Chapter 6.
Traffic Management Decisions in the United Kingdom

6.1 Introduction

In contrast to the US, application-specific traffic management was pervasive in the United Kingdom during the 2000s. The UK's earliest broadband providers adopted traffic management techniques aimed at controlling peer-to-peer and other non-web applications when broadband was still a nascent product with few subscribers, and in the years that followed, those techniques became entrenched and diffused throughout the industry. By 2010, more than 75% of UK residential subscribers were subject to some form of application-specific traffic management. Many of the techniques employed were blunt instruments that severely constrained the performance of the affected applications for large portions of the day, as opposed to targeting specific instances of instantaneous congestion. Even as network conditions improved or evolved, application-specific traffic management techniques often did not.

The most common form of application-specific management observed in the UK – rate limiting peer-to-peer file-sharing traffic – was initially adopted for the combined benefits of improving performance and controlling variable costs (predominantly backhaul capacity costs). The significant competition that ensued after Ofcom and BT took steps to make local loop unbundling viable had two effects that reinforced the utility of this approach: a significant decline in retail prices, and mounting pressure to increase broadband speeds, which required further investment. To meet marketplace demand for increasing speeds at decreasing prices, operators were incentivized to do whatever they could to keep costs down, including rate limiting high-volume applications.

Throughout the decade, Ofcom's actions and rhetoric created a culture of flexibility around traffic management. Ofcom was infused from its creation with the culture of a competition regulator and its embrace of functional separation was one of the most interventionist steps it

could have taken to stimulate broadband competition. Having taken that step, it demonstrated a deep commitment to the disciplining power of competition in the marketplace, including in its policy formulation around net neutrality and traffic management. At the same time, Ofcom officials were adamant that the management of peer-to-peer applications was essential for networks to function. This logic gave operators nearly as much freedom as they could have desired in managing applications as they saw fit.

Consumer concerns about application-specific traffic management (to the extent that they existed) were muted, in part because of the perceived association between peer-to-peer applications and illegal copyright infringement. Affected users would complain or inquire about their operators' traffic management practices in online forums, but these concerns rarely garnered the attention of media outlets, consumer groups, or Ofcom.

Because neither consumers, nor Ofcom, nor the industry drew attention to application-specific traffic management, during the mid-to-late 2000s there was no nuanced public debate in the UK about the relative merits of application-specific and application-agnostic approaches or about whether the approaches in place could or should be more targeted or consumer-friendly. As a result, the relatively blunt approaches to managing applications that some operators adopted early on became ossified, hardly evolving over the decade despite massive evolution in network performance and speeds.

This chapter analyzes how the UK's wholesale market structure, the intense effects of competition, and a regulatory culture of permissibility combined to foster a British broadband landscape in which application-specific traffic management became pervasive. Marketplace dynamics bolstered operators' incentives to adopt or perpetuate application-specific management while countervailing public pressure never materialized, yielding a vastly different traffic management experience than that of the United States.

6.2 Early Adoption of Application-Specific Management

Application-specific traffic management was a fixture in the UK dating back to the earliest years of broadband itself. Numerous early broadband operators turned to application-specific traffic management strategies that relied on port-based application classification or deep packet inspection equipment to identify and manage high-traffic applications on the network. Application-specific management was adopted as part of larger strategies that often involved application-agnostic approaches as well, most commonly tiers of service with explicit volume caps or fair usage policies (FUPs) that prescribed limits for how much data customers could consume.

Peer-to-peer applications were the most common and widespread applications targeted, with operators limiting the capacity available to them either on a per-user basis or in the aggregate (for example, allowing the applications to consume no more than 1% of the network's overall peak bandwidth). Other non-web applications, including newsgroups, were at times subject to rate limits as well. Providers commonly imposed these limits across their entire user base, and at times did so with minimal or buried disclosure. Different providers imposed the limits over different time periods, with some applying them only during a few evening hours or at times of congestion, and others applying them for longer stretches of the day or at all times. Some operators, most notably Plusnet, adopted more complex application-based prioritization schemes, offering broadband packages where many more application classes were identified (web, email, VoIP, gaming, etc.) and prioritized according to their time-sensitivity or other factors.

This section explores the key drivers for adoption of application-specific management in the early years of broadband, prior to the fundamental market changes that accompanied the roll out of local loop unbundling (LLU) beginning in 2005 and 2006. It analyzes the most common case – DSL operators' management of high-volume applications – and contrasts that approach to those of the cable operators and Plusnet, whose early strategies helped to shape

the traffic management landscape for the remainder of the decade. This section shows how application-specific management commonly met DSL operators' needs to control both performance and costs while incurring little attention from the public, the press, or Ofcom.

6.2.1 Wholesale Network Model for DSL

Broadband Internet service arrived in the UK around the turn of the 21st century. Early broadband was predominantly provided by cable operators and the retail arm of BT. With the launch of BT's wholesale bitstream products in 2000, bitstream operators proliferated, led by a number of companies that would all eventually come under the TalkTalk brand (Carphone Warehouse, Tiscali, Pipex, Nildram, and numerous others). Because the cable networks' footprint is geographically limited, DSL has been the dominant form of broadband access since 2003 (Ofcom 2009b).

Backhaul Prices

There are several features of the bitstream cost model that played an important role in operators' decisions to adopt both application-specific and application-agnostic traffic management in the first place. Figure 3 shows the key components of a typical bitstream operator's network (and BT Retail's network). Bitstream operators would contract with BT's wholesale division to run their customers' access lines. Although these access network costs typically accounted for more than 50% of the total monthly cost of each customer's line, the costs were bandwidth-independent. Access links are provisioned at a specific line rate per customer, and fluctuations in bandwidth usage up to that limit based on each customer's activity do not incur any additional charges. Nor do such fluctuations impact other customers in the neighborhood since each customer is connected to the exchange on a dedicated access line.

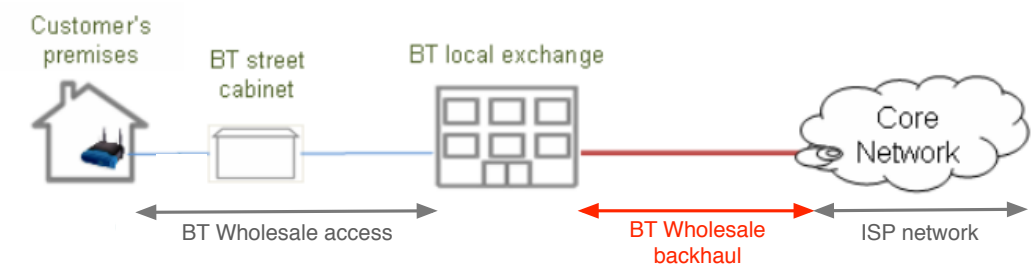

Figure 3. Typical UK bitstream ISP network in the early 2000s.

The backhaul component, in contrast, is both bandwidth-dependent and shared among many subscribers. Bitstream operators would purchase a certain amount of capacity from BT Wholesale to connect their customers from the local exchange to their own networks and on to the rest of the Internet. The amount of capacity purchased was based on a number of factors, including available budget, network utilization, and performance targets. If peak time bandwidth demand exceeded available capacity, operators could either purchase more bandwidth or allow performance to suffer as each customer's traffic competed for space on constrained links. In describing an early broadband network prior to the adoption of traffic management tools, one network engineer explained that "congestion was what actually controlled the amount of bandwidth and the experience the customer had."

The most popular wholesale products consumed by DSL operators in the early days were known as Datastream and IPStream with BT Central Plus. These products, whose prices were regulated by Ofcom based on a finding that BT had Significant Market Power in the wholesale market, were priced to average out the fixed and ongoing network costs that BT incurred in maintaining the nationwide wholesale network. As described by Plusnet's Dave Tomlinson (2008), UK operators saw a "huge price premium" for backhaul because the uniform national price was set to cover the cost of "links to some of the smallest exchanges in the country that perhaps only serve a couple of hundred customers, some of which are probably very uneconomical." Backhaul bandwidth was generally the most expensive bandwidth segment on UK operators' networks. Indeed, in the mid-2000s, the price per Mbit/s of IPStream (in the £120-£200 range for most ISPs' networks) was five to ten times

the median price per Mbit/s that ISPs would pay for transit links that connected them with the rest of the Internet (Telegeography 2011).

