISPs began engaging in one of two practices. They either pay for transport, directly or via a transit arrangement with a large ISP that is able to peer with Tier

1 providers, or they collaborate by using an Internet Exchange, such as the London Internet Exchange (LINX). In this way, the very largest ISPs continue to peer, and certain smaller ISPs might do the same thing, but much of the traffic on the Internet ‘pays its way’ through the network. Even until recently, many analysts did not fully appreciate that so much traffic on the Internet was monetized and had to pay its way. Brett Glass, owner-founder of Wyoming-based ISP Lariat Networks, states that pricing and regulating overtly anti-competitive behaviours include ‘price squeezing, price gouging in the “middle mile” by large ISPs and refusal to deal by backbone owners’.¹⁴ Small ISPs regularly rehearse this complaint. The Internet is, for the most part, therefore not free for most traffic: all but the very largest ISPs charge for traffic to be carried over their networks, and this really matters.

The average Internet packet ‘hops’ across a dozen or more network routers before reaching its intended recipient. All the delays at each router add up, so that if the packet really does take twenty hops with an average 10 ms delay, it is a meaningful – for video and computer gaming in real time, actually disastrous – 200 ms. For several other time-sensitive applications, it could be a fatal delay.¹⁵ For instance two-way voice communication, as in a telephone call, can break down with delay, as can streaming video, as can interactive gaming. That is why most rich media applications of this type include ‘buffering’ as a means of protecting against delay. De Beer reminds us that: ‘No internet transmission is truly instantaneous, so the issue of time delay is really a relative question. And traffic shaping necessarily requires an implicit or explicit judgment about the relative time-sensitivity of the delayed telecommunication.’¹⁶

The five most common types of delay are ‘dropped packets’, delay, jitter, out-of-order delivery and error. Dropped packets might occur where the routers fail to deliver (drop) packets if their buffers are already full, as a result the receiving application may ask for this information to be retransmitted, possibly causing severe delays in the overall transmission. There is straightforward delay where the packet is held up in queues, or takes a circuitous route to avoid congestion, which can make useless an application such as Skype or online gaming. Jitter occurs where packets from the same source reach the destination with different delays, which can vary unpredictably (depending on the queues at the routers involved for each packet). This is critical to the quality of streaming audio and video, whereas buffering can cope with delay. In out-of-order delivery the packets arrive in a different order, which

requires special reordering protocols, again very important for video and VoIP streams. Error takes place when packets are misdirected, combined together or corrupted en route. The receiver has to detect this and ask the sender to repeat itself, as with dropped packets. The list of delays can be terminal for some applications, and this explains why rich media, especially video, transport is unreliable using the Internet.

It is important at this stage to discuss the sceptical view, and not only because, at least in relation to today's consumer Internet, I largely agree with it. The primary voice of caution is Odlyzko,¹⁷ not least because he called out WorldCom's nonsensical claim that Internet traffic was doubling every hundred days in the late 1990s (it was always a stock market-scaring ruse to increase the stock price of the ultimately criminally doomed executives). Odlyzko continues to worry that ISPs will waste resources trying to become cable companies as in the late 1990s, estimating that carriers do not need added revenue from price discrimination to finance residential broadband infrastructure.¹⁸ Odlyzko argues that carriers have systematically over-valued one-way entertainment including video over two-way services (telephony, instant messaging, texting or social networking sites) that use less bandwidth. Expanding networks to handle a premium video ZettaFlood would be a mistake. He also argues that video should use P2P distribution and other efficient methods, rather than replicate broadcast with real-time streaming. Criticizing the study by Clarke,¹⁹ Odlyzko states 'resists giving broadband operators leeway for price discrimination (except by offering consumers different connection speeds) or increased ability to vertically integrate into content or applications to facilitate price discrimination.' He believes that the operators' own shareholders may benefit from network neutrality rules that would limit such wasteful ventures and force them to stick to their core competence – providing basic connectivity. He also identifies the potential for price discrimination and discrimination against network neutrality at any Internet bottleneck, for instance via Google, the leader in online search advertising.

