

Obren, Mark; Howell, Bronwyn

Conference Paper

The tyranny of distance prevails: HTTP protocol latency and returns to fast fibre internet access network deployment in remote economies

8th Asia-Pacific Regional Conference of the International Telecommunications Society (ITS): "Convergence in the Digital Age", Taipei, Taiwan, 26th-28th June, 2011

Provided in Cooperation with:

International Telecommunications Society (ITS)

Suggested Citation: Obren, Mark; Howell, Bronwyn (2011) : The tyranny of distance prevails: HTTP protocol latency and returns to fast fibre internet access network deployment in remote economies, 8th Asia-Pacific Regional Conference of the International Telecommunications Society (ITS): "Convergence in the Digital Age", Taipei, Taiwan, 26th-28th June, 2011, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/52331>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**NEW ZEALAND INSTITUTE FOR THE STUDY
OF COMPETITION AND REGULATION INC.**

**The Tyranny of Distance Prevails:
HTTP protocol latency and returns to fast fibre internet
access network deployment in remote economies**

November 21 2010 (replaces version of October 2010)

Mark Obren and Bronwyn Howell

Address for Correspondence:
New Zealand Institute for the Study of Competition and Regulation Inc.
Victoria University of Wellington,
PO Box 600, Wellington, New Zealand.
Email bronwyn.howell@vuw.ac.nz

Acknowledgements: The authors wish to acknowledge the helpful comments of Mike Belshe, Lewis Evans, Dave Heatley and Anton Nannestad. The views in this paper solely reflect those of the authors, and do not necessarily represent those of the institutions with which they are affiliated or their constituent members. Any errors or omissions remain the responsibility of the authors.

Abstract

As public policies seek to advance deployment of enhanced broadband infrastructure as a means of acquiring economic advantage, the issue has arisen of the extent that additional economic performance accrues from increases in headline bandwidth speed in locations that are physically remote from the infrastructure hosting time-critical information services. For time-dependent applications, latency (the time delay in accessing data across a network) is correlated with the effective bandwidth (the actual speed of access), and thus impacts upon the economic performance of the application to the user. We extrapolated data for interactive web-based applications from Belshe (2010), where latency was found to substantially reduce the effective bandwidth available to the user of a typical web-based application, to estimate the effective bandwidth over a range of headline bandwidths and latencies typical of web-based transacting patterns in New Zealand. We find that the decreasing returns on effective bandwidth as headline bandwidth increases are further exacerbated by the higher levels of latency experienced as a consequence of New Zealand's distance from the bulk of the global infrastructure supporting web-based applications. The benefits of enhancing headline bandwidth through new forms of faster infrastructure were substantially reduced by the impact of the latencies typically experienced by New Zealand users accessing remote web based applications, and thus the economic benefits expected from investment in infrastructure in accessing enhancing those applications most impacting economic performance is likely to have been exaggerated; providing an insight into a constraint upon cloud computing and other web-enabled information systems.

JEL classification: O18; O33; R11; R12

Keywords:

Internet, broadband, fibre-optic cable, regional economic development, HTTP protocol

The Tyranny of Distance Prevails

Internet technologies have been widely claimed to herald an end to the “tyranny of distance”¹ that has proved costly for small, remote economies whose economic performance has historically relied upon the trade of (often bulky) physical goods. Consequently, the governments of remote economies (such as Australia, New Zealand and South Africa) have been urged to develop sophisticated internet-based “knowledge economy” policies as a means of not just countering historic distance-based disadvantages, but as a means of ‘leapfrogging’ ahead of international competitors in developing and capitalising upon internet-enabled economic growth opportunities². For example, in a submission to the New Zealand Government commissioned by the Minister for Information Technology’s IT Advisory Group, Ernst & Young (1999:7) assert:

“For New Zealand, one of the single most important aspects of the Information Age is the “death of distance” (Cairncross, 1995). Distance no longer determines the cost of communications. This will be one of the most dynamic shaping forces for New Zealand. Patterns of international trade, concepts of national borders, and the basis of decisions about where people live and work will be altered in ways that are only dimly imaginable.”

“In this new view, New Zealand is at the centre of the world. Our white-collar workers can compete on price and quality with those in London or California. New Zealand will be able to retain graduates who until now have emigrated in search of higher salaries. Work-related travel will decline. People will no longer have to live in cities to work; instead, they will be able to work from wherever they choose to live. ICT will enable New Zealand not only to overthrow the tyranny of distance but also to mitigate the disadvantages of our small population and low population density.”

The substantive message of these claims has been remarkably consistent over time, and has been used to support almost all ICT-related policy endeavours in both countries³. Not surprisingly, similar claims have been used in support of recent policies adopted in both Australia and New Zealand committing the respective governments to invest substantial sums in the deployment of ultra-fast ‘fibre to the home’ (FTTH) broadband access networks⁴. For example, the New Zealand government claims its investment creates:

“a step-change in the provision of broadband services. The expansion and development of broadband is a vital component of New Zealand's economic growth, productivity improvements and the government's wider strategy to increase New Zealand's global competitiveness, particularly compared to other OECD countries.”⁵

¹ The term the ‘tyranny of distance’ is widely attributed to the Australian historian Geoffrey Blainey, whose 1966 work *The tyranny of distance: how distance shaped Australia's history* explored how physical distance shaped the development of Australia’s economic and social institutions.

² For a discussion of the role of economic geography on economic performance in OECD countries, see Boulhol, de Serries & Molnar (2008). Evans & Hughes (2003) discuss the relevance for New Zealand. Several studies have examined the implications for Australia, including Battersby (2006), Battersby & Ewing (2005), Wilkie & McDonald (2008) and Rahman (2005).

³ See, for example, Gregor *et.al.*, (2004); Ovum (2005); Muller & Williamson (2005).

⁴ In Australia, this is the A\$43 billion National Broadband Network in New Zealand it is the NZ\$1.5 billion Ultra Fast Broadband initiative). For a comparative evaluation of both proposals, see Heatley & Howell (2010).

⁵ http://www.med.govt.nz/templates/StandardSummary_40551.aspx accessed 20 October 2010

Whilst acknowledging the use of the FTTH network for “IPTV and 3D media are other applications that may become popular with residential users”, the New Zealand Telecommunications Minister continues to assert that the economic step-change is predicated upon the use of the new super-fast access network to address commercial disadvantages arising from the “tyranny of distance that’s hampering businesses here compared to ones in the US that have access to a vast internal market”⁶. Indeed, in both countries, it has been suggested that the new, ultra-fast local access networks could underpin the development of new information processing industries whereby firms based in Australia or New Zealand can become ‘information hubs’, storing and processing information generated by clients in physically distant foreign locations⁷.