The resulting backhaul prices also provided DSL operators with limited economies of scale in their purchase of capacity: a linear increase in capacity necessary to accommodate a growing traffic base was met with a roughly linear increase in the price of backhaul. For those using the IPStream product, for example, the price per Mbit/s when purchasing 622 Mbit/s of capacity on a BT Central Plus link was roughly the same as the price per Mbit/s when purchasing 155 Mbit/s on the same link once the annual rental for the link had been paid (BT 2006). Thus as traffic growth outpaced customer growth, operators were faced with backhaul costs that rose faster than revenues from customer subscriptions.

Retail Networks' Lack of Control

The fact that the wholesale network was run by a distinct entity from the retail networks also reduced the control that bitstream operators had over their networks in a number of ways. The lead time for deploying additional backhaul could be in the three-to-four month range, meaning that if operators' quarterly forecasts for traffic growth were too low, or if the process of getting new capacity installed incurred extra delays, there was no way to augment capacity on short order to mitigate performance problems. The kinds of "emergency" capacity increases described by US cable operators could not be deployed by UK bitstream operators.

UK ISPs were unabashed about the drawbacks of being beholden to BT's upgrade schedule. For example, the bargain bitstream operator FairADSL was explicit in explaining to its customers the source of performance problems on its network in late 2002:

> [W]e placed upgrade orders with both BT and our outgoing bandwidth providers in plenty of time to keep up with demand. Unfortunately due to ineptitude and bureaucracy on behalf of these parties, the bandwidth upgrade did not operate correctly and the BT Central Pipe upgrade has not been implemented in time. We are very annoyed about this, and extremely angry on behalf of our customers. (Jackson 2002)

ISPs felt the need to provide such explanations because the performance implications of delayed capacity upgrades or incorrect bandwidth demand predictions were obvious to subscribers. Their application performance suffered as a result. This issue was exemplified by the experience of the bitstream operator Nildram in January 2003, as reported on the ThinkBroadband blog (then known as ADSLGuide):

> Nildram, a well known ADSL ISP . . . is experiencing difficulties with a growing user base and lack of capacity on their central pipes. In the past few days, users of latency-sensitive applications such as online gaming and telnet have experienced problems at peak times when the central pipes have been overloaded. Orders for two more 155 Mbps pipes have already been placed with BT some time ago and were due to be on service by December 17th but Nildram now expect them not to be available until January 17th at the earliest. (Lahtinen 2003)

Degradations in service quality brought on by delayed capacity upgrades were a recurring theme in these early years.

Furthermore, when congestion arose operators had little visibility into its causes on the wholesale network prior to the interconnection point with the operators' networks. Because of the wholesale/retail split, operators lacked visibility into what was taking place on that first segment of the network where multiple subscribers' traffic came together, as one engineer explained: "the reality was that if [customers] had problems, you had very little that you could do with them because . . . all you knew is that they went through the wholesale network and somewhere along the line as far as [we were] concerned they magically connected to the Internet." In sum, operators lacked control over both the costs of their networks and the performance their users experienced, with little ability to diagnose or solve either problem.

6.2.2 Traffic Management to Control Cost and Performance

These problems came to a head soon after broadband became widely available. Subscribership grew rapidly in the early years, with well over 50% growth in both DSL subscribership and broadband uptake overall each year between 2002 and 2005 (Ofcom 2009b). With more subscribers, and more high-traffic subscribers, network operators were seeing increased complaints from customers about the speeds they experienced at peak times.

Keeping up with demand required increased investment in network capacity, with operators at times doubling their entire network capacity in a single upgrade. But given the pace of growth and the backhaul pricing structure in place, operators were searching for mechanisms to control the amount of traffic on their networks without needing to provision for an increasingly high and difficult-to-predict peak traffic rate. Backhaul costs were becoming a growing concern, and "writing a blank check for demand" to BT Wholesale (as one engineer put it) was not seen to be a sustainable solution.

In assessing their options, operators observed that a small minority of heavy users were creating the majority of the traffic, as has been common not just in the UK but across European and international broadband networks (Mooyaart 2012; Sandvine 2010a). Because high traffic meant high cost, these heavy users became an important target for traffic management. Volume caps and fair usage policies were introduced to help operators control heavy users' traffic in coarse-grained ways. Some operators would send warning letters to users when they had exceeded daily or monthly volume caps, hoping to encourage them to moderate their usage. Others instituted overage fees, slowed users' speeds, or disconnected them altogether when they exceeded their volume limits, attempting to create even stronger incentives for heavy users to moderate their usage while also increasing revenue or reducing traffic-related costs in the process.

While these approaches were aimed at reducing the bandwidth demand of a voracious minority, they did not give operators more precise ways of controlling or predicting that demand. Even operators that adopted a daily volume cap that resulted in slower speeds when breached could not know if, when, or how many subscribers would reach the limit on any particular day. In some cases users could choose to face the consequences of exceeding a cap – receiving a letter or paying a fee – without changing their behavior. Thus even with these approaches in place, operators were still looking for more fine-grained ways of predicting and controlling the demand for capacity.

Operators were aware that peer-to-peer traffic was quickly becoming the dominant type of traffic on the network, representing more than 50% of network traffic in some cases, as reported by interviewees and elsewhere (Ferguson 2002; Lunden 2005; Mooyaart 2012; Sandvine 2010). Figure 4 provides further evidence of the trend, showing the ratio of upstream-to-downstream traffic on one large UK network since 2004. Because peer-to-peer traffic tends to be roughly symmetric, a large presence of peer-to-peer traffic on a network implies a higher upstream-to-downstream ratio than would be observed in situations where web browsing or streaming are the dominant application types. The above-50% ratios in 2004 and 2005 demonstrate the dominance of peer-to-peer traffic.

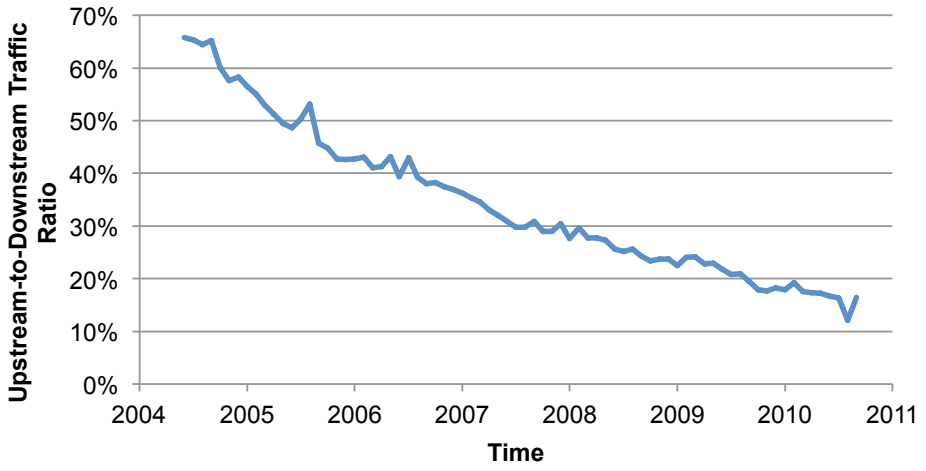

Figure 4. Monthly upstream-to-downstream traffic ratio on a large UK network. Reproduced from Cooper, Jacquet and Soppera (2011).

Using application-specific traffic management equipment to identify and control the capacity available to peer-to-peer file-sharing (and other high-volume applications such as newsgroups) seemed like a natural choice to operators under these circumstances. While application-agnostic approaches were broadly aimed at changing the behavior of a subset of customers, approaches focused on high-volume applications could very specifically and precisely stifle the prime sources of traffic on the network, giving network operators far more control over their bandwidth investments than they had previously. Management of peer-to-peer and newsgroup applications reduced the variability of bandwidth demand, as one

engineer explained: “You could basically make your network grow, if you want, at a much lower rate than it would grow otherwise, you could control your costs that actually were real costs. That’s why you have seen the retailers in the UK . . . the majority that are consuming BT Wholesale products have deployed that type of solution.” In a market where operators were beholden to BT Wholesale’s prices and upgrade schedule, these tactics gave them back some element of control.

But managing high-volume applications was not only about controlling the timing of bandwidth investments – it was also about reducing overall costs while safeguarding the user experience for the majority of subscribers. The use of application-specific traffic management as a way to reduce costs was a constantly recurring theme among the observed research team, and numerous interviewees were unequivocal about the effectiveness of peer-to-peer management as a cost reducer. The following is a sample of responses provided when interviewees were asked why this form of traffic management was deployed on ISPs’ networks:

- “You save real money.”
- “In order to save . . . from paying [BT Wholesale] loads of money in that interconnection point.”
- “To save money! (laughing) Of course! Peer-to-peer traffic is a lot of traffic.”