I explained to the European Commission in a study on the potential future of video consumption²⁰ that the great majority of the population will be 'time shifting linear consumers', who use broadband mobile and in-home devices, time-shift their media to suit their schedule instead of that of the broadcaster and fundamentally alter the media landscape. Either consumers choose to 'stream' linear video at alternative on-demand schedules, or as non-linear

interactive users, they adopt an entirely 'pull'-based behaviour, in which content is downloaded as if in a peer-to-peer music file-sharing network, for offline consumption on demand. We are no closer to understanding whether they stream (try to experience a real-time broadcast, subject to buffering) or download content for viewing later, than we were four years ago.²¹ What remains true is that the difference is largely artificial, because streamed content is actually buffered and played with delay, so it is essentially a fast download with less delay, not a 'live' programme. Try streaming live radio programming on the Internet alongside the broadcast version: it will rapidly become obvious which is broadcast and which is cleverly disguised downloading for later consumption.

Odlyzko is joined in his sanguine analysis of the prospects for high QoS streamed video and other one-way traffic by a 2009 study for UK regulatory Ofcom by Analysys Masons, the pre-eminent European telecoms consultancy, who conclude that the local exchange and currently commercially available access technologies can cope with the increase in bandwidth required, without the need for QoS to discriminate between traffic types, but that backhaul to the Internet backbone will have to become significantly more efficient:

Even with aggressive predictions of user demand, the cost implications may not be excessive if future [backhaul] pricing benefits from the economies of scale that fibre-based backhaul products offer unbundlers today, and if there are technological advances that further reduce the [bandwidth] cost.²²

In other words, there is sufficient potential capacity in the new IP networks to deal with more video traffic, but that capacity is currently too expensive for ISPs to access. If it becomes cheaper, due to efficiency savings (and regulated pricing), then video can be supplied cost-effectively to the local exchange. Once in that exchange, it can be delivered using upgraded local connections, either fibre or new generations of cable, DSL (telecom) or mobile/wireless links. The possibility of net neutrality regulation preventing some kind of fast-lane access is discussed: 'innovative business models might be limited by regulation: in the event that the ability to develop and deploy novel approaches was restricted by new regulation, this might curb the potential for growth in online video services.' Analysys Masons also provide some estimates of historical and future growth in bandwidth provisioning, based

on their confidential discussions in preparing their report. They estimate that from 2001 to 2008, ISPs increased their peak-time bandwidth per household from 4 to 24 kbps, a 23% annual growth rate.²³ They estimate that in future the growth of peak-time video watching may increase this from 24 up to 300–2,000 kbps by 2018. This would be a major, indeed exponential, increase in cost for ISPs, but is related to backhaul, not local access. Note that ISPs must aggregate all their local exchange customers, so 24 kbps for 1,000 users would indicate a theoretical minimum of 24 Mbps required at the local exchange, a significant ‘pipe’ for the Internet data required. Analysys Masons estimates that the average UK household used 5.25 GB of data per month, based on industry figures.²⁴ Those used to Internet backbone prices using dark fibre should note that the price of a 155 Mbps BT central circuit was in 2008 almost £200,000. Regulated local access is a very different pricing environment.

In the United Kingdom in mid-June 2009, this question of backhaul cost was thrown into public view when BT acknowledged that it had been throttling BBC services as well as Google’s YouTube.²⁵ Under the ‘fair use’ policy of BT’s Option 1 broadband package, BT cuts video streaming from 8 Mbps down to 896 kbps between 5 pm and midnight, which BT said was sufficient to watch BBC iPlayer, but obviously nothing else concurrently. John Petter of BT Retail stated that:

We can’t give the content providers a completely free ride and continue to give customers the [service] they want at the price they expect ... [its] much bigger issue than the BBC iPlayer, it’s true of all forms of video content coming across the web. It’s becoming a more and more pressing issue ... If it wasn’t a significant sum, we wouldn’t be focused on it.²⁶