Without doubt the internet has enabled some of the ‘connectedness’ and competitiveness gaps to be closed between Australian and New Zealand firms and their counterparts in the rest of the world. However, just as with the historic ‘freezer ship’, with which the internet has been frequently compared as an innovation with the potential to revolutionise the two nations’ economies, technological limitations mean that the benefits cannot be expected to accrue equally across all firms and industries. Whilst the freezer ship opened up opportunities for Antipodean⁸ farmers to supply British and other distant markets, the accrual of benefits was confined to agricultural products and the development of competition was constrained by the extent of substitutability of the frozen and long-keeping Australian and New Zealand variants for fresh and perishable northern hemisphere goods. Likewise, with internet transacting, the market opportunities opened up will be contingent upon the relative disadvantages still faced by Australian and New Zealand firms given the nature of the markets in which they are transacting and technologies they must employ compared to those used by competitor firms located much closer to the end customer (Howell, 2006).

Whilst internet technologies have lowered the costs and shortened the time taken to transport information, and increased the quantity of information that can be moved in a given time period, two important facts still remain: first, Australian and New Zealand firms still sit at a vast distance from the parties with which they seek to transact; and second, the countries both lack scale. Combined, this means that for truly international firms relying upon the trade

⁶ <http://computerworld.co.nz/news.nsf/news/lets-have-some-excitement-around-ufb-joyce> accessed 20 October 2010.

⁷ For example, the New Zealand Capital Markets Development Taskforce has suggested New Zealand could utilise its strengths in agriculture to position the country as a capital market hub in this area, and to develop upon its existing expertise as an exporter of high value middle and back office financial market services for this enterprise as well as for funds management firms (CMD, 2009).

⁸ In the [British Isles](#), “the Antipodes” is often used to refer to [Australia](#) and [New Zealand](#), and occasionally [South Africa](#) and [Zimbabwe](#), and “Antipodeans” to their inhabitants.^[2] Geographically the antipodes of the British Isles are in the [Pacific Ocean](#), south of New Zealand. Wikipedia <http://en.wikipedia.org/wiki/Antipodes> (retrieved October 31 2010).

of information goods⁹ or the timely transmission of information as a significant component of their income-earning potential, the majority of their customers will have multiple alternative potential suppliers sitting at the end of very much shorter cables than those over which the Antipodean firms must transact. These factors tend to suggest that place-based policies endeavouring to generate localised relative economic advantages will succeed only to the extent that the technologies they rely upon militate against the nonlinear tendencies that attend the agglomeration of economic activities in locations where (absence of) distance and scale advantages already exist (Glaeser & Gottlieb, 2008).

Technological and Economic Limitations to ‘Closing the Gaps’

It is apposite therefore to examine the economic natures of the internet technologies employed to enable Australian and New Zealand firms to both compete and transact with foreign parties, and assess the extent to which they ameliorate or exacerbate the problems of distance and (lack of) scale. This necessitates a consideration of both the transport technologies used (i.e. fibre-optic cables analogous to the freezer ships of the past), and the applications and the relevant markets in which they are deployed (analogous to the differentiated cargoes carried on the freezer ships).

Fibre-optic cables enable large quantities of data to be transported at the speed of light¹⁰. However, distance matters: the light degrades within the cable, so the signal must be boosted regularly. Every time the signal is boosted, it is converted from light to electric signals, put through electronics and converted back into light for the next leg of the journey. Each conversion adds significantly to the time taken for the journey. The longer the journey, the more boosts are required and the greater is the delay. The Southern Cross Cable linking New Zealand to North America has some 500 repeaters spaced at intervals of between 40 and 70 km¹¹. The average time taken for a round trip (RTT) of data from Wellington to a platform in the United Kingdom (the historic locus of New Zealand trade transactions) is 375 milliseconds. The RTT to East Asia (the new trade locus) is around 350 ms. The RTT to near neighbour Australia is around 120ms (Table 1). This compares to a worldwide average RTT to Google servers of slightly over 100ms, and to United States Google servers of between 60 and 100ms (Belshe, 2010). Lack of scale means that currently New Zealand is

⁹ For a definition of information goods, see Shapiro & Varian (1999). For a more specific definition of ‘digital goods’, see Quah (2002).

¹⁰ The speed of light in fibre-optic cable is approximately 2/3 of the more commonly quoted speed of light in a vacuum. Data can therefore be transported along a fibre-optic cable at a maximum theoretical speed of approximately 200,000 km/sec. At this speed, it would take 200ms to circle the globe at the equator.

¹¹ Southern Cross (2006) Network <http://www.southerncrosscables.com/public/Network/default.cfm> (accessed October 22, 2010).

primarily served by only one sub-oceanic cable system linking it to the worldwide internet¹². This cable directs most traffic first across the Pacific to the west coast of the United States, and then onward to its ultimate destination. A smaller volume of traffic is directed to Australia, with interconnections through hubs in Guam and the Philippines, to Asia and beyond.

If the data transported is not time-sensitive, then it may not matter much that it takes longer for data to travel to and from Australian and New Zealand firms (analogous to frozen meat being viewed as a close substitute for fresh by its end consumers). This is likely the case for the transmission of large volumes of static data contained in music and video files. However, where the benefits of the application rely upon direct interaction between transactors at either end of the cable, then the time delays do matter. Most of the interactive internet applications commonly used today rely upon web pages constructed using the HTTP protocol¹³. These HTTP-enabled¹⁴ applications are commonly used by firms interacting (transacting) with each other electronically across continents and oceans – for example, for participating in electronic marketplaces or placing orders through supply chain management applications.

HTTP uses short, bursty connections between the remote user and the platform on which the application is based. Specifically, for every packet of data sent during an HTTP session, an acknowledgement of receipt is expected by the transmitting system. Until the acknowledgement is received, the transmitting system assumes that transmission has not been successful, leading to delays in sending subsequent data packets. In a simple case, doubling transmission times (e.g, due to distance) quadruples the time waiting as messages have to travel both ways. This is further compounded by increased likelihood of packet loss as a consequence of systems ‘timing out’ and retransmitting the packet¹⁵. Together, these factors lead to real reductions in performance the greater the distance the packet must travel. This effect can be exacerbated if software development tools are calibrated for delays shorter than those typically experienced by New Zealand users. Thus, the further away the internet user is from the platform, the larger is the number of nodes the message must pass through, and the

¹² Although plans have been announced to construct a second cable to compete with Southern Cross, albeit with the same United States destination as served by Southern Cross. See Pacific Fibre <http://www.pacificfibre.net/> (accessed October 22 2010).