6.2.3 Operational Drivers of Application-Specific Management

While high-volume applications may have been prime contributors to congestion, it was the operators’ view that they did not suffer its effects in the same way as other applications, providing a further rationale for management. Congested networks can indiscriminately delay applications of all types whether or not they are interactive. Suppressing peer-to-peer traffic and newsgroups helped to reduce the likelihood that more popular applications like web browsing would suffer from congestion. As one operator explained, “if you didn’t apply some level of control, if we just allowed it to congest like we’d done previously, it was very clear

you could see browsing issues, you could see performance of VoIP issues," because those were the kinds of applications where users would notice the delay caused by increased congestion. The Nildram example provided above illustrates this as well: users of latency-sensitive applications were the ones to notice the effects of constrained capacity.

Peer-to-peer and newsgroup applications were also viewed by operators as less time-sensitive and interactive than other traffic. By limiting the capacity available to these applications at peak times, file transfers would be delayed, but they would still eventually complete. Interactive applications like web browsing, VoIP, and gaming, by contrast, were viewed as having limited delay constraints, otherwise they could become unusable.

As with application-agnostic approaches, approaches focused on high-volume applications were appealing to operators because they were likely to have detrimental effects on only a minority of heavy users, to the benefit of the majority of users. On many networks, peer-to-peer file-sharing followed a similar power law distribution to broadband usage overall, with only 5% or 10% of a network's subscribers creating the peer-to-peer traffic that accounted for the majority of traffic on the network. Thus application-specific management was viewed as a safeguard for the majority of users' experiences. As O2 later noted in a filing to Ofcom, operators that implemented these kinds of strategies did so "in the interests of all users (and society) even though this may inconvenience a tiny minority of users who are consuming a disproportionate volume of scarce resources" (Telefónica 2010, 10).

As in the US, DPI-based solutions were also attractive to operators for the insights they gave about the causes of network problems and growth. Prior to the deployment of DPI, some operators used port- or IP-based information to guess at how much traffic could be attributed to various applications, but DPI provided the additional benefit of granular data about individual application usage, in the aggregate or on a per-customer basis. DPI gave operators a platform for understanding which applications were driving growth on their networks, which applications heavier users tended to use compared to light users, and a variety of other

insights that could be used for capacity planning and design of new products (Cooper, Jacquet, and Soppera (2011) provide examples of such insights).

6.2.4 Alternative Approaches

Although DSL operators imposing limits on peer-to-peer applications was the most common type of application-specific management in the early years, alternative approaches pursued by Plusnet and the cable companies played important roles in shaping the overall UK traffic management landscape. Plusnet pioneered the explicit use of traffic management for product differentiation, while early experiences with complex application-specific tools in the cable industry created a lasting belief there that traffic management solutions should be simple to administer and explain to consumers.

Plusnet: Productizing Traffic Management

While other early adopters of DPI-based traffic management were quietly applying blanket restrictions to high-volume applications across their networks, Plusnet was spearheading an alternative approach that put traffic management at the center of its product offerings. In the early days of broadband, Plusnet bore the impact of BT Wholesale pricing perhaps even more so than most other ISPs because Plusnet had a policy of operating broadband as a profitable business. It would generally only buy capacity that customers were paying for and not more (Wyse 2008).

To accommodate growing demand while remaining profitable, the company initially experimented with an application-agnostic approach in which the traffic of the network's heaviest users was gathered together on the same link, forcing them to contend with each other rather than interfering with lighter users. However, based on a backlash from heavy users, conversations with other large ISPs where DPI was already employed for traffic management, and a desire to manage traffic in a more sophisticated and efficient way, Plusnet soon shifted to application-specific management that made use of DPI equipment (Imtech Telecom 2005).

As former Plusnet CTO Alistair Wyse (2008) has explained, Plusnet began by instituting an aggregate limit on peer-to-peer traffic capacity, but found it to be too much of a "blunt tool," treating heavy and light users of peer-to-peer applications the same. Instead, the company opted to develop specific usage profiles based on how their customers were using different applications across all hours of the day and to base its product offerings on those profiles. The result was a set of broadband packages that combined rate limits for high-volume applications with prioritization schemes that took effect when congestion occurred. For example, a lower tier package first offered in 2005 included (among other limits) rate limits of 100 Kbit/s for peer-to-peer and newsgroups, 2.5 Mbit/s for YouTube, and 2 Mbit/s for all other streaming from 6:00pm to 10:00pm. The rate limits differed during other hours of the day, with no limits from 12:00am to 4:00am. This was complemented by a prioritization scheme where VoIP and gaming received top priority, streaming and browsing were in the middle, and peer-to-peer and newsgroups received the lowest priority during times of congestion (Plusnet 2012a). Plusnet offered a series of higher-tier packages with higher rate limits and different prioritization schemes at higher price points, topped off by a package that prioritized VoIP and gaming but otherwise included no application-specific management (Plusnet 2012b).

Although Plusnet was a small ISP with a savvy user base, its "productization" and transparency around traffic management signaled to other ISPs that application-specific management could be a legitimate foundation for broadband product offerings. Numerous interviewees praised the Plusnet approach as a "really good solution" that was "more intelligent" than existing solutions and that other ISPs "definitely tried to learn from." By deliberately and openly segmenting its products based on specific application limits and priority schemes, it differentiated itself and set an example that other ISPs later looked to as they were considering traffic management options. Plusnet showed that to some degree and with a specific audience, customers would accept application-specific management. It could be used to sell broadband products, not just control costs and performance in the background.

Cable's Preference for Simplicity

Although the UK cable networks' architectures were obviously very different from those of the DSL providers, many similar factors played into the cable operators' traffic management decisions.

In the early days of broadband there were two primary companies offering cable broadband service, ntl and Telewest. They merged in 2005, began operating as ntl:Telewest, and were re-branded as Virgin Media in 2007 following the combined company's acquisition of Virgin Mobile. Prior to 2010, application-agnostic management aimed at heavy users was the primary approach used to control costs and performance. Like the DSL operators, cable operators adopted these strategies on the assumption that they would receive few complaints since only a minority of users would see an adverse impact.

Not long after launching broadband, the cable networks were seeing the same kinds of growth and traffic patterns as the DSL operators, with a small fraction of users and applications creating the majority of network traffic. As with the US cable companies (and unlike the DSL operators), one of the major expenses associated with surges in traffic was the need to manage demand in the access network, where subscribers shared capacity and a single heavy user could significantly degrade the experience for other users sharing the same Internet Protocol port. One cable product manager aptly captured how this dynamic influenced his company's traffic management decisions:

> [A] tiny, tiny percentage of customers, you know, single percents, were absolutely blasting the network. And not only costing us a lot of money but actually ruining the peak time experience for other customers on that port. So you actually solved for two things. You're able to better manage those heavy users and that's going to help our capital profile, which is of course an objective. (You know, we don't necessarily shout about that on [our] web site.) But it's not all about that. It absolutely protects customers who are normal users from congestion. And congestion on cable is obviously a very important metric.

Finding ways to control the need to invest in re-segmenting the network to reduce the demand on shared access infrastructure became imperative.

The cable operators' earliest approaches to managing traffic were application-agnostic measures aimed to limit the impact of the minority of heavy users on the rest of the user base. In 2003, ntl instituted a 1 GB/day download usage cap and began contacting customers who persistently exceeded this threshold (T. Richardson 2003). Telewest had a similar policy for its lowest tier service (Telewest 2004), but also began trialing a more fine-grained system. During certain daytime and evening periods, network routers would measure the volume of traffic that each subscriber was sending and receiving. If a subscriber exceeded a pre-defined threshold in either the downstream or upstream direction, his or her line speed would be significantly reduced (up to 75%) until the end of the period. This functionality, known as "subscriber traffic management" (STM), was a feature of the Cisco routers already in use on the network.

In 2006, before the two networks were fully integrated post-merger, application-specific measures were briefly added to the ntl network, reducing the capacity available to peer-to-peer and other high-volume applications at peak times. The rate limits for different applications were set differently in different parts of the country depending on the extent of the use of high-volume applications. In areas with large student populations, for example, capacity for these applications was more severely limited.

This deployment was short-lived, however. With Telewest staff primarily taking responsibility over the product management functions in the combined company, the decision was made to remove the application-specific measures and apply STM across the whole network. The technology used to do the application shaping was "incredibly complex," "expensive to run," and "had obviously been a little bit unloved on the ntl network" according to one interviewee. The fact that it operated differently in different parts of the country made it difficult to explain to consumers and to assess network capacity plans overall.