This caused much speculation as to the transfer-pricing within BT, as their wholesale executives had claimed repeatedly that there was plenty of capacity in their network. I quoted Matt Beal in the introduction, for instance. However, if BT Retail is paying the same regulated price as other ISPs (which under regulation, it must be) then the cost structure appears to be challenged. Burstein, a continual critic of ISP throttling, stated in response that:²⁷

the entire problem would disappear if Ian Livingston [BT CEO] actually asked his network people ‘how much does all this bandwidth really cost?’ The answer will be 2–5% of our charge, going down as often as up. The

right answer, based on data from a dozen large carriers, will be ‘well under a pound/month/subscriber’. One particularly efficient European, not far in size from BT Retail, has told me their figure is 55 cents.

If BBC and Google will be forced to pay excessively for access to the end-customer by an ISP insisting on QoS (or just negative neutrality throttling), do they have choices? Let us explore local caching options.

How to Avoid the Middle Mile

Note two means by which networks currently manage traffic: caching close to the end-user to save on costs and service delays on the wider network, and capping bandwidth for end-users. First, consider CDNs. The Wall Street Journal on 16 December 2008 claimed Google wanted to hire faster capacity, when it actually wanted to try to cache more locally, to push content to the edge.²⁸ Saul Hansell explained in the wake of the article:

Sometimes the issue is framed as a total bandwidth egalitarianism, when that’s not really what they want. There is a huge fight here, not over whether there will be first class and coach seats, but how those seats will be priced and who will pay for them. Google and others are saying that, in effect, every seat in the same class of service should have the same price, and that Internet providers can’t add surcharges to companies they don’t like or give discounts to those they do.²⁹

Do caches breach net neutrality? Kevin Baughan from Virgin Media (the UK cable provider) recently stated that Virgin is experimenting with both edge/network caching, and in the Ashford 50 Mbps trial area, provisioned a 10 Gbps link directly to Level3 (a Tier 1 backbone provider), in an attempt to answer the question of whether storage trumps transmission, or the reverse. One intriguing idea he floated was that perhaps ISPs should be building their own internal CDN capabilities, striking deals with other CDNs and content players to mirror the most popular content.³⁰ Virgin already has taken its two major sources of network congestion offline, by hosting BBC iPlayer (catch-up TV) services on its cable channels, and by aggressively targeting BitTorrent and other P2P traffic via DPI.³¹

Two elements of this are important – if Akamai is an ‘SME solution’, based on the idea that the big content providers like Yahoo!, Google and BBC

‘Project Canvas’ install their own servers with ISPs, then does local caching by proprietary networks solution fall foul of net neutrality? First, Akamai’s CDN is in an increasingly competitive market. Second, should there be FRAND regulation of these caches? The idea is that if Google can build such servers then so should everyone else, if they wish. This should be paid for by content providers to ensure equality: Disney can’t persuade BT to pay it for the content or the server. This is a well-known principle in for instance LLU regulation, so use it for caches too. The non-profit content providers such as Wikipedia (and you) will continue to rely on the open Internet and usual speeds, so that that does not remove the need for net neutrality regulation at all.

Second, ISPs can simply cap the total monthly capacity for its customers. That does not affect their peak-time usage and congestion problems, but does ensure that ISPs save money on backhauling the content from the open Internet to the local telephone exchange. Time Warner Cable was an early leader in capping customers, in its Beaumont Texas trial allowing 5–40 GB caps with \$1 per GB excess fee.³² The pricing plans ranged from \$29.95 a month for a 5 GB cap and 768 kbps download speeds to \$54.90 for a 40 GB cap at 15 Mbps. AT&T was set to impose bandwidth caps in Reno, Nevada of between 20 and 150 GB for new customers in 2008, extending to existing customers in 2009. It was relatively transparent about its plans, and its announcement resulted from discussions with FCC Chair Kevin Martin. However, Burstein wrote:

I’m one of the few who said Comcast’s 250G cap was reasonable in 2008, because there is a cost of bandwidth. But every policymaker in D.C. knows that the 20, 40, and 80 gig caps are a blatant attempt to discourage access to content of your choice.³³

The question of what level of monthly cap is acceptable is one that regulators must grapple with, especially as mobiles currently offer caps as low as 1 GB. Moreover, caps must suggest a decision not to prioritize such services, as high-priority service would mean large volumes of video being downloaded which would make the cap redundant (or dangerously low). If you download a series of ‘House’ in HD video, a 20 GB cap would make no sense.

There is unquestionably a regulatory issue that is vital in the delay of media transmissions, whether you agree that it is middle or last mile. It will affect the content provider, who should at this point be introduced to what has thus

far been a rather dry technical discussion. Decisions regarding traffic and QoS on the backbone are intimately connected to the experience enjoyed by the end-user, so it is properly the place of regulation, however 'light touch'. As De Beer puts it:³⁴

Access providers should not be empowered to exert influence over citizen's ability to telecommunicate by making judgments about which information from which sources is transmitted more or less quickly. That is unacceptable in any free and democratic society.

He further identifies the core of the problem as ignored by the Canadian regulator CRTC:

... control over content could be exercised other than editorially, but the Commission did not seem to consider that possibility. A distinction was drawn between controlling the speed of telecommunications and the content itself. The Commission held that simply delaying a telecommunication does not alter its meaning or purpose. That is also an odd conclusion to reach, because a delay could very much affect the purpose of communicating via the internet.

As we saw with BT's recent action, throttling speeds beyond about 800 kbps for BBC iPlayer would mean that it could no longer be properly accessed during peak time. Would that appear to be network rules making a decision for BBC licence fee payers? If rules are introduced for the middle mile, they will need careful drafting. Faratin, Clark et al. state:

the design of such a rule will be both complex and informationally demanding... While there may be opportunities for abuse by providers with excessive bargaining power, the complexity of what is in place today, and what seems to be working today, would argue that the best way to address any potential concern would be to focus on the sources of bargaining power and identify anti-competitive opportunism, rather than to impose ex ante restrictions on the range of bilateral contracts.³⁵

They plump for more transparency, unsurprising given the obscurity of many Internet peering arrangements, and the lack of regulatory oversight. Indeed,

it can be argued that only the Bell mergers of 2006, leading on from the WorldCom and AOL mergers of 1999–2001, really produce an opportunity to examine these arrangements. The type of ‘emergency regulation’ that is possible in merger cases is ill-suited to the careful balancing act that may be needed to ensure competition is fair and effective across Internet peering arrangements. They state:

If it were possible to bring a ‘best practice’ or ‘common practice’ in interconnection out from the non-disclosure agreement and into the light, this might also help reduce bargaining costs, but in a more flexible way than might be achieved via regulatory constraints. An industry forum that tried to discuss this openly (and which was given a clear mandate for how to behave so as to avoid anti-trust concerns) might offer a substantial contribution to efficient operation of this asymmetric world, and might mitigate the sorts of fears that have prompted calls for more direct regulation.

I agree with their final warning: ‘Regulating the middle mile will be a difficult and time-consuming business that will require great skill by regulators in ensuring the right degree of co-regulation and industry involvement.’ So far, regulators have shown no sign of appetite to enter into examination of this critical arena for negotiation and interconnection. They will need to do so to have real understanding of middle-mile problems.

Public Service Broadcasters

The net neutrality/QoS argument in Europe is particular because PSB occupies a position of strong bargaining power with legislatures and regulators. What would happen if the PSB and the monopoly carved up the market? What if the incumbent refused to allow video content from the PSB? The former may be quite likely, the latter in several countries unthinkable. The argument can therefore be characterized as: will Net Neutrality apply only to PSBs, or to other/all content providers? I consider a different set of content provider preferential treatments below. The doyen of PSBs, the BBC, formalized its position in February 2008:

The BBC considers that [proposed EC] rules would be sufficient for the time being. The key issues are transparency, the capability to identify problems promptly, and the ability to impose remedies in a timely

way ... NRAs should require service providers to disclose to consumers promptly and in full the 'blocking' or any significant manipulation of any services ...

