

Quantifying the broadband access bandwidth demands of typical home users

Warren Harrop, Grenville Armitage
Centre for Advanced Internet Architectures.
Swinburne University of Technology
Melbourne, Australia
{wazz, garmitage}@swin.edu.au

Abstract-We take a closer look at the probable demand for bandwidth that is presumed to motivate the deployment of new broadband access technologies. We construct plausible estimates of what a 'typical home' might need given two different service provision scenarios – managed bandwidth and best-effort statistical multiplexing. We estimate a household of five people requires between 58 and 113Mbit/sec if bandwidth is managed on a per-application basis. If statistical multiplexing is used, the same family's bandwidth requirements jump to between low hundreds of Mbit/sec to almost a Gbit/sec.

Keywords- Gigabit, home networks, bandwidth requirements, Internet applications

I. INTRODUCTION

There is a common and intuitively reasonable belief that dramatically higher-speed access technologies are needed to meet the bandwidth requirements of emerging, more demanding consumer entertainment devices. Many estimates argue for 100s or 1000s of megabit/sec rates to the home [1] [2] [3]. Nevertheless, there are a number of back-pressures acting against the deployment of high-speed optical and wireless last-mile technologies. Clearly there are costs associated with physically deploying new access technologies. In addition, the evolution of compression technologies is pushing bandwidth demands down, consumers are sensitive to pricing of internet service, and there are practical upper bounds on just how much content a typical family home can consume at any one time. Internet service providers (ISPs) also face the challenge of scaling their core networks to cope when hundreds of thousands of individual homes have 100Mbit/sec or 1Gbit/sec access links.

In this paper we take a closer look at the probable demand for bandwidth that is presumed to motivate the deployment of new access technologies. Our goal is to construct plausible estimates of what a 'typical home' might need given two different service provision scenarios – managed bandwidth and best-effort statistical multiplexing. By 'managed bandwidth' we mean that priority queuing (or similar) schemes are actively utilised to protect real-time interactive services from the fluctuating bandwidth demands of uncontrolled non-interactive services. By 'best-effort statistical multiplexing' we mean the internet's current regime of benign neglect, where burstiness is handled by queues or out-right over provisioning of link capacity. The former typically requires far less physical layer bandwidth (but more complexity within the network) than the latter. Our estimates

can then be used to create realistic proposals and expectations for roll-out of broadband services in the consumer market.

It is not sufficient to simply identify applications or circumstances that might, in principle, consume lots of bandwidth. It is necessary to identify applications that will cause a majority of people in a given area to take up (or demand) and pay for, broadband service. Put another way, it is not acceptable to roll-out gigabit/sec physical layers to an entire suburb or town if only a niche market (or early-adopters) place high financial value on the service. For this reason arguments such as 'some people will demand uncompressed high definition TV content' does not provide a sound economic argument for rolling out new access links and backend infrastructure to everyone.

Human history is replete with instances where the majority are happy with fairly modest content delivery. A modern example is Apple's iTunes [4]. Although an audiophile might decry the quality of compressed sound reproduction, the iTunes music/video store's success [5] argues that there are many people willing to trade perfect audio for other attributes (such as convenient portability and accessibility). Technical superiority is rarely a good indicator of likely value to (or acceptance by) the market. Betamax died and VHS survived, Laserdiscs did not displace the DVD, Super Audio CD (SACD) did not replace the CD and so on.

Consequently in this paper we focus on estimating the bandwidth demands of applications and usage patterns that are likely to reflect an incremental evolution of today's consumer behaviours. We steer clear of speculating on the bandwidth demands of non-existent applications and services that have no analog in services currently provided by other means.

The rest of the paper is structured as follows. In section II we identify three general usage patterns and in section III summarise currently available visual and aural content delivery techniques. In section IV we present calculations of bandwidth demand for a household of between one and five persons, in section V we discuss future work, and in section VI we conclude the paper.

II. THREE USAGE SCENARIOS

When estimating the bandwidth required to support a particular application it is important to decouple the processes of content creation, content delivery and content rendering. Content may be created in digital form for immediate dissemination or stored centrally for later retrieval. Content may be transferred for immediate rendering (such as video

display or audio playback) or to local storage for later rendering. Finally, the local storage and rendering device may be continuously connected to the home's broadband connection or connected only temporarily.