The STM functionality, on the other hand, was "pretty low maintenance and pretty easy," as the interviewee put it. It was viewed as a simpler approach that could be explained in the

same way to customers everywhere. The technology was well understood within the combined company and did not require additional devices on the network beyond the routers that were already necessary to route traffic. STM was therefore viewed as a simpler and more cost-effective solution that helped to solve a largely similar problem that application-specific management was designed to address. The value of simplicity in traffic management was internalized at the company and continued to influence its decisions in later years.

6.2.5 Lack of Public Attention to Application-Specific Management

Role of Perceived Copyright Infringement

The widely held assumption that most, if not all, peer-to-peer file-sharing consisted of illegal copyright infringement created an important subtext in which both DSL and cable operators made their decisions about how to manage traffic from the very beginning. By the early 2000s, peer-to-peer applications had been widely associated with copyright infringing activity given the content industry's pursuit of a number high-profile lawsuits and messaging campaigns since the late 1990s. This association had two key effects on operators' traffic management decisions: it created an additional rationale for targeting peer-to-peer file-sharing with traffic management, and it gave operators some comfort that their customers would be unlikely to complain about traffic management practices aimed at peer-to-peer applications.

The extent to which illegal uses of peer-to-peer applications helped to legitimize their management varied from operator to operator, but across the broadband industry operators acknowledged its role in shaping traffic management decisions, even if that role was small. One former cable executive explained that when considering traffic management solutions, the thinking was that "if you're going to traffic shape, [de-]prioritize [peer-to-peer] first because it's illegal and it's all the heavy users." While others were not quite so emphatic, interviewees generally agreed that the perception of illegality made peer-to-peer applications

a more justifiable target for rate limiting. One policy executive explained that peer-to-peer was simply "politically disfavored traffic."

Perhaps more importantly, the perception of illegality provided operators with some assurance that customers were unlikely to complain about peer-to-peer management, either by calling up customer service agents or contacting the press or regulators. They expected customers to keep quiet for fear of retribution for their infringing activity. This dynamic underscored the attractiveness of peer-to-peer management because it was unlikely to increase either the public attention paid to traffic management or the volume of customer service calls, which can create significant costs for broadband operators. One former DSL strategist explained this vividly, illustrating why such a customer service call would be unlikely to take place:

> You know everybody used to just choke the hell out of peer-to-peer at peak. Just almost squeak it to nothing. And that was okay because who's going to complain about that? Who's going to phone up [and say], "I'm trying to download these films for ages and I can't get them"? "Alright kid, would you like me to forward you to Warner Brothers?" You know, it's not going to happen.

From the operators' perspective, consumers had a clear disincentive to complain. Some operators took this one step further by intimating to heavy users that given the volume of traffic they were downloading, they were likely to be engaged in illegal activity. In this case, the operators themselves were raising the specter of legal consequences for infringing activity, further reinforcing the message to consumers that drawing attention to themselves by complaining about traffic management practices was a potentially risky activity.

When they did complain, users primarily limited themselves to online forums and bulletin boards where their identities could be protected and where they could learn from other similarly situated users about ways to evade traffic management limits or switch providers. Operators have long been aware of these forums and the complaints they contain; many operators even host their own or have staff who participate on the forums and address concerns as they are raised. But awareness of forum-based complaints never created any

particular imperative against the use of peer-to-peer management. One engineer explained this as follows:

> I mean drop the peer-to-peer on the floor for two hours, I don't think that anyone is going to complain. And it represents quite a significant portion of traffic on the network. And tends not to drive too many calls when people phone up and say, "my BitTorrent isn't working." Generally they're not actually phoning up to do that. They might be bitching on a message board somewhere, but it's not such – the cost of that and the reputational impact it causes – it's a case still to be proven.

In short, the fact that peer-to-peer applications were widely associated with copyright infringement added an additional element of justification to some operators' traffic management decisions while also insulating them from the kind of backlash, both public (press or regulatory) and private (customer service calls), that could create financial and reputational costs.

Regulatory Focus Directed Elsewhere

With user complaints about application-specific traffic management maintaining a low profile, questions about consumer harm and discriminatory practices related to traffic management failed to materialize in the halls of Ofcom (and Oftel before it) during the early years of broadband. Instead, Ofcom was concentrated on creating a more competitive marketplace, where faster broadband speeds could be garnered at lower prices across more of the population.

UK broadband penetration was just 10% in 2004, about average across all of the OECD countries (OECD 2006). The provision of broadband by LLU operators was scant. Improving that situation was one of Ofcom's first priorities after it was created. The regulator laid out its plans for a revised broadband regulatory framework in 2004 (Ofcom 2004d) that, together with the broader set of regulatory changes being formulated as part of the strategic review of telecommunications, would allow LLU-based broadband to become a widespread service offering. That review on its own was a massive undertaking, and combined with the LLU market review, Ofcom had plenty of work to do to make broadband more competitive and

widely available. Traffic management practices that were potentially discriminatory or anti-consumer simply were not on Ofcom's agenda.

Blunt Instruments of Traffic Management

With neither Ofcom nor consumers putting pressure on ISPs early on to change their traffic management practices, operators had the flexibility to deploy application-specific management as a means of controlling costs and performance in the background, rather than as a central product feature. One DSL product manager explained that peer-to-peer and newsgroup management "was put in out of necessity as opposed to anything in terms of differentiation of product," although there were occasional exceptions such as Plusnet.

This flexibility accorded to the ISPs meant that operators could be aggressive – "pretty draconian" or "pretty blunt" in their own words – in configuring when and by how much they reduced the capacity available to high-volume applications. Limits on high-volume applications would commonly apply to all customers regardless of the package to which they were subscribed. Some limits were in place 24 hours per day, while others applied for shorter but still significant portions of the day. Under the "just squeak it to nothing" approach described earlier, the targeted applications had their speeds reduced dramatically, often down to single-digit percentages of users' headline speeds. With few customers and no market or regulator providing incentives to make application-specific management more targeted to congestion, more consumer-friendly, or less blunt in general, operators generally sought as much cost and performance control as they could without blocking high-volume applications altogether.

6.2.6 Summary of Early Years

Traffic management, both application-agnostic and application-specific, was a feature of UK broadband since the inception of broadband itself. The structure of the DSL market – with retail providers as customers reliant on BT Wholesale's nationally averaged, linearly priced backhaul – spurred DSL operators to seek out solutions that would allow them to control both

costs and performance. Because the majority of traffic could be attributed to heavy users and high-volume applications, those became logical targets. The cable operators saw similar usage trends, but early experience with a complex application-specific approach created a desire to maintain a simpler application-agnostic approach going forward.

The widespread management of applications on DSL networks was further rationalized on the basis that peer-to-peer and newsgroup applications were more amenable to having their performance reduced than other applications on the network while also being in use by a smaller segment of the subscriber base. The application awareness involved in conducting application-specific management benefitted operators by providing them with insights into how their networks were being used. Operators expected few customer complaints, publicly or privately, given consumers' perceived hesitancy to be associated with illicit activity taking place on peer-to-peer networks.

Taken together, all of these factors drove the adoption of traffic management strategies that would become the norm across the industry. By 2004, a selection of large and small DSL providers – serving more than a quarter of all broadband users – had taken up application-specific traffic management. Fairly blunt and unsophisticated management of peer-to-peer and other high-volume applications was most common. The cable networks and smaller operators like Plusnet differentiated themselves in ways that would have consequences in the market going forward.

While Ofcom was keenly focused on broadband issues, traffic management was not among them. The regulator's attention was diverted to larger efforts, the results of which contributed to a reshaping of the UK broadband landscape in the second half of the decade.

6.3 Competition Entrenched and Diffused Application-Specific Management

The UK broadband market entered a period of upheaval in 2005. Ofcom conducted a complete review of telecommunications regulation, culminating in a significant set of reforms adopted to spur competition in the broadband access market. Shortly thereafter, as net neutrality discussions in the US intensified, Ofcom began to formulate its position regarding discriminatory treatment of Internet traffic, beginning from the notion that the competitive framework in place should serve to discipline providers' behavior. At the same time, competition itself was intensifying some of the pressures that had led to adoption of application-specific traffic management in the first place, requiring network operators to invest in increasing capacity even as retail prices declined.

This section explains how these events created a climate in which application-specific traffic management became legitimized, entrenched, and diffused across the bulk of the UK broadband industry. Early broadband providers that had already been using application-specific approaches by and large left them in place untouched. Application-specific strategies had become so well accepted that BT briefly began rate limiting video traffic, including traffic generated by the BBC's popular iPlayer service, in 2009. Later adopters of application-specific management expanded the use of traffic management in new ways and rationalized their choices differently than early adopters had. The DSL operator Sky was a notable exception to these trends, showing that a major ISP could compete on price and performance without managing applications.