As a public service broadcaster, the BBC depends on the internet as on other distribution platforms for access to its viewers. This access must be unhindered ... If the internet moves away from non-discriminatory treatment of all content and broad neutrality, the resulting environment will undermine the EU goals of developing creative content online in the single market and potentially threaten cultural diversity and media pluralism.

The BBC then suggests what I term a co-regulatory solution:

we would suggest that the Commission should consider stimulating (or facilitating) a stakeholder forum, possibly involving NRAs as well as industry, to address issues related to non-discriminatory access to the internet; for example: transparency about service management at user level; best practice exchange about prompt identification of abuses; monitoring impact of application of EU rules in remedying abuses.³⁶

The BBC began streaming television content as a mainstream proposition to the Internet in 2006.³⁷ It created huge controversy as it encoded content at 800 kbps and that content was subject to extreme peak-time congestion: it was live coverage of the European Championships football tournament. The BBC helpfully informed consumers that if they could not access the content properly on their broadband connections, they should contact their ISPs and suggest specific traffic management changes! The ISPs' response was predictable: they began blocking BBC services. The BBC then launched a service called iPlayer on a P2P service, which was controversial not only for creating large peak-time congestion problems, but also because it was initially only available to Windows Media Player users, enraging various Apple and Linux home users. Despite these controversies, the BBC dedication to streaming programming to all licence fee payers has continued.

James Enck reports³⁸ that Anthony Rose, BBC Head of Digital Media Technology, stated his hope that the industry could now move on to consider issues around ISP incentivization and monetization. The iPlayer team resisted playing out HD content until 2009, due to concerns that the

experience would be unsatisfactory for many consumers, due to contention rates on DSL connections. HD content on the iPlayer is encoded at 4 Mbps. The minimum threshold bitrate for true HD is above 3 Mbps, which would be challenging for a lot of broadband connections and would risk high buffering levels. There is a necessity to ensure the experience is comparable with what HD TV viewers have come to expect.

Rose made a case for the BBC assembling the puzzle pieces for ‘others’ (presumably ISPs) to build a business model around the iPlayer. He expressed an interest in working with ISPs to develop tiered service offerings to more closely align costs with revenues, as well as to cooperate on technology-based strategies to alleviate pressure on networks. The BBC suggests tiered pricing for live TV, as the obvious solution to bandwidth constraints on the iPlayer service – ISPs should charge end-users more. Rose stated that:

The future lies in tiered services. What we need to do is to create the iPlayer services at different quality levels and then let ISPs offer different bandwidth propositions to users... Of course, nobody should get a worse experience than today. For example, the user can get a good quality iPlayer service for, say, £10 a month but for £20, a much better iPlayer quality would be available. This can lead to win-win situations and ISPs will see video services as a profit centre rather than a cost burden.³⁹

Video is as much a ‘bandwidth hog’ in the United Kingdom as elsewhere, and the free availability to BBC licence holders of the company’s video offerings from early 2007 resulted in serious peak-time congestion issues where several subscribers in each exchange were watching the BBC simultaneously via the iPlayer. Ofcom had carried out a Market Impact Assessment (accompanying the BBC Trust’s ‘Public Value Test’) on the iPlayer idea, to estimate its impact on ISP traffic.⁴⁰ The later joint distribution network between BBC Worldwide Ltd (the BBC’s commercial arm), Channel 4 and ITV, ‘Project Kangaroo’,⁴¹ was referred by the Office of Fair Trading (OFT) to the Competition Commission.⁴² The concept behind Project Kangaroo was that the CDN for live broadcast video over the Internet be provided via a proprietary P2P client, which would be more efficient than using BitTorrent, for instance, and might provide more ISP cooperation in its distribution. It was effectively vetoed by the Competition Commission following investigation over 2008–9, stating