¹³ Strictly speaking, the HTTP protocol is layered on top of the more general-purpose TCP protocol. Most of the characteristics examined in this paper are discussed in this paper are features of the TCP protocol.

¹⁴ The authors acknowledge that alternatives to HTTP exist in the HTML5 specifications, with WebSockets being an example of data communications capabilities for browser-based applications without requiring the use of HTTP protocol. However, HTTP remains the dominant form of access to web-based applications in 2010, and thus is a primary area of focus for current policy development.

¹⁵ It is noted that that the TCP/IP protocol on which HTTP relies was designed to carry non real-time data traffic. It thus contains a high level of redundancy in order to achieve error-free transmission, so is not optimized for low levels of latency required by real-time applications.

greater is the number of interactions required to complete each transaction, the greater the impact of these delays will be¹⁶¹⁷.

The combination of distance and HTTP protocol leads to some substantial disadvantages for Australian and New Zealand firms relative to their counterparts sitting much closer to the platforms hosting such time-sensitive applications. Although the volume of internet data exchanged increases as the volume of non time-sensitive video, music and other large content downloads increases, the variety of applications based upon time-sensitive HTML5 and CSS standards is increasing exponentially. Increasingly, these are applications supporting economically important e-commerce transactions between firms and in markets. Furthermore, as a consequence of the economics of scale and density, there is ongoing diffusion of interactive international applications onto cloud computing platforms that are located in major population regions such as North America, Europe and Asia and not small, remote ones such as Australia and New Zealand. As the unique time-sensitive data required to support such interactive transacting is not amenable to solutions such as caching¹⁸ (storing a copy on a local server for distribution more speedily to a local population) or Content Distribution Networks (CDN) - both used to militate against the tyrannies of distance for static content - the inevitable conclusion is that for those Australian and New Zealand firms transacting electronically and internationally via HTTP-enabled applications, the tyranny of distance is still very much a factor disadvantaging them relative to their distant foreign counterparts.

Local Access Bottlenecks or Latency Barriers?

This observation leads back to a consideration of the credibility of claims made that policies based upon the deployment of government-funded ultra-fast local access broadband networks will lead to a step-change in local economic performance as a consequence of increased global competitiveness of firms based in the small, distant Antipodean economies. At the very least, this claim would appear to be highly contingent upon the types of products and transacting the firms undertake.

As Figure 1 illustrates, increasing the bandwidth available in the ‘last mile’ bandwidth bottleneck within Australia and New Zealand will have no effect upon the

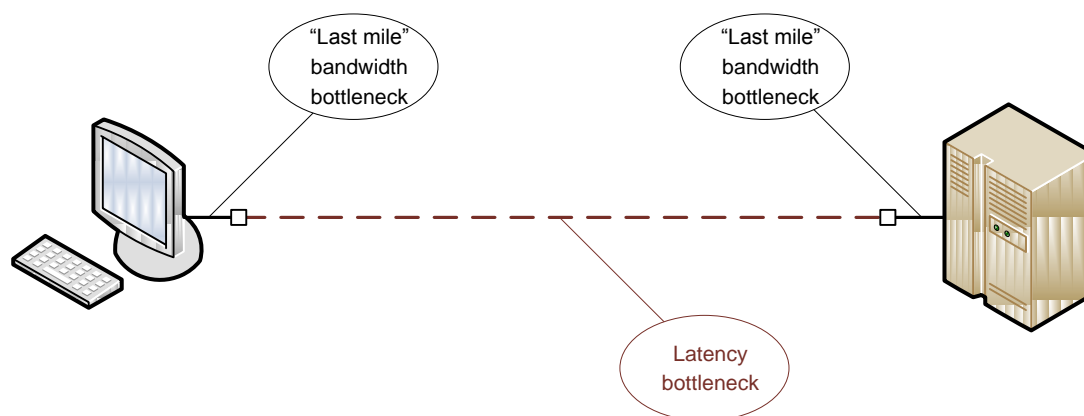
¹⁶ It has long been recognized in the technical literature that there is a trade-off between latency and bandwidth that must be considered in network design. See, for example, Lakshminarayanan & Padmanabhan (2003).

¹⁷ The authors acknowledge that other factors exist that impact upon the user experience, such as congestion, HTTP pipelining, and selective acknowledgement. It is recognized that techniques exist that enable enhanced application performance, but these are applied on an application by application basis by individual firms, rather than across an industry so are less amenable to policy intervention. Whilst these and other factors may affect the user experience, our data set is based upon a consistent application of these tools so does not alter our findings regarding the correlation of latency and effective bandwidth.

¹⁸ Whilst it is possible to construct individual sites to cache active data content, the bulk of sites that New Zealanders access are internationally located and target an international audience, and thus would not be calibrated to optimize the experience for New Zealand users.

‘bottleneck’ that occurs as a consequence of the length of the cables connecting the user and the platform. However, it is apposite to examine the extent to which an increase in the available local bandwidth affects the performance of specific web-based applications for Australian and New Zealand users. If local bandwidth is actually a performance constraint for local firms using HTTP-enabled applications, then there may be benefits to such an investment. However, if the major source of delay is actually the distance from the platforms on which applications are based, then investment in local bandwidth will have minimal effect upon the competitiveness of, and hence the economic potential offered by, Antipodean firms transacting electronically in international markets using HTTP-enabled internet applications.

Figure 1: Which Bottleneck Matters Most?



Source – Dave Heatley.

To this end, we draw upon the work of Belshe (2010) and extrapolate it to take account of the particular parameters prevailing in the New Zealand context. In his United States-based study, the author systematically varied both the bandwidth available and the length of time taken for a message to pass from the user to the platform and back again to assess the effect of increasing bandwidth on the time taken to load web pages for a sample of 25 of the most popular pages on the internet. Section 1 describes his methodology and his key finding that an increase in bandwidth from 5Mbps to 10Mbps led to a 5% improvement in page load times, and that the relationship between the ‘effective bandwidth’ of the web page download as the raw bandwidth available increases exhibits decreasing returns. In Section 2, we extrapolate and calibrate Belshe’s findings to assess the effect of increasing local access bandwidth from a range of capacities currently used in New Zealand to the proposed 100Mbps speed for the government-funded Ultra-Fast Broadband Network and the RTT times routinely experienced by New Zealand internet users communicating with platforms in New Zealand, Australia, Asia, the United States, Europe and Africa. Consistent with Belshe’s findings, we conclude that the effective increase in performance in web page loading speed as

a consequence of an investment in 100Mbps connections is very modest indeed. Section 3 makes some policy conclusions arising from these findings.