6.3.1 Ofcom's Shaping of the Traffic Management Landscape

One of Ofcom's first and largest tasks when it was founded was a complete review of existing telecommunications regulation, known as the Telecommunications Strategic Review (TSR). Announced in late 2003, one of the key goals of the TSR (and of Ofcom more broadly) was to promote competition in telecommunications services. As Ofcom noted in launching its first

TSR consultation, “despite nearly 20 years of regulatory activity intended to promote competition” (Ofcom 2004b, 2), BT remained in a position of significant market power (SMP) in a number of telecommunications markets, including wholesale broadband markets (Ofcom 2004c).

The undertakings to which BT agreed at the conclusion of the TSR helped to change that. Local loop unbundling had been possible in the UK since 2000, but only in 2005 did it become a widely viable possibility for competitive operators. As part of the undertakings, BT agreed to functionally separate its access network division (which became Openreach) from the rest of the company and to provide equivalent wholesale prices, terms, and service guarantees to all ISPs, including its own ISP, BT Retail (Ofcom 2005a). Separately, BT also made commitments to lower access network prices, maintain its bitstream prices at sufficient levels to allow LLU operators to compete, and make improvements in provisioning of service systems used by competitive operators (Ofcom 2005b).

Together, these commitments – and the oversight structure that Ofcom and BT agreed to use to enforce them – provided the foundation for an explosion in LLU, as depicted in Figure 5. From 2005 to 2006, the number of unbundled lines grew by more than 500%, instantly giving UK users increased broadband choices (Ofcom 2011a). By 2009, 35% of all broadband subscriptions were via unbundled lines and 85% of broadband customers had access to at least one LLU entrant (Nardotto, Valletti and Verboven 2012), in addition to BT, bitstream providers, and, in many areas, cable.

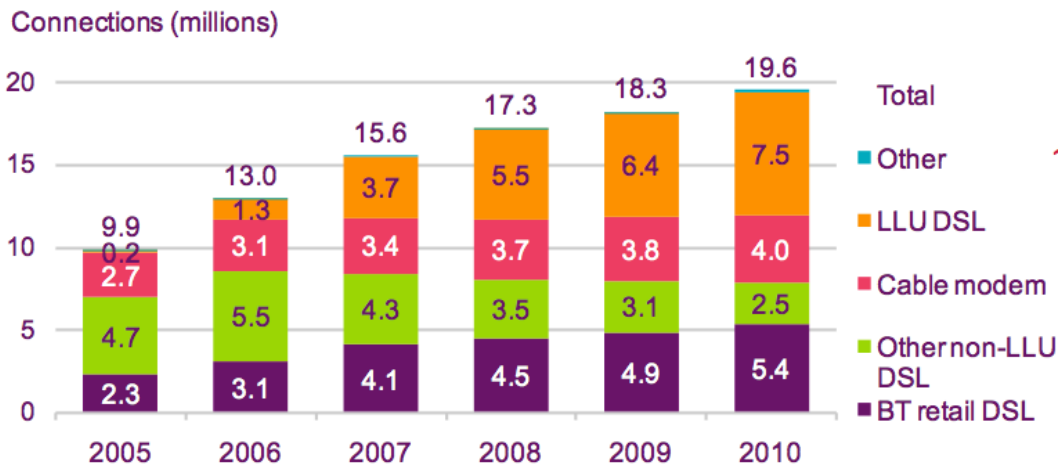

Figure 5. Number of UK residential and small business fixed broadband connections by connection type, 2005-2010. Reproduced from Ofcom (2011a).

Ofcom's Belief in Competitive Discipline

The TSR was a "huge exercise" (as one former Ofcom official put it) for the newly constituted regulator. The TSR also resulted in some of the boldest steps to promote competition that any European telecommunications regulator had taken since the beginning of liberalization – a "radical intervention," as Ofcom's former head of strategy described it (M. Brown 2010). Although other national regulators in Europe and elsewhere had intervened to spur the development of LLU and other means of competitive entry, Ofcom was the first to require functional separation of the incumbent (Tropina, Whalley and Curwen 2010). The former official characterized the undertakings as "a very heavy regulatory regime" with "nondiscrimination rules on steroids" to ensure that alternative providers would have the chance to compete.

It was obvious from the beginning of the review that one important measure of Ofcom's success would be whether competition emerged in the broadband market. In the years following the conclusion of the TSR, Ofcom repeatedly highlighted the emergence of competition and the associated consumer benefits of lower prices, higher speeds, and more choices, as evidenced by the agency's annual reports:

> Genuine competition in the fixed-line telecoms markets, at the deepest levels of infrastructure, creates a virtuous circle of new investment in emerging technologies and innovation in services and price competition. This benefits consumers and helps maintain the competitiveness of the UK economy as a whole. (Ofcom 2006a, 18)
>
> Towards the year-end, the UK passed a significant milestone with more than half of all households now on broadband, encouraged by cheaper prices, higher bandwidths and an increasing range of new services. . . . Ofcom's role in wholesale broadband regulation has been critical to the roll-out of Local Loop Unbundling, which passed the two million lines mark during the year . . . . (Ofcom 2007a, 4)
>
> Citizens and consumers are already reaping the benefits of competition. Over half of all households have broadband from one of more than 500 different providers. (Ofcom 2008a, 6)

Ofcom's belief in the promise of competition and its devotion to ensuring competitive entry had profound effects on the regulator's approach to traffic management in the context of net neutrality. As the net neutrality debate attracted increasing attention in the US in 2006, Ofcom began developing its own position as to whether regulatory intervention was necessary to prevent discriminatory practices, including application-specific traffic management. While in the early years of traffic management in the UK Ofcom had been focused elsewhere, the growing attention devoted to net neutrality abroad drew Ofcom's attention to the issue and created an imperative for the regulator to react.

One of the key tenets of its early positioning on the issue was that competition among broadband providers in the UK was sufficient to discipline providers' urges to discriminate. One former Ofcom official explained that "because it's a competition regulator it sees [net neutrality] as a competition problem." In its first public writing on the issue, Ofcom declared that "[f]or those operators without SMP in the relevant market, we consider that efficient working of a competitive market will address the risks posed to consumers from non-network neutral approaches. An effectively competitive market at the retail level, with relatively low barriers to entry, means that customers have a range of choice in their ISP" (Ofcom 2006b, 17). As a competition-focused regulator, Ofcom instinctively approached the issue with the belief in the ability of competition to restrain operators' behavior.

In public addresses over the years that followed, Ofcom officials reiterated this argument, emphasizing that the UK has "such strong retail competition" (Ingram 2006); that "competition itself can provide a constraint" (Kiedrowski 2007); that degrading application performance would make an ISP "a less attractive proposition to consumers" (D. Scott 2007); and that "[i]n a more competitive environment, there is less inherent problem with traffic management and prioritisation" because "[i]f network operators get these calculations wrong, consumers will switch to another provider" (Richards 2008). The belief in competition existed at the highest levels of the agency.

Underlying these arguments was not only a belief in the logic of competition, but also the need for Ofcom to prove that its interventions to promote competition were justified. To those who had crafted the extensive competition framework that emerged from the TSR, it was unconscionable that a problem could so soon emerge in public discussion that could not be remedied by the competitive market. One former Ofcom official described this internal agency logic:

> Providing open access is a pretty big competition intervention. You don't do it for fun. Therefore, hang on a minute, surely, the threshold we would want to set would be so high that we can't even begin to see there's going to be a problem.

To Ofcom, it was essential to demonstrate that competition delivered on its promises. Having expended tremendous time and resources to achieve the post-TSR regulatory framework, it was not conceivable that discriminatory traffic management was a problem that required a solution beyond competition. Not only did Ofcom believe in competition, it was institutionally imperative that competition be shown to work for consumers.

Competition was therefore a constant refrain for Ofcom clear through to the end of the decade. Spurred on by the conclusion of the review of the EU telecommunications framework and burgeoning net neutrality discussions in Europe, Ofcom issued a public discussion document in 2010 in which it repeatedly affirmed the argument it had been making since 2006:

> Without market power there is a strong presumption that no anti-competitive effects and consumer harm will arise. This is because consumers will tend to punish attempts at exclusionary behaviour by simply shifting their business to an alternative provider who does not engage in the same exclusionary practices. At present, both the fixed and wireless retail broadband markets in the UK are considered effectively competitive. In principle, therefore, there ought to be sufficient choice of provider to discipline firms' behaviour. (Ofcom 2010c, 26)

Ofcom's Claims about the Necessity of Application-Specific Management

In the first several years after the conclusion of the TSR, Ofcom officials also took the position that application-specific management was essential to safeguarding network performance. In 2006, Ofcom explained to the European Commission that ISPs were "routinely degrading" peer-to-peer applications "given the implications this traffic can have on other users of the network" (Ofcom 2006b, 17). Ofcom officials argued publicly that peer-to-peer management was an "example of discrimination on today's internet that is essential for its smooth operation" (Ingram 2006), that without it "the net would be in a sorry state right now" (Ingram 2006), and that peer-to-peer traffic "would otherwise cause the network to gridlock" if its performance were not degraded (Kiedrowski 2007).