Focusing on the bandwidth demands placed on the home's broadband link leads to three usage scenarios:

- Content is downloaded at real-time
- Content is downloaded at slower than real-time
- Content is downloaded at faster than real-time

('Real-time' is the natural rate at which an particular application would transfer data to achieve its content delivery goals.)

A. Content is downloaded at real-time

In this scenario we focus on the fact that the content is consumed (rendered) as it is receive, thus delivery occurs at essentially the same rate as rendering/playback by the home user's equipment. This applies to content being created in real-time, or content streamed from a central storage facility. It is analogous to existing broadcast television and radio services, and covers interactive services such as video/audio conferencing and online computer games. Anytime the end-user receives and consumes content at the same rate it is being generated the network must support the application's natural bandwidth requirement (regardless of when the user chose to initiate the content delivery).

B. Content downloaded at slower than real-time

This scenario is likely when the home user plans ahead. The home user nominates the content they desire ahead of the time at which they wish to experience the content. Their local content rendering device (such as a video player, personal computer or audio player) would be programmed to download and buffer (store) the content for later replay (rendering). Because immediate play-back of the content is not required, the content may be downloaded at slower than real-time.

Planning ahead may be deliberate (such as a user-controlled video-on-demand system) or subscribed (such as an RSS feed [6] that pushes content out to subscribers on a schedule dictated by the content provider). Just how fast or slow the download can be depends on when replay will begin (since reply will, by definition, consume the locally stored data at 'real-time'). The primary requirement is that the download stays ahead of the playback (otherwise the rendering will stall).

An underlying assumption here is that a user's local storage device is connected to the internet for the entire period between initiating the download and initiating playback.

C. Content is downloaded at faster than real-time

Our third scenario may occur when content download is initiated to a local storage device that will be disconnected from the network at some time prior to content playback. The bandwidth requirement depends primarily on how long the storage device can remain connected to the network. Download must be faster than real-time if the time between deciding to download the content, and disconnecting the local storage device, is shorter than the content itself. A typical

example would be downloading a few hours of video or audio clips into a portable device (such as 'podcasts' into a portable mp3 player [4]) when one is hurrying to leave the house in the morning.

In principle there isn't really an upper bound on 'how much faster than real-time' is desirable or necessary. If the user wishes to download and then disconnect from the network for later playback, they'll consume whatever excess bandwidth we can provide in order to speed up the download process. However, if extra bandwidth costs extra money, users are likely to adapt their demands to some modest multiple of 'real-time' download speed. (Note that for video downloads, portable devices typically utilise lower resolution content than full-sized display systems, thus reducing the number of bytes downloaded in this scenario.)

III. LIKELY APPLICATION BANDWIDTH REQUIREMENTS

Of the previous section's three scenarios, 'slower than real-time' is encompassed by 'at real-time' delivery, whilst 'faster than real-time' is difficult to assess without evaluating customer sensitivity to price. Thus for this paper we will focus on the requirements of typical application under the assumption that content is being delivered 'at real-time'. We have broken down major bandwidth using applications into categories and will discuss each in turn. (Unless otherwise stated, all bandwidths are quoted without transport, link layer or any form of encapsulation protocol overheads.)

A. Visual Content

A typical household will consume one or more video channels for movies, television and interactive conference calls. Bandwidth requirements depend on the user's desired video quality the available compression schemes.

Judging from the market today's consumer is fairly happy with DVD-quality video rendered at PAL or NTSC quality. With well established codecs a PAL or NTSC-equivalent channel would easily consume less than 3Mbit/sec.

Of greater future interest is the bandwidth requirement imposed by 'High Definition' (HD) video. HD has been used to describe resolutions from 1280x720 to 1888x1080, substantially greater than PAL or NTSC resolution and much more suitable for large screen televisions and displays. We will make bandwidth estimates based on 1888x1080.