1. Belshe Methodology and Findings

Belshe, a software engineer working at Google Labs¹⁹, undertook his study to evaluate the effect of bandwidth and RTT on the loading of web pages as an illustration of the fact that, despite the intuitive focus of users on local access bandwidth speed, it isn't the only factor to take into account when assessing the performance of the internet. He summarises “given the way HTTP uses short, bursty connections, it turns out that round-trip-times dominate performance more than bandwidth does”. He conducted his experiments in the United States in April 2010, by systematically evaluating the effect of page load times for a varying range of bandwidth speeds at a constant distance (proxied by Return Trip Time – RTT), and then varying the RTT for a given bandwidth.

In order to separate out the effect of bandwidth from the distance over which the message must pass, Belshe chose a base round trip time of 60ms and packet loss of 0% and simulated page load times via HTTP for 25 of the web's most popular web pages²⁰, varying the bandwidth from 1Mbps to 10Mbps in 1Mbps increments. His results are shown in Figure 2. Latency (page load time) exhibits strongly decreasing returns as average bandwidth increases (left panel – Belshe, 2010:2), with a clear concave pattern evident in the percentage improvement from adding extra bandwidth (right panel – Belshe, 2010:3). Belshe then records the effective bandwidth of the web page download as the raw bandwidth capacity is increased (Figure 3 – Belshe, 2010:3). At 1Mbps, web pages can be downloaded at around 69% of total bandwidth capacity, but this rapidly tails off to only 16% at 10Mbps.

Figure 2: Page Load Times and Bandwidth: 60ms

¹⁹ <http://www.belshe.com/about/> (accessed October 22 2010).

²⁰ These are not stated, so we are unable to replicate his findings directly. However, the findings are consistent with experiments done on sites with which we are familiar.

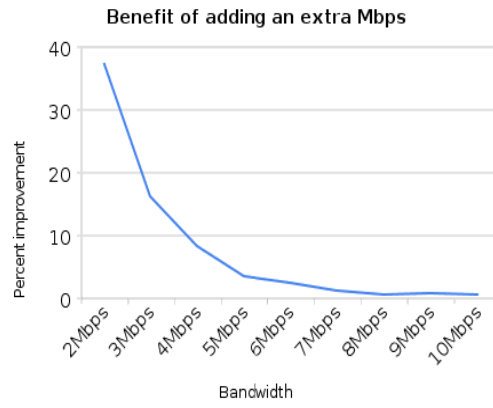
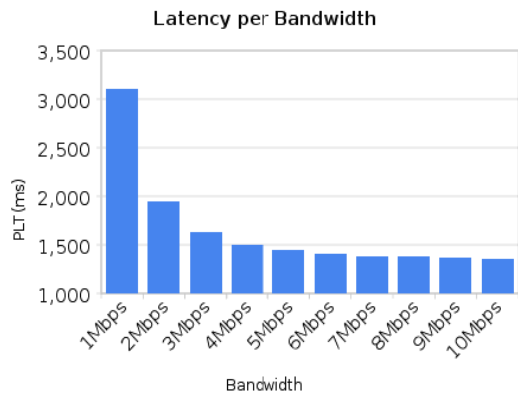
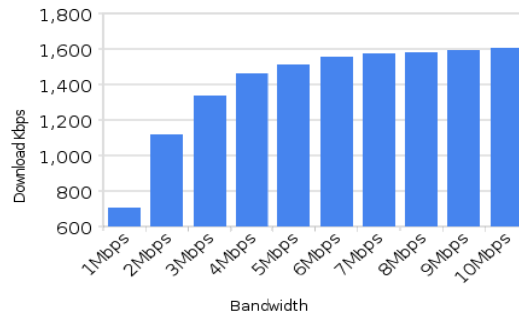


Figure 3: Effective Bandwidth of HTTP: 60ms



Belshe then proceeds to vary the RTT from 0ms to 240ms in 20ms increments for a constant bandwidth of 5Mbps. The reduction in page load times is broadly linear as RTT decreases (Figure 4 – left panel – Belshe, 2010:4), with the percentage reduction varying from 7% to 15% (right hand panel – Belshe, 2010:5). The effective bandwidth as the RTT decreases is shown in Figure 5 (Belshe, 2010:5). With a high RTT, the effective bandwidth was as low as 550kbps, or only 10% of the theoretical 5Mbps throughput of the connection. Overall, this relationship too exhibits decreasing returns (Figure 5).

Figure 4: Page Load Times (PLT) and Varying RTT: 5Mbps

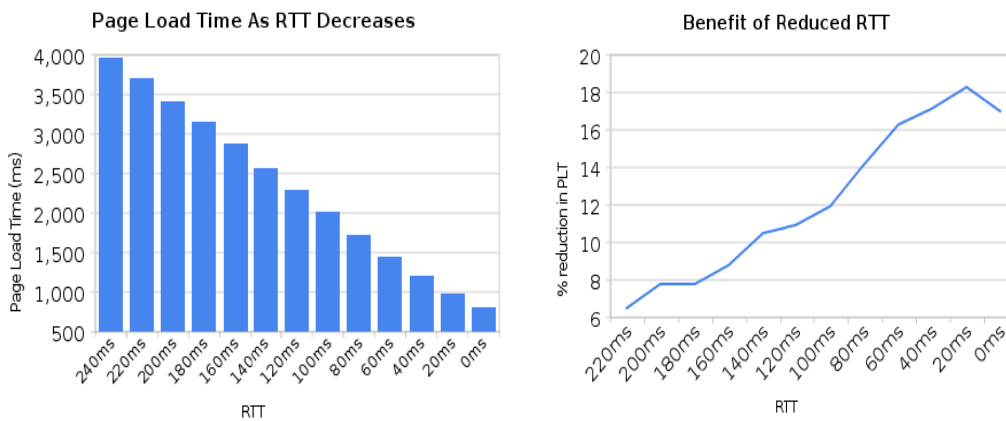
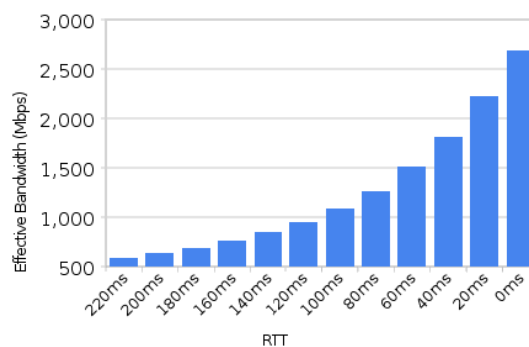