These arguments are remarkable in their insistence about the necessity of the specific technique that network operators had most widely adopted, especially given that when network operators took the occasion to discuss traffic management themselves, even they were not publicly framing the issue in such dire terms. Nildram, for example, explained to its customers in 2006 that high-volume use of peer-to-peer applications "sometimes has a knock on effect for other users on the network" (Ferguson 2006). When ntl briefly introduced peer-to-peer limits in 2006, it was "to maintain an excellent quality" service in light of the fact that peer-to-peer applications "may have a detrimental affect [sic] on other users' services" (Towny 2006). While the impact of peer-to-peer users on others was noted, operators were not claiming that the measures taken were absolutely necessary for the network to continue functioning. By comparison, Ofcom officials insisted that peer-to-peer traffic would render the network unusable unless specifically degraded.

Perhaps because large operators had been managing peer-to-peer traffic for several years, Ofcom took for granted that this was a necessary tool to safeguard performance. The regulator failed to acknowledged the cost-based rationales that accompanied operators' decisions, as if the tradeoff between the demands that peer-to-peer traffic put on the network and the ability to meet these demands by deploying additional capacity did not exist. Nor did Ofcom even hint at the possibility that application-agnostic traffic management approaches could achieve the same goals, or that application-specific approaches could be more or less tailored, invasive, or blunt depending on how they were designed. Instead the regulator appears to have observed what was going on in the marketplace and assumed that it was indispensible.

Ofcom's initial and enduring position was that traffic management, whether discriminatory or not, was acceptable and that problematic practices could be remedied by competition (or, as a last resort, by Ofcom's existing regulatory powers). By focusing on competition as a disciplining force and rationalizing peer-to-peer management as a performance necessity, Ofcom made it obvious to operators that their application-specific traffic management approaches would not meet with regulatory resistance. The regulator took no issue with the approaches that were already in place, and if operators did take up traffic management that consumers found problematic, they were expected to "be punished in due course through the market mechanism, not through anything that the regulator might do," as one former Ofcom official explained. The threat of regulatory backlash was essentially nonexistent.

6.3.2 Market Effects of Competition

After the TSR was completed, it was not long before increased competition started having profound effects in the marketplace, most notably on price and speed. The year 2006 became the year of the "free" broadband offer, with TalkTalk, Sky, and Orange all offering free broadband as part of a bundle with other services (Ofcom 2007b). From 2005 to 2008, the

average cost of broadband fell by over 40%, with Ofcom attributing around half of that decrease to LLU take-up (Ofcom 2009b).

At the same time, operators were investing to rapidly increase connection speeds. Average headline broadband speeds more than doubled in 2006, from 1.6 Mbit/s to 3.6 Mbit/s (Ofcom 2007b). By 2009 they had nearly doubled again, to 7.1 Mbit/s (Ofcom 2010d). In short, the arrival of LLU made faster speeds available to a larger portion of the population at lower prices. Figure 6, reproduced from Analysys Mason (2010), demonstrates this aptly: as the fraction of broadband subscribers on LLU rose, prices declined sharply, making 8 Mbit/s service affordable where it had previously been prohibitive while also driving prices down for slower offerings.

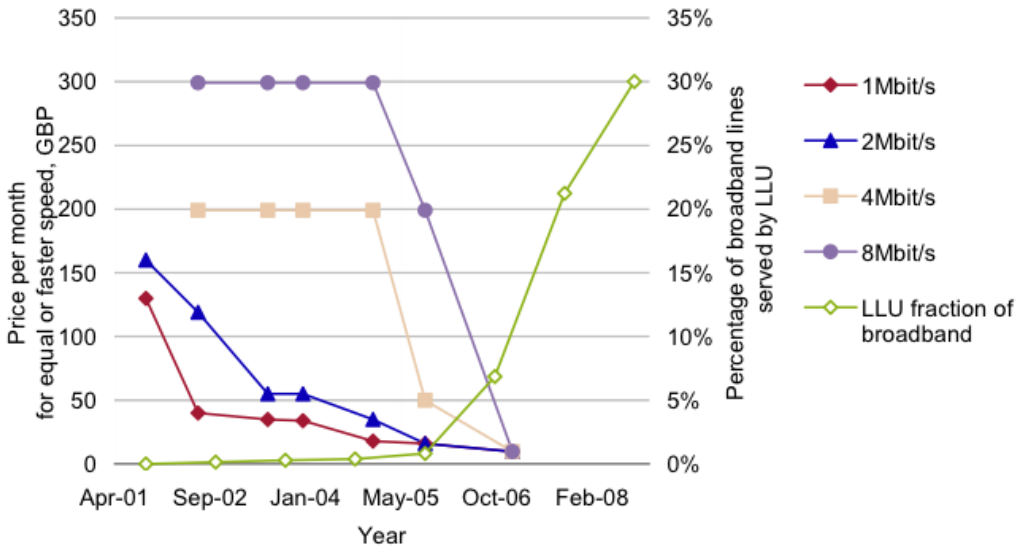

Figure 6. Lowest market price for a UK broadband product of specified speed, 2001-2008. Reproduced from Analysys Mason (2010).

As faster headline speeds became available, Ofcom increased the pressure on operators to continue to deploy new capacity and to deliver the speeds that they advertised. Faced with complaints from consumers and Ofcom's own Consumer Panel (2007), the regulator worked with industry in 2008 to create a voluntary code of conduct requiring signatories to disclose,

among other things, estimated actual speeds in addition to headline speeds (Ofcom 2008c). That same year, Ofcom initiated a nationwide speeds testing initiative that used specialized hardware placed within thousands of residences to obtain statistically significant samples of broadband performance across the nine largest fixed retail providers (Ofcom 2009a). Since 2009, Ofcom has reported the speeds testing results annually and continued to add new tests, most of which focus on web browsing performance.

The simultaneous decrease in prices and pressure to increase speeds worked in tension. Operators felt pressure from Ofcom and the market to increase network capacity, which required up-front investment. But offering broadband for free, as operators felt pressure from the market to do, was not much of a strategy for raising the capital necessary for serious network expansion. One DSL product manager described this dynamic aptly in puzzling over a competing ISP's pricing structure: "It's just crazy! We look at our pricing and just can't work out how we do it. And then they're doing it at half the price again . . . this is crazy." These pressures were direct results of the newly introduced competition spurred by LLU.

6.3.3 Application-Specific Management As a Legitimized Practice

The climate of permissibility that Ofcom established around traffic management during its early forays into net neutrality policy combined with competitive forces in the marketplace to entrench application-specific traffic management where it already existed and diffuse it across the industry, even to some operators that initially resisted such measures. The extent to which application-specific management had been legitimized was demonstrated differently by different factions of the largest operators, but by the end of the decade they had all adopted application-specific management in some form, with the key exception of Sky.

Early Adopters Left Application Management Largely Untouched

For operators that had deployed application-specific management in the early-to-mid 2000s, the competitive dynamic and Ofcom's reluctance to intervene reinforced their original traffic management decisions. In the newly competitive environment, expending budget on

bandwidth-independent costs made sense: every provider paid the same to hook up an access line, and more access lines meant more customers. But anything providers could do to reduce bandwidth-dependent expenses helped them to compete on price in the marketplace. As Everything Everywhere explained to Ofcom in 2010, "[w]ith both speed and costs driven by the level of peak time capacity provided, networks must optimise the balance between peak time speeds and the impact on retail charges to arrive at a competitive position" (Mooyaart 2012, 2).

This logic applied whether providers were unbundling exchanges (as was the case with TalkTalk, Tiscali, and Orange, for example) or not (as was the case with BT Retail). The cost model for LLU backhaul was more attractive than for bitstream. Operators could deploy their own backhaul and enjoy the associated economies of scale without being beholden to BT Wholesale's nationally averaged bandwidth prices (Analysys Mason 2008). But the competitive nature of the market post-2005 gave even LLU operators that were transitioning customers from bitstream a good rationale to keep their peer-to-peer management in place for all customers regardless of whether they were being served from an unbundled exchange or not. Peer-to-peer traffic was a large fraction of overall traffic; managing it down continued to help operators save on bandwidth costs. The market dynamics that resulted from the growth of LLU simply gave DSL operators more reason to keep their existing traffic management strategies in place.