HD content delivery is most commonly achieved with the Digital Video Broadcasting Project (DVB) [7] standards. DVB defines the standard method for terrestrial transmission of digital television. For example, Australian television stations achieve approximately 20Mbit/sec from their 7MHz analog bandwidth allocations [8]. This can be used entirely for one channel of HD resolution MPEG2 [9] encoded video, at 20Mbit/sec. We will use 20Mbit/sec as our upper limit for a HD channel.

Apple are promoting a more advanced H.264 codec [11] capable of providing movie trailer content at 'full' HD resolutions of 1888x1080 [10]. A survey of the Apple content available shows average video data rates of around 9Mbit/sec. We will take 9Mbit/sec as a middle value for HD content delivery (on the basis that future codecs may reduce this rate even further).

Home video conferencing might well occur using TV-quality channels, and perhaps rarely with HD quality channels.

B. Audible content

In terms of two channel music-based audio, common lossy-codecs such as 'mp3', 'aac' and 'ogg' support a wide range of bitrates with perceived quality degrading as bit rate reduces. The point at which bit-rate reduction creates unacceptable audio is subjective. However, it is worth observing that Apple's iTunes online music store has created a successful market with the 128kbit/sec Advanced Audio Coding (AAC) format (over a 1 billion tracks downloaded as of February 2006 [5]). Due to the wide acceptance of this codec and bitrate, for our calculations we will take 128kbit/sec to be our base rate for stereo 'CD quality' music-based audio.

The current standard codecs for home and commercial cinema multichannel audio are DolbyDigital 'AC-3' (DD) [21] and DTS [13]. When used on DVD titles, DTS has bitrates of 768 to 1536Kbit/sec compared to DD's rates of 384 to 448kbit/sec. It is interesting to note that cinema standard bitrates for these two codecs has been 1103kbit/sec for DTS and 320kbit/sec for DD for many years. Millions of cinema goers world wide have been more than satisfied with 5.1 channels being presented at these (relatively) low bitrates.

Both organisations have multiple next generation codecs defined, either advanced high bit rate codecs or lossless-compressed formats. For example, the next generation DD codec has a maximum of 13.1 channels at a bitrate of 6Mbit/sec. A DTS HD codec using lossless compression provides up to 8 channels in 18Mbit/sec. Even if a format such as DD at 6Mbit with 13.1 channels is accepted into homes, it is unlikely that the average household could achieve more than one (if even this many) of these elaborate multi-speaker setups due to the space requirements.

Due to these limitations we will assume only one multichannel audio stream is requested by a home at one time, and assume all other streams are two channel. For our optimistic calculations we will assume that acceptable audio can be provided in 448kbit/sec while our pessimistic calculations will use the upper limit of DD's next generation codec running at 6Mbit/sec.

Voice over IP (VoIP) bandwidths are low even by current broadband standards because the single audio channel is restricted to human voice frequencies. Figure 1 shows that the majority of current voice codecs consume far less bandwidth than a plain, uncompressed G.711 codec (at 128Kbit/sec in each direction). VoIP could be expanded to include high quality audio using advanced codecs, in which case the raw data rate converges on the 128Kbit/sec mentioned earlier for 'CD quality' audio. This is a very pessimistic upper bound as 128kbit AAC also assumes stereo full frequency source content, where VoIP is clearly not. Even so, we will use 128Kbit/sec as our pessimistic upper limit VoIP bitrate.

For video enabled VoIP calls our upper limit assumes HD video is included along with a voice stream. A pessimistic estimate would use 20Mbit/sec in each direction for HD video conferencing. A more optimistic video conference would use PAL or NTSC quality video at 3 Mbit/sec.

C. Games

For performance reasons current online, multiplayer graphical computer games tend to be client-server in nature. Rapidly changing game-state is calculated centrally on a server and the resulting images rendered locally by each client. In general, even the fastest paced games (the first person shooter, FPS) sit comfortably under 200Kbit/sec per client in each direction [14]. This will constitute our optimistic bandwidth requirement for games.

Increasing complexity in games may drive bandwidth requirements higher. However, there are natural limits. Games must appeal to mass markets (and hence the lowest common access speeds across the market) and server hosts must support (and pay for) the aggregate bandwidth to all clients. (An extreme version would have the server rendering each player's video and streaming it out across the network. At most this would consume HD-quality data rates, but in practice the added encoding/decoding latency would make this approach unacceptable for interactive gaming.)