that it would result in ‘substantial lessening of competition in the supply of UK TV VOD content at the wholesale and retail levels’, and the PSBs gave five-year undertakings not to replicate its activities.⁴³ The Commission on 4 February 2009 had concluded that: ‘a relevant merger situation would be created if the Joint Venture proceeded; Kangaroo would be likely to lead to a substantial lessening of competition in the supply of UK TV VOD content at the wholesale and retail levels; prohibition of Kangaroo was the only remedy that would address the issue.’ As a result, the BBC has embarked on exploration of non-commercial alternatives paid within the licence fee funded part of the BBC, code-named ‘Project Canvas’. Canvas will supply catch-up programming via digital terrestrial and satellite TV, and will be open to any other operator on FRAND terms, avoiding the competition problems that scuppered Kangaroo.⁴⁴

Note that the political economy of UK communications is such that it is unthinkable that PSB streams cost extra for the BBC and Channel 4 to distribute with ISP-induced QoS, though the commercial content might be more debatable. Both are funded by a type of universal service tax and other government subsidies (in kind and otherwise) paid directly by the taxpayer. For ISPs to try to breach net neutrality and charge extra or throttle this service would be to deny the taxpayer their public service content. ISPs effectively have to swallow the pain of distributing the content, and therefore there is a radically different environment than, for instance, in the United States.

Commercial Content Providers and Net Neutrality

Let us accept that the ISPs may not be the bad guys. Just as we will see that copyright holders, security agencies, police forces, libel lawyers, anti-paedophilia campaigners and others are pushing ISPs to regulate content that passes over their networks, so content providers may try to make ISPs pay for premium content that their subscribers desire. The street is two-way: net neutrality is as likely to be circumvented by content owners – and indeed governments – as by ISPs. If a provider of live football, baseball or cricket tells ISPs that they must bid for exclusive live rights, just as TV networks do today, the winning provider will presumably enter into agreements for distribution and billing to make that arrangement as agreeable to the end-user as possible. That includes the end-user who is not a baseball fan and does not want her local loop congested by the baseball backhaul. In that case, would an ISP arrange to host as much content as

possible on local servers or even to broadcast to that server by satellite to avoid backhaul congestion? This is no fairytale: ESPN has secured carriage agreements on 60 US ISPs,⁴⁵ while several European incumbent ISPs have signed exclusive sports rights deals. Small US ISPs have claimed that discriminatory treatment should be outlawed⁴⁶ – a judo move on net neutrality turning it against the major content provider! The cable model is arriving, and net neutrality is being eroded often to the largest content providers' benefit and upon their strategic decisions. Is this a major problem for net neutrality if it boosts the size of pipes for everyone? I would argue that it is not, but will it really boost everyone's pipes, or let you eat your neighbour's congestion?

Take an example from the United Kingdom. Pay-TV operators are in a constant bargaining game with content providers, in which their relative strength depends on the competition between the operators' platforms and the substitutability of the content offer for the end-user. If you need Premier League football and your platform has not struck a deal, then you will not subscribe to that platform, demonstrating the insubstitutability of the product and the bargaining powers involved. It is complicated where the content offer has been purchased by a rival platform, as when Sky TV buys the Premier League highlights. In this case, rival platforms complain that this vertically integrated provider is in a position to exploit its position to their cost. Hence, Sky TV has had disputes recently with BT, Virgin Cable and ITV, in all cases over the supply of premium programming and Sky's perceived domination of the market. The BT dispute was the sole 'Internet' TV dispute and it is worth examining.

BT is obliged to supply backhaul at the same cost to its retail arm as to other ISPs, which means that IPTV is either an expensive option or a non-interactive option (to butcher the definition). BT broadcasts digital terrestrial TV to a set-top box with a hard drive, enabling users to record programmes as required and download specifically requested premium content. It is an ingenious and bandwidth sensitive method of delivering 'near-IPTV'. It is close to a linear continuum, which given the price its own wholesale arm charges is just as well.