Operators had so much flexibility to manage applications as they wished that a number of them made virtually no major changes in how they managed applications over the course of the decade even though the capacity of their networks changed radically during that time. One large operator, for example, had instituted a per-user rate limit for peer-to-peer traffic at peak times that in the early 2000s equated to about 3% of a typical user's overall headline speed. Since that rate limit remained largely unchanged – and it always applied to all users regardless of the package to which they were subscribed – by the end of the decade users' peak time peer-to-peer traffic limits were closer to 0.05-0.5% of headline speed depending on

their chosen packages. Because there was little external pressure to increase the limit and plenty of cost savings to be had by keeping it as low as it was to begin with, this approach to traffic management did not evolve as the network evolved. At the ISP where participant observation was conducted, traffic management researchers' attempts to convince broadband product managers to relax or reduce peer-to-peer limits were met with significant, ongoing resistance. Some providers did reduce the number of hours per day during which applications were managed, but the mechanisms used for management and the limits in place did not necessarily evolve. Many of the "pretty blunt" and "pretty draconian" approaches that were instituted early on only became more so as time wore on.

A 2009 decision made by one early adopter, BT Retail, provided further compelling evidence of the extent to which application-specific management had been legitimized in the UK. That year BT Retail quietly began limiting video streaming to less than 1 Mbit/s at peak time on its low-end "Option 1" package, which at the time offered an 8 Mbit/s headline speed. BBC staff soon noticed this change and the impact that it was having on the popular iPlayer streaming service (Cellan-Jones 2009a). The dispute between the two corporate giants escalated in the press, with BT publicly seeking payment from video providers to carry their traffic and the BBC accusing BT of hiding its practices from consumers (Watson 2009). After intense scrutiny in the press, BT abandoned the practice in less than a year (the ramifications of this decision are further discussed in Chapter 7).

Despite reversing course, BT's mere attempt to manage traffic in this way was an extremely bold move. BT Retail was the nation's largest ISP. Option 1 was among its most popular broadband products. Video streaming was rapidly becoming one of the most popular applications on the Internet, with Sandvine (2010a) estimating 80% growth in peak-time streaming traffic in Europe from 2008 to 2009. The iPlayer was exploding in popularity (BBC 2010). The fact that BT believed, or even hoped, that imposing limits on such popular applications for so many subscribers would be accepted in the marketplace was a testament to the culture of legitimacy that had developed around application management.

Later Adopters Rationalized Application Management in New Ways

Not all operators entered the post-Undertakings era with application-specific management in place. In the cases of O2 and Virgin Media, the companies had strong reasons for refraining during the early years of serious LLU growth. By the end of the decade, however, they both developed rationales for adding application-specific management to their networks, further demonstrating the diffusion of application-specific management across the industry.

O2

O2 purchased the nation's only pure ADSL2+ network, Be, in 2005 and launched its own line of consumer broadband products in 2007. When it came to traffic management, Be had always been something of a "purist brand," as one interviewee described it, with no management or restrictions of any kind on its network; the company's motto was "life uninhibited" (Be 2005). In the early years after O2's acquisition, this philosophy persisted and was bolstered by the fact that the network was flush with capacity, obviating the need for the kind of cost and performance control mechanisms used by earlier traffic management adopters. With up to 24 Mbit/s service offered across 150 unbundled exchanges at purchase, O2 had a clear speed advantage over competing ISPs whose top-of-line products maxed out at 8 Mbit/s. It also had a smaller user base, yielding minimal contention on the network. As a result, there was no need to invest in traffic management solutions.

The situation changed as the network and the market matured. As with many other LLU providers, O2 offered both "on-net" LLU service and "off-net" broadband service that relied on bitstream access from BT Wholesale. For O2, the disparity in the performance of the two kinds of service offerings was more severe than its competitors since the LLU network provided faster-than-average speeds. Congestion on the off-net network was becoming a serious problem, resulting in the decision in 2009 to introduce application-specific management for its off-net product that combined prioritization of streaming, gaming, and other time-sensitive applications with rate limits for peer-to-peer and newsgroups (Ferguson 2009b), not unlike the approach taken with Plusnet's product offerings years earlier. That

move improved customer satisfaction while slowing the pace of bandwidth investment necessary to meet demand.

Impressed with the results of application-specific management on the off-net network, O2 moved to shed all vestiges of Be's purist approach, adopting a strategy that put traffic management at the center of its product offerings. In 2010, O2 unveiled a wholly new product strategy that involved three different packages, each of which had different rate limits for peer-to-peer and streaming applications at different times of day (Ferguson 2010). The low-end package capped peer-to-peer traffic at 50 Kbit/s per user at peak and 100 Kbit/s off-peak and included an 800 Kbit/s limit on streaming video. The only limit on the high-end package was a 250 Kbit/s limit on peer-to-peer at peak times. The theory behind this approach combined the usual rationales for application-based rate limits – cost and performance control – with the notion that traffic management could be used to segment the market, attracting customers who were willing to knowingly buy a product with degraded application functionality at a lower price point.

This was a bold step from O2, which was the sixth largest fixed ISP in the UK at the time. Even after BT Retail had met with resistance and ultimately abandoned its effort to impose streaming video limits on its lower tier packages the previous year, O2 accepted the risks associated with taking the same approach. In doing so, it became the largest ISP to use traffic management as a product differentiator – not just as a means of background cost and performance control – and to have the public accept that strategy as legitimate. Smaller ISPs like Plusnet had operated on a similar model for years. But the fact that O2 – a highly recognizable brand with a fixed network supporting more than 600,000 customers – was able to sustain a product strategy with traffic management at its center demonstrated how legitimized application-specific traffic management had become.

Virgin Media

Virgin Media, meanwhile, was benefitting from a tried and trusted application-agnostic traffic management system together with cost effective network upgrades that increased capacity without requiring significant new investment. The subscriber traffic management (STM) system that Telewest had deployed years prior continued as an effective tool for moderating the behavior of heavy users on a daily basis. Like O2, Virgin offered an "off-net" DSL product that relied on bitstream backhaul and had an application-specific priority scheme that gave higher priority to VoIP, gaming, and streaming than to other applications. But the core cable product that served the vast majority of the user base had no application-specific management.

The years that followed the TSR saw the launch first of 20 Mbit/s offerings and then 50 Mbit/s offerings made possible by the introduction of DOCSIS 3.0 technology on the network. Compared to the up-front investments made by the LLU providers and BT's investments in fiber, the capital expenditure necessary to deliver increasingly higher speeds on the cable network was small. As Virgin Media CEO Neil Berkett explained in 2009, "iPlayer was launched in January of 2008. All of the DSL ISPs sort of went how terrible is this, it is costing us capital. We felt it, but it didn't cost us capital" (Berkett 2009, 2). Moreover, the deployment of DOCSIS 3.0 that year required only "incremental tens of millions of pounds" from which the Virgin network derived "significant upgrades in terms of capability and capacity" (Berkett 2010, 10).

However, even the combination of the well-functioning STM system and relatively inexpensive capacity upgrades were not enough to defend against competitive pressures to minimize operating costs through application-specific management. In 2010, Virgin Media introduced aggregate rate limits for peer-to-peer and newsgroup applications across all of its packages, just as numerous DSL ISPs had done years prior. By applying limits in the aggregate, they affected all users similarly and thus did not suffer from the complexities of the ntl system that had been in use years prior.

Unlike in earlier years, Virgin Media's decision in 2010 was not a matter of finding the applications creating the most traffic on the network and ratcheting them down. Just two months after the application management was introduced, Virgin Media CTO Paul Buttery presented the graphs shown in Figure 7 to analysts, showing that the network had experienced 300% traffic growth in the previous three years and that streaming and web traffic combined to represent more than half of peak time network traffic, while peer-to-peer and newsgroups accounted for less than one quarter.