D. Web surfing / File transfer / Email

Web surfing is an elastic application – much of the text and static image content may be downloaded at variable speeds, making use of whatever excess available capacity exists on the network. Text and images often compress well and viewing web pages is, at a bandwidth level, 'bursty'. Long-term bandwidth requirements are low. Instantaneous bandwidth needs may be quite high to download page elements (content) in reasonable time. With many webpages measured in the 10s of Kbytes we will assume that web content transfers can be interleaved with other non-real time concurrent file transfers, allowing us to set aside one bandwidth block for both in our calculations.

File transfers can tolerate both long Round Trip Time (RTT) on individual packets, short term delays and in the extreme, resumed transfer if needed. To download HD content at 20% of real time speed, 4Mbit/sec is required. This still allows a reasonable speed for file transfers and makes the average web page download in a small number of seconds.

Email is a relatively delay tolerant application and most emails are far less than 10Mbytes long. To transfer a 10MByte email within 1 minute, a speed of 2Mbit/sec will suffice.

E. Future applications

In this section we touch on future applications that have not yet seen mass take up and attempt to quantify some of their possible bandwidth requirements.

Body telemetry

A possible future application with wide community application is remote body telemetry monitoring. Even in the absence of data aggregation at the personal collection device the bandwidths required are small. Pediatric ECGs should be sampled at or slightly higher than 1000Hz [16] to retain diagnostic value while digital sample rates at 250Hz with 10

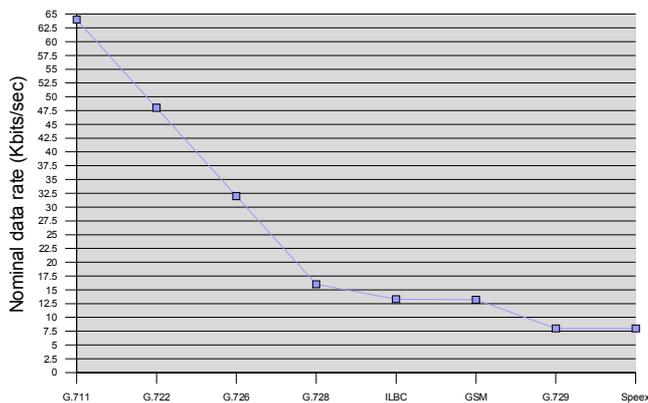


Figure 1 Bandwidth requirements for voice codecs

bit samples is suggested for diagnosis of EEGs [15]. Even quantizing and oversampling these signals aggressively leaves us with bandwidths less than that of a standard phone call. (We could imagine more extreme scenarios where invasive body monitors generated Mbit/sec of traffic. However, such applications would ultimately be limited by the bandwidth capabilities of the low-powered, battery-operated wireless device required to link the customer to the network.)

Home telemetry

The highest bandwidth application we have defined is that of HD video content, and its use in home telemetry is limited. Home security applications may need or benefit from HD video for intrusion detection, this can easily be achieved with on-site video processing. High bandwidth would only be required when initial sensors are set off and off site processing in real time is needed. When this occurs the HD monitoring would pre-empt all other applications in the house, so we do not need to consider the monitoring channel's bandwidth requirements to be in addition to everything else. All other monitoring applications, such as environmental conditions appliance communications are either non-real time (and can use off peak times to transfer data) or have trivial bandwidth requirements.

Ultra HDTV

Research into 7680 x 4320 pixel Ultra HDTV (UHDTV) is underway in Japan [17]. While it is speculated that the bandwidth of UHDTV will be in the realm of 400Mbit/sec the question begs if this is 'too much' resolution. The bandwidth of a human eye-brain connection has been recently estimated at 10Mbit/sec [18]. While this does not mean that any video presented at 10Mbit/sec or above will exceed the acuity of a human eye with 20/20 vision it does indicate that there are certain limits to the human mind's ability to absorb content and moving beyond these provides no perceived benefit to the watcher. If there is no perceived benefit to the watcher then they are unlikely to pay the premium to receive it. [19] suggests that the optimal viewing distance for current HD content (at 1888x1080) is 3 to 4 times the distance of the height of the screen. With UHDTV resolutions and average sized family rooms it is an open question as to if the visual acuity of the human eye is not only reached but exceeded for the average viewer. Technologies that waste visual resolution beyond average human perception could be treated as those that waste audio resolution. There have also been questions raised about the reaction of people to such high bandwidth displays and how wide spread the possibility of motion

sickness related feelings might be when viewing this form of content.