BT is unable to negotiate terms for a carriage agreement of Premier League football – famously, it has been said of Englishmen that they follow three main sports: 'football, football and football'. Moreover, Rupert Murdoch described football as the 'battering ram' of his pay-TV business in the early 1990s. It will not be BT's battering ram, not least because Murdoch's Sky owns

the rights (News Corporation owns 36% of Sky, giving it effective control) and it has not yet negotiated a terrestrial deal, let alone for BT's service. BT argues that because Sky is dominant in pay-TV, with a master-and-servant relationship with cable, that therefore BT is caught in regulatory asymmetry: Sky can unbundle BT's local loop for its highly successful Sky Broadband service, but BT cannot access Sky's equally dominant pay-TV content for its platform, because content is not regulated under EC competition law, only under UK law.⁴⁷ Briefly, this is because culture is a national competence jealously protected from EC intervention, and the EC regulates the electronic communications networks, not the pay-TV content. This will become very messy when net neutrality comes of age – is it a content or a carriage issue? De Beer is pulling on these tangled threads in the Canadian context, as Gibbons and Woods have done in the United Kingdom.

Content providers with market power can carry out three actions that impede consumer enjoyment of their broadband connection:

1. Refuse to license their technology and prosecute or persecute those who put the content online vicariously – this has been the past decade's policy
2. Cut exclusive deals with pay-TV operators or ISPs for all content, thus preventing the majority of users from accessing the content at any price
3. Distribute the content without regard for network traffic management, thus flooding the local loop and preventing users' enjoyment of their or anybody else's content during peak hours, if not continuously.

Note that each action is motivated by entirely different beliefs and that the third action contributes to net neutrality but not to realistic settlement of the issue. Pipes are narrow and end-users have to be rationed, if not forever, certainly now. The flooding of the ISP network by the BBC iPlayer content is not sustainable in the United Kingdom, and it has acknowledged this by supplying iPlayer content to Virgin Cable by using a TV channel, not the Internet connection. Sensible decisions should be made to ensure that video content reaches end-users without flooding the capacity leased out to all local loop ISPs that are not cable. If not, cable wins by virtue of not having unbundled competitors (a logic employed by telcos in BandX and other decisions in 2003–5 to secure equal deregulatory treatment).

The second option, which is pursued by several European ISPs and content carriers as well as ESPN in the United States, is standard premium content/cable TV negotiation. Is it better to guarantee bandwidth and limit audience

to one network while gaining revenue from an exclusive deal? That is a decision for the content provider, not the regulator, as long as that provider is not dominant. We are in for an extended period in which the business rules for premium content online will be established and it is not my intention to suggest that net neutrality should do more than make clear that the widest possible access is the most advantageous policy for consumer welfare.

Conclusion: Termination Fees for Content Providers

Since broadband ISPs have a termination monopoly or duopoly⁴⁸ over the end-user, they can use that to charge termination fees to those who wish to get access to the user. This behaviour is familiar to the cable TV industry, where only large content providers can secure free or even profitable carriage, whereas smaller content providers with less contracting power are forced to pay the cable TV operator for access. The fear is that a similar model will be imposed on the Internet, where only large content providers with sufficient negotiating power and those with political influence to secure favourable carriage terms will secure free carriage. The argument in Europe is particularly pernicious because PSB occupies a position of strong bargaining power with legislatures and regulators. The argument can therefore be characterized as whether net neutrality applies only to PSBs or to other/all content providers? The street is two-way: net neutrality is as likely to be circumvented by content owners – and indeed governments – as by ISPs. In Chapter 2 in the case of governments, and in Chapter 4 in the case of copyright and child protection policy, it is discussed that ISPs are under constant pressure to investigate their traffic and that this threatens to introduce liabilities (and potential commercial benefits) to their business model, with drastic implications for net neutrality.