Figure 5: Effective Bandwidth as RTT Decreases: 5Mbps



Belshe concludes that if users double their bandwidth without reducing the RTT significantly, the effect on web browsing will be only a minimal improvement. He suggests that speeding up the internet (and hence user experiences) will be better achieved by reducing RTTs (e.g. content nearer to users), redesigning web pages so that they require less communication between the client and the platform (e.g. fewer 'handshakes' when establishing the communications in the first place), and technological innovation in fibre-optic cables so that the trips themselves can be made faster.

2. Extrapolating from the Belshe Findings

Belshe (2010) uses a set of parameters for his model that sit at the lower end of those which we need to assess in the New Zealand context. Firstly, he focuses on relatively low bandwidth at the local access level. The proposed New Zealand and Australian ultra-fast broadband networks have specified a target speed of 100Mbps – ten times that of Belshe's maximum of 10Mbps, where the effective bandwidth at which pages were downloaded was only 16% of the capacity available, even at the relatively short RTT of 60ms (Figure 3). Extrapolating Belshe's relationship for even higher bandwidth at 60ms gives us the relationship illustrated in Figure 6.

Figure 6 confirms that the sharply decreasing relationship between actual bandwidth and increases in page load time performance mean that there is very little additional benefit to be gained from investment in local access network speeds above 10Mbps, even for HTTP-enabled applications communicating with local platforms (e.g. within the same region – where RTTs of 60ms might be expected). To put the benefits of 100Mbps bandwidth in context, at the current point in time, most New Zealand broadband users have access to fixed line ADSL services with headline speeds in excess of 5Mbps (the capacity specified for rural broadband connections under the Government's Rural Broadband Initiative – RBI). Many in metropolitan areas already enjoy headline speeds of 8Mbps or higher, with a recorded actual average speed of 3.3Mbps (Akamai, 2010). Following the completion of the Chorus cabinetisation programme in 2011, 80% of New Zealanders will have access to broadband connections with headline speeds between 10Mbps and 20Mbps²¹. Some in metropolitan areas will have access to ADSL2+ services with speeds up to 24Mbps²². Mobile broadband connections using 3G technology claim effective speeds in excess of 1Mbps²³ and headline speeds of up to 7.2Mbps²⁴, while 3.5G wireless technology suppliers in New Zealand claim

²¹ <http://www.chorus.co.nz/enhancing-the-broadband-network>

²² ITU G.992.5

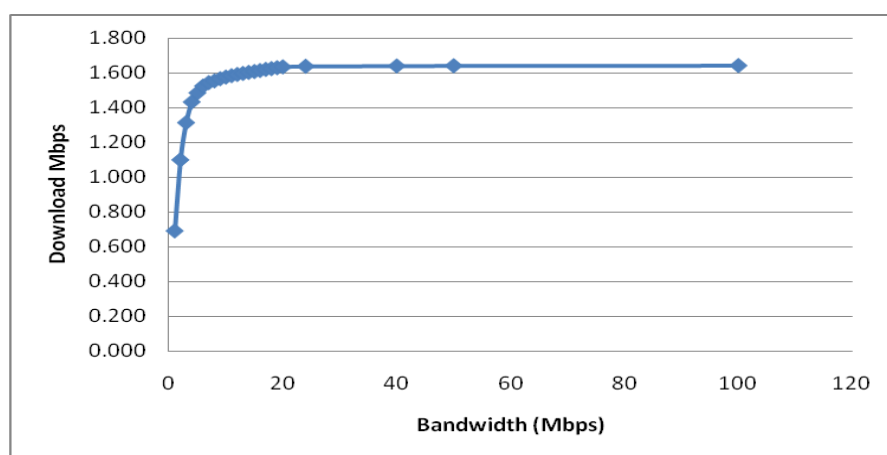
²³ <http://www.vodafone.co.nz/mobile-broadband/speed.jsp?q=mobile%20speed>, confirmed by Akamai (2010), where the average speed actually achieved is reported as 1098kbps – that is, 1.07Mbps.

²⁴ <http://www.nbr.co.nz/opinion/chris-keall/telecom-xt-launch-day-speed-tests>

headline speeds of 28.8Mbps²⁵ and effective speeds of around 20Mbps will be available in the very near future²⁶.

Given the likely performance of applications at the local broadband access speeds currently available or anticipated for the near future, Figure 6 suggests that most of the benefits from increased bandwidth for local HTTP-based applications will come from existing technologies without the need for investment in ultra-fast local access networks.

Figure 6: Effective Bandwidth of HTTP at RTT=60ms



We move now to consider the effects of distance. For New Zealand firms, the platforms they will be communicating with will likely be based in locations with significantly longer RTTs than the sub 220ms samples used by Belshe. To confirm this we sampled the RTTs for a variety of overseas locations that are likely to be typical of the New Zealand internet-trading firms’ transacting patterns. The specific sites sampled were selected because they have access to reasonable local bandwidth, are not cached in either New Zealand or another closer location (verified using www.geobytes.com’s IP locator tool), and allow RTT data to be collected. Table 1 summarises a sample of over 2000 RTTs for each website from a location near Wellington, New Zealand in September 2010.

Table 1 confirms that, as Belshe found within the United States, the average RTT for sites within New Zealand is between 65 and 85ms. The average RTT for sites in eastern Australia is between 120 and 135ms. This is broadly consistent with his finding for the average worldwide RTT to Google sites. Thus, in the same time that it takes an average worldwide internet user to communicate from anywhere with Google, the average New Zealand firm can communicate with platforms only as far away as Australia. It takes over

²⁵ http://www.telegeography.com/cu/article.php?article_id=35052 accessed on November 21st 2010

²⁶ Vodafone New Zealand sponsor’s presentation at the 1st Asia-Pacific Regional International Telecommunications Conference, Wellington New Zealand, August 26 2010.