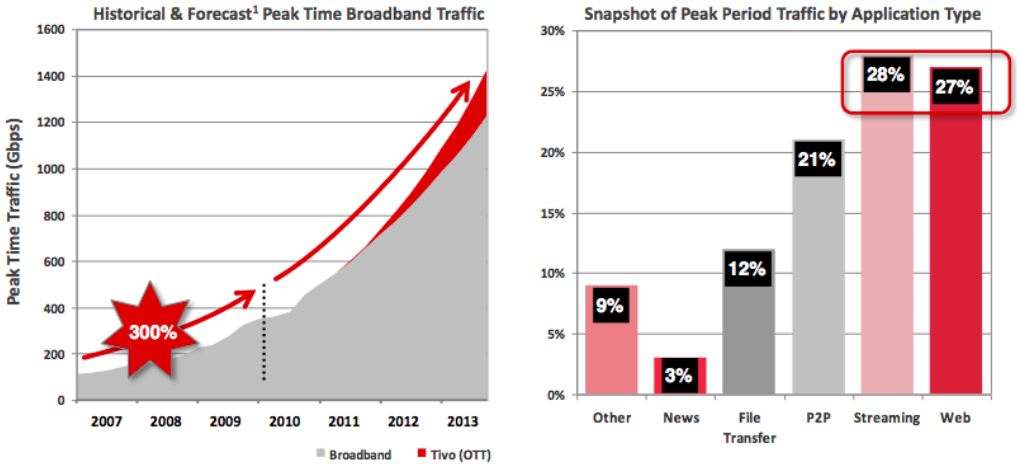

Figure 7. Peak period traffic growth and volume by application type on the Virgin Media network. Presented December 2010. Reproduced from (Buttery 2010).

Buttery explained this data as follows:

> [I]f you look at the things that are driving that growth – what we've got there is a snapshot of peak traffic, and we've broken it down into the major types there. If you'd been looking at that chart three years ago, that would have been a picture that was dominated by peer-to-peer. And what we're seeing now is because of the quality of our network and because of the drive towards over-the-top TV, people are taking more and more video streaming from us, and that streaming traffic is continuing to grow and grow. We've also seen web traffic rising. (Berkett 2010, 7)

Clearly, although traffic growth had been substantial, it was not being driven by the applications that Virgin Media started managing in 2010. When those applications were dominant years prior, Virgin had refrained from managing them. In the interim, application-specific management was legitimized by Ofcom even as competition forced providers to offer

higher speeds at lower prices. The result was that in 2010, Virgin had the flexibility to take a step that, while not aimed at its biggest traffic source, helped to reduce peak time demand even further than the STM system could.

Although the idea of managing the highest volume applications may have no longer held, many of the other rationales that providers had long used to justify their traffic management decisions did. Peer-to-peer applications were still viewed as less time-sensitive than other popular applications, their management still seemed unlikely to draw customer complaints, and limiting the traffic of modest peer-to-peer users – those whose peer-to-peer usage was higher than average but not enough to trigger the STM thresholds – could still improve performance for other customers on shared links.

Unlike O2, the particular ideals upheld internally at the company created an additional incentive for Virgin to seek ways to manage traffic beyond STM. For Virgin, being able to claim that all of its packages were "unlimited" had always been viewed as extremely important from a marketing and product differentiation perspective. One of the key criteria that had been established through years of advertising campaign disputes before the Advertising Standards Authority (ASA) was that packages with monthly caps could not be advertised as "unlimited." As a result, Virgin was committed to not introducing monthly caps on any of its cable products. There was an ethos about unlimited products that ran deep in the company.

Having taken monthly caps off the table, and with the STM system already in place, Virgin had limited further options if it wanted to reduce peak time traffic demand. Adopting an application-agnostic solution similar to the one Comcast had developed after the FCC intervened was considered to have the same problems as the old ntl application-based management: because it was only triggered when links became congested, the traffic management would affect different customers in different parts of the network differently. Establishing an aggregate rate limit for peer-to-peer and newsgroup applications allowed the

products to be described as "unlimited," with a traffic management explanation that applied the same to everyone. These public-facing considerations plus the opportunity to reduce capacity costs spurred the nation's second-largest ISP in 2010 to follow the path that so many of its competitors had adopted long before.

Sky As the Lone Major Differentiator

Similar to O2, Sky entered the broadband market by purchasing a network (Easynet) with large excess capacity, few existing customers, and many unbundled exchanges. Also like Be, Easynet had operated under a corporate philosophy that did not support application-specific management. This had been a key underpinning of the high-end unlimited broadband packages that Easynet offered under the UKOnline brand and was believed internally to be a powerful differentiator in the marketplace. At the point of acquisition, accommodating relatively few customers on an underutilized network made it simple to perpetuate the unmanaged approach from a demand management perspective.

But even as take-up soared and Sky surged from tens of thousands of users at launch to 3.5 million in 2011, that strategy remained unchanged. Sky held firm to its approach of dimensioning the network to accommodate demand, and as a result likely transitioned its backhaul network to higher capacity links sooner than its competitors. Sky offered an off-net product with application-specific management just as its competitors did, but that technology was not applied to the core on-net broadband offering that most of its customers purchased.

Unlike most of those competitors, Sky had a television product that provided it with a major source of revenue. Sky customers were not able to purchase broadband without a TV subscription until 2010 – the first year that broadband was profitable for the business (Darroch 2010). Although the exact figures are not public, it is widely assumed that Sky was able to continue to compete on both speed and price without traffic management or obvious performance problems because it subsidized the cost of broadband capacity with TV revenues.

This was in stark contrast to competitors such as Virgin, whose CEO acknowledged in 2010 that "the most gross margin for us is produced by broadband, followed by fixed line telephony, followed by our B2B and our mobile business, followed by video" (Berkett 2010, 9). Sky did not feel the need to manage cost in the same way as the rest of the market because it was able to cross-subsidize its broadband service. A number of much smaller providers such as Zen and Andrews & Arnold likewise refrained from managing applications, but they did so at a significantly higher retail price point than Sky given that they had no way to subsidize the cost of unmanaged traffic demand.

Sky's unmanaged approach belied Ofcom's claims about the necessity of peer-to-peer management while signaling the extent to which other operators' decisions were at least partially grounded in concerns about cost. By maintaining a network whose performance was good enough to continue to attract new customers, Sky showed that focusing on capacity upgrades rather than traffic management could allow its customers to achieve acceptable performance on the DSL network without the operator needing to take action against peer-to-peer applications. In other words, Sky's existence and growth demonstrated that there was nothing inherent about the impact of peer-to-peer applications on the network that required some response other than capacity upgrades.

Sky also held firm to the belief that lack of traffic management was a marketable quality to consumers, often emphasizing its lack of traffic management in marketing materials. Figure 8 shows an example from an advertising campaign that included the claim that Sky "never slows you down at peak times no matter how much you use" (Hermes Project 2011).

Figure 8. Sky Broadband advertising campaign from 2011.
Reproduced from Hermes Project (2011).

It is difficult to know how much the unlimited, unmanaged aspect of the product has contributed to Sky's growth, but having maintained that philosophy even as every major competitor decided otherwise, the company continued to demonstrate its belief in the virtues of running a nondiscriminatory network. As will be discussed in the next chapter, Sky provides an important practical example to test theoretical arguments about how to safeguard nondiscrimination. If the goal of a regulatory policy is to ensure that most users have access to a nondiscriminatory offering, then the growth of Sky indicates that supporting a competitive retail marketplace may be one strategy for achieving that goal. But if the goal is to ensure wide nondiscriminatory access to applications offered by independent application providers, the UK case should be more accurately viewed as a disconfirming example of the power of competitive forces.

6.4 Conclusion

In stark contrast to the US, the UK broadband landscape has been dominated by traffic management, with application-specific management becoming increasingly pervasive over time. The DSL wholesale market structure and the nature of high-volume applications provided incentives for DSL operators to limit those applications early on, while early experiences with complex application-specific management on cable networks inspired cable operators' preferences for simpler solutions. Over time, competitive pressure, the low risk of backlash from customers, and Ofcom's insistence on both the market's disciplining power and the necessity of peer-to-peer management firmly ingrained application-specific management as an acceptable practice industry-wide. As a result, both DSL and cable operators toward the end of the decade rationalized its adoption and extended its use in new ways. The key exception was Sky, which used lack of management as a product differentiator and found other ways to pay for the costs associated with its customers' unfettered traffic demand.

The implications of the UK experience are profound. The UK case reveals that in markets with a wholesale/retail split, the structure and bandwidth prices built into the wholesale market can have significant effects on traffic management decisions. It demonstrates how competitive forces – in absence of any countervailing pressure from customers or regulators – can intensify operators' incentives to manage applications. It provides a testament to the influence of the regulator not only in assuring the industry of its regulatory restraint, but also in bolstering competitive conditions. And it shows how much inertia can surround traffic management technology once it has been deployed in an environment where neither customers nor regulatory authorities demand – or even discuss – how the technology might evolve as networks evolve. The next chapter draws out these implications for the net neutrality policy debate by comparing and contrasting the US and UK experiences.