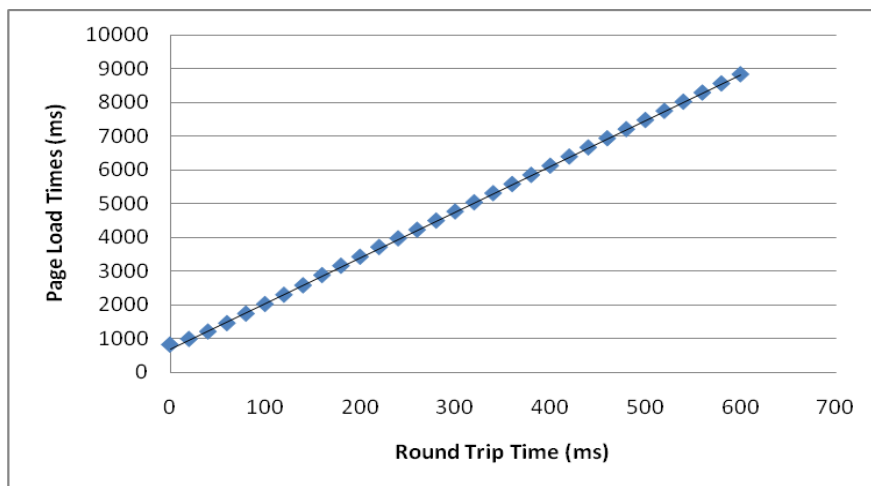
200ms to communicate with platforms in the west coast of the USA – and as most other internet traffic must go there before being routed to its final destination, most other times are consequently greater than this. For example, the east coast of the USA and Canada are between 200 and 300ms away, the United Kingdom and Europe around 380 ms and South Africa nearly 600ms. Whilst the primary cable is United States-focused, the subsidiary route via Australia delivers traffic to Asia. Singapore is around 200ms away (comparable to the west coast of the United States), whilst Eastern Asia – the new focus of New Zealand’s agricultural exporting activity – is around 350ms away (only slightly better than Europe) . The effect of distance is clearly going to have a further substantial effect upon the page load times of HTTP-enabled applications.

Table 1: RTTs From Wellington, New Zealand, September 2010

Country	Website	RTT (ms)		
		Minimum	Average	Maximum
New Zealand	www.trademe.co.nz	60	65	204
	www.stuff.co.nz	80	86	234
	www.nzherald.co.nz	68	83	229
Australia	www.whitepages.com.au	103	134	398
	www.foxtel.com.au	101	119	268
Singapore	www.changiairport.com	203	208	416
	www.singtel.com	200	206	415
Hong Kong	www.thestandard.com.hk	308	315	473
China	www.chinamobileltd.com	362	375	534
	www.chinaunicom.com	311	353	507
East Coast, USA	www.nytimes.com	216	244	402
	www.wsj.com	256	261	416
	cbs5.com	291	304	463
West Coast, USA	www.nordhavn.com	192	197	342
	www.visitcalifornia.com	214	240	401
	lawa.org	193	208	416
United Kingdom	www.bbc.co.uk	346	375	534
	www.whitepages.co.uk	373	389	543
Germany	despiegel.de	353	374	547
South Africa	yellowpages.co.za	584	595	877

To assess the effect of RTTs on page load times, we used linear regression to extrapolate Belshe’s page load time graph for varying RTTs at a constant bandwidth of 5Mbps (Figure 4) to find the relevant times for typical New Zealand-based international transactions. The results are contained in Figure 7. We are now able to calculate the percentage reduction in page load times and the effective bandwidth as the RTT decreases, as per Belshe’s calculations illustrated in Figures 4 and 5. It is noted that the linear assumption arrived at from Belshe’s sampling at lower RTTs may or may not hold for the longer RTTs. However, as greater distances are more likely to result in both a greater number of signal repeats and packet losses, and hence even longer page load times, our use of a linear extrapolation leads to us finding a ‘best case’ scenario for page load times.

Figure 7: Extrapolated Page Load Times for NZ RTTs, Bandwidth=5Mbps



We now assume that the linear relationship between distance and effective bandwidth applies equally across all bandwidths. This assumption is plausible for the New Zealand case, as beyond the local transport network, most traffic is subject to the constraints of the Southern Cross cable, at least as far as the west coast of North America. By applying the same RTT percentage changes to each of a range of relevant bandwidths at the 60ms level, we can estimate the likely effective bandwidth for each of the RTTs representative of New Zealand transacting patterns (60, 100, 200, 240, 300, 360, 600ms) for each of a range of local bandwidths representative of those currently available or possible in the near future (Belshe’s 1Mbps to 10Mbps, plus 12, 20, 24, 50 and 100Mbps). The results are shown in Figure 8.

It is now possible to estimate the likely improvement in page load times arising from the New Zealand Government’s investment in a 100Mbps FTTH local access network, relative to a range of bandwidths either currently available or likely to be available in the near

future. These results are displayed in Table 2. Whilst the FTTH network offers substantial improvements over basic broadband (1Mbps) regardless of the destination of communication, the benefits rapidly decrease as the existing capacity increases. For example, relative to a network with speeds of 8Mbps currently available widely, the 100Mbps FTTH network will offer a page load performance improvements ranging from 19% for sites in New Zealand and 13% for sites in Australia to 5% for sites in East Asia and Europe and 3% in South Africa. These benefits are only very slightly better than those possible from ADSL2+ speed connections likely to be widely available following completion of the Chorus cabinetisation programme in 2011 (11%, 8%, 5% and 3% respectively).

Figure 8: Estimated Effective Bandwidth, New Zealand Parameters

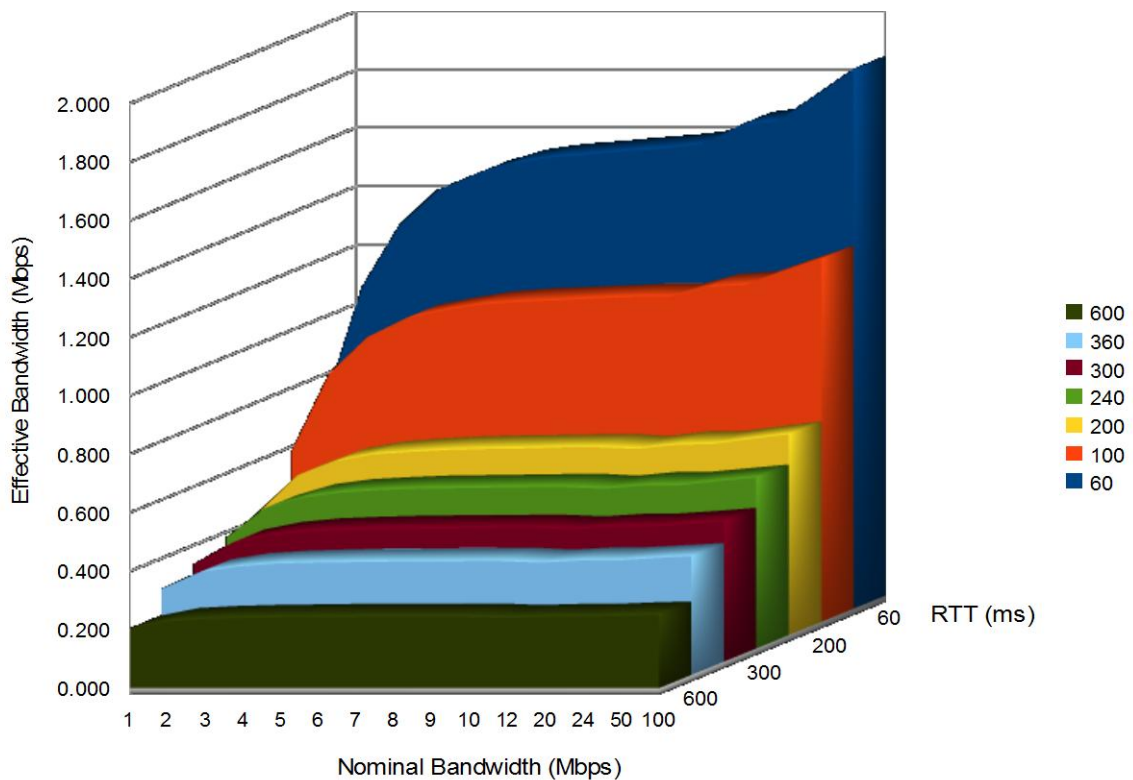


Table 2: Percentage increase in effective bandwidth of 100Mbps Fibre Optic over indicative alternative network capacities

Network Type	Bandwidth of Network Type (Mbps)	Theoretical speed improvement 100Mbps fibre if 100% scaleable	Location of Internet Platform						
			New Zealand	Australia	Singapore, USA (West Coast)	USA (East Coast)	East Asia	Europe	South Africa
			60ms	100ms	200ms	240ms	300ms	360ms	600ms
3.5G Wireless	20	400%	12%	8%	5%	5%	4%	4%	3%
ADSL2+	24	317%	11%	8%	5%	5%	4%	4%	3%
ADSL	8	1150%	19%	13%	7%	6%	5%	5%	3%
Rural	5	1900%	24%	17%	9%	8%	7%	6%	4%
Basic Broadband	1	9900%	168%	113%	63%	54%	44%	37%	23%

Table 2 confirms that, unsurprisingly, the greatest benefits from deployment of the 100Mbps network accrue to communications between sites within New Zealand, and to a lesser extent, between New Zealand and Australian sites. The benefits for New Zealand firms communicating via HTTP-enabled web pages with platforms in the USA, Asia, Europe and South Africa are very modest indeed. These findings suggest that the claims made that government-funded FTTH networks in Australia and New Zealand will serve to substantially increase the ability for local firms to compete with their foreign counterparts, and thereby enable a step-change in local economic growth relative to the potential of existing access networks, are not credible.

3. Policy Implications

The preceding calculations indicate that, to the extent that the development of commercial internet applications is resulting in an increasing number of applications based upon the HTTP protocol, the deployment of ever-faster local access networks will have an exponentially decreasing effect on page load performance (and by extension, direct productivity gains from using the application) for web page users the further the distance between them and the platforms with which they are transacting. The bulk of the benefits of super-fast networks pertain to communications between users within the same local access network environment. Investment in FTTH networks in Australia and New Zealand will thus yield the greatest benefits to firms transacting within the same region. There is nothing to

suggest that these networks alone will serve to substantially increase the competitiveness of local firms in an international context, especially if the firms in the foreign locations are also increasingly connected to ultra-fast networks to communicate with each other²⁷.

The inescapable fact for Australian and New Zealand firms is that the tyranny of distance prevails to the extent that technological constraints will almost inevitably compromise their ability to transport goods to distant markets, regardless of the nature of the goods (physical or information). Investment in local access networks will not ‘solve’ this problem as it cannot move the rest of the world any closer. Whilst the effect of ‘moving the internet closer’ to Australia and New Zealand in the form of extensive local caching or Content Distribution Networks may mitigate some of the problems of distance, its effect is confined to the distribution of static data. Where the information required is unique, and needs to travel across the world to complete the transaction, and timeliness matters, distance still imposes relative cost disadvantages.

To the extent that the constraints for HTTP-enabled transactions across fibre-optic cables are of a technological nature, is it feasible to look to technological ‘solutions’ to ‘level the commercial playing field’? One possible solution would be to increase the performance of signal ‘boosting’ on fibre-optic cables. While there may be some gains, foreign firms will also benefit from this technological change. The relative difference will decrease, but the Antipodean firms will never be on an equal footing with their rivals. Alternatively, a different approach to web page design could be adopted for Australian and New Zealand firms that minimises the number of international data movements. Even here, however, such customising may be adopted by competing firms as well if it serves to preserve their relative advantage. It is noted that changes in technologies may result in the foreign firms ‘recalibrating’ applications on the distant platforms to suit new local conditions. Antipodean firms lacking scale may have no choice but to use the foreign applications hosted on the foreign systems calibrated to optimise performance for their competitors.

The implications are, just as with the freezer ships, that competitive advantage lies not in the transportation technology itself, but in the fact that the economy in question has a competitive advantage in relation to the production of a good or service that is amenable to distribution via that transportation technology. For remote economies, the benefits derived from that advantage must be relatively impervious to the effects of transportation times. As long as there is a commercial advantage arising from faster transacting, firms located closer to the platforms will be more competitive in the markets for that good or service (i.e. it will

²⁷ We note that Grimes *et al.* (2009) show in the New Zealand context that there is no apparent improvement in productivity for New Zealand firms using fast internet connections relative to those using standard connections. Howell and Grimes (2010) offer a number of explanations based on the nature of applications as to why this finding might be plausible. This paper adds a further, technological, explanation for the Grimes *et al.* findings

always cost a New Zealand firm more to conduct a transaction, as the transaction times are longer).

Furthermore, the shift to cloud-based computing poses a restriction to New Zealand's ability to gain full benefit from future productivity-enhancing information systems. As more New Zealand-based transactions shift to cloud-based platforms to gain access to economies of scale in processing, there is a trade-off against the increasing costs of latency that come from these platforms being located in major off-shore population centres benefitting from agglomeration effects. As it becomes more cost-effective to develop cloud-based applications than bespoke applications shared amongst New Zealand firms, the risk is that both data and applications will move from New Zealand (with short transaction times enabled by either existing or planned fast broadband network access and low trunk transport demands) to more distant platforms, where fast access networks offer no material advantage and latency delays become more economically significant. In addition, such transitions will likely lead to further 'hollowing-out' effects as key resources such as human and physical information processing capital migrates to these distant information hubs, in a pattern already identified in many other sectors of the New Zealand economy (Sweet & Nash, 2007).

4. Conclusion

The overriding economic policy lesson for the Australian and New Zealand governments is that scale and distance cannot be ignored, even in respect of internet-enabled transacting. The economic and policy challenges in these countries are necessarily different from those in other localities. When analysing the benefits arising from investment in infrastructure, it is imperative that there is clarity about the extent to which that investment facilitates transacting within the local economy and where it facilitates international transacting. For small distant economies, investments focused upon improving local internet access infrastructure will have a material effect upon firms' economic performance only to the extent that those transactions are confined within the local economy – including on the basis of distributing locally-cached international data. Time-dependent interactive transactions will continue to be disadvantaged by distance. The claims that Australian and New Zealand government investment in ultra-fast local broadband access will facilitate a step-change in national economic performance fail to take account of the fact that whilst information might exhibit different economic characteristics, the technologies which transport and process it are still subject to decreasing returns arising from the unrelenting economics of scale, density and distance.

References

- Akamai (2010). *The state of the internet*. Second Quarter, 2010 v3(2). Available on <http://www.akamai.com/stateofhteinternet> retrieved 25 October 2010.
- Battersby, B. (2006). *Does distance matter? The effect of geographic isolation on productivity levels*. Treasury Working Paper, 2006-03, April, Treasury, Canberra.
- Battersby, B. & Ewing, R. (2005). *International trade performance: the gravity of Australia's remoteness*. Treasury Working Paper, 2005-03, June, Treasury, Canberra.
- Belshe, M. (2010). *More bandwidth doesn't matter (much)*. Google. Retrieved on October 22 <http://docs.google.com/a/chromium.org/viewer?a=v&pid=sites&srcid=Y2hyb21pdW0ub3JnfGRldnxneDoxMzcyOWI1N2I4YzI3NzE2>
- Blainey, G. (1966). *The tyranny of distance: how distance shaped Australia's history*. Sydney, Australia: Sun Books.
- Boulhol, H, de Serres, A & Molnar, M. (2008) *The contribution of economic geography to GDP per capita*, OECD Economic Department Working Papers, no. 602, OECD, Paris.
- Capital Market Development Taskforce (2009). *Capital Markets Matter*. Available on http://www.med.govt.nz/upload/71047/MDV6220_CMD_TombStone_04c.pdf retrieved 20 October 2010.
- Ernst & Young (1999). *The Knowledge Economy*. Retrieved October 20 2010 http://www.med.govt.nz/templates/MultipageDocumentPage_17263.aspx
- Evans, L. & Hughes, P. (2003), *Competition policy in small distant open economies: some lessons from the economics literature*. New Zealand Treasury working paper 03/31, Wellington.
- Glaeser, E., & Gottlieb, J. (2008). The economics of place-making activities. *Brookings Papers on Economic Activity*, Spring 2008: 155-239.
- Gregor, S., Fernandez, W., Holtham, D., Martin, M., Stern, S., Vitale, M. & Pratt, G. (2004), *Achieving Value from ICT: key management strategies*, Department of Communications, Information Technology and the Arts, ICT Research Study, Canberra.
- Grimes, A., Ren, C., & Stevens, P. (2009). *The need for speed: impacts of internet connectivity on firm productivity*. Motu Economic and Public Policy Research Working Paper 09-15. Available on http://motu-www.motu.org.nz/wpapers/09_15.pdf . retrieved 22 October 2010.
- Heatley, D., & Howell, B. (2010). *Price discrimination, structural separation and the diffusion of fibre broadband: policy implications for Australia and New Zealand*. Paper presented at the 1st Asia-Pacific Regional Conference of the International Telecommunications Society, August 26-28, Wellington, New Zealand.
- Howell, B., & Grimes, A. (2010). Productivity questions for public sector fast fibre network financiers. *Communications and Strategies* 78: 127-45.
- Howell, B. (2006). Submission on the Telecommunications Bill. Presented to the Finance and Expenditure Select Committee, 11 August 2006. Available on request from iscr@vuw.ac.nz.
- Lakshmirnarayanan, K. & Padmanabhan, V. (2003). *Some findings on the network performance of bandwidth hosts*. Proceedings of the 3rd ACM SIGCOMM conference on internet measurement. Miami Beach, Florida USA.
- Muller, G. & Williamson, A. (2005). *The New Zealand ICT sector: economic impact*. IDC Market Research report for the Hi-Growth Project.
- Ovum. (2005). *Productivity and Organisational Transformation: optimizing investment in ICT*. Commonwealth of Australia: National Office for the Digital Economy.
- Quah, D. (2002). *Technology dissemination and economic growth: some lessons for the New Economy*. Centre for Economic Policy Research. <http://www.cepr.org>.
- Rahman, J. (2005). *Comparing Australian and United States Productivity*. Economic Roundup Autumn 2005, Australian Government Publishing Service. Available on <http://search.informit.com.au/fullText;dn=371199774428891;res=IELBUS> (retrieved October 31 2010).

- Shapiro, C., & Varian, H. (1999). *Information Rules: a strategic guide to the network economy*. Boston, Massachusetts: Harvard University Press.
- Sweet, A., & Nash, M. (2007). *Firms' decisions on the location of economic activity – a survey*. A report for the New Zealand Treasury. <http://www.treasury.govt.nz/publications/informationreleases/locationeconomicactivity/fire-zus-econactivity-sep07.pdf> (retrieved 25 October 2010).
- Wilkie, J., & McDonald, T. (2008). *Economic geography and economic performance in Australia*. Economic Roundup 3. Australian Government Publishing Service. Available http://www.treasury.gov.au/documents/1421/PDF/01_Economic_geography_economic_performance_Australia.pdf (retrieved October 31 2010).