Tele*

Telepresence, immersive teleworking and virtual reality applications are in their infancy at the current time and are difficult to speculate on, but even with future developments they are still limited by the bandwidth of the human senses. We have already found that these requirements are not great even under current codec implementations and the utility of higher quality implementations is questionable.

Other senses

There is currently research into remote touch applications [20] but the bandwidth requirements of such systems at this point in time are low. Physical touch, taste and smell senses are likely to have even lower requirements than audio stimuli, and hence bandwidth requirements can be estimated as lower than 128Kbit/sec (even without knowing the details of the future touch applications).

Because all of the future applications we have defined above have either questionable take up possibilities and/or very low-bandwidth requirements in comparison with a channel of HD content we will leave them out of our worst case calculations. We will assume that if they or similar applications are wildly deployed that their bandwidth requirements will be dwarfed by other application's bandwidths.

IV. DISCUSSION

A. Family bandwidth requirements

We can now combined the usage scenarios and application bandwidths defined above to create a set of instantaneous bandwidths that provide a pessimistic upper bound of bandwidth consumption for a larger than average family, with larger than average bandwidth requirements, at a peak usage time of day. We also create a formula for a more optimistic view, where where we assume usage of next generation codecs. HD content is by far the most bandwidth intensive application we have defined and what we will assume to be our worst case application that all household members are primarily using.

For our pessimistic scenario (calculations using current codecs bandwidth requirements) we dedicate 20Mbit/sec for a HD video stream, it's associated audio stream is 6Mbit/sec for the first user and 128kbit/sec for subsequent users. 2Mbit/sec is allocated for email communication and 4Mbit/sec for file sharing and traditional web browsing. For our optimistic calculations using future codecs) we dedicate 9Mbit/sec for a HD video stream with other bandwidth allocations staying the same.

When we extrapolate the bandwidths to multiple members of a household. We will assume household members ranging in number from 1 to 5 - all at an age where they can consume HD video content. We also take their use at its peak, with each household member individually using bandwidth and not sharing content with any other members of the household. This now gives us the following formulas.

$$BW_{\text{Pessimistic}} = 20x + .13(x-1) + 6 + 2 + 4$$

$$BW_{\text{Optimistic}} = 9x + .13(x-1) + 6 + 2 + 4$$

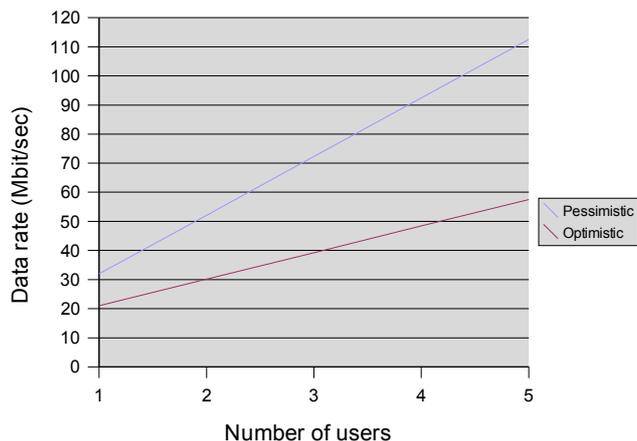


Figure 2 Household users vs bandwidth consumption using current (pessimistic) and future (optimistic) codec rates.

Figure 2 is the graph for these formula showing number of household users against bandwidth with both pessimistic and optimistic usage scenarios. Our highest projected bandwidth with absolute worst case situation is 113Mbit/sec, well below 1Gbit/sec. This drops to 58Mbit/sec when we assume only 9Mbit/sec HD content.

These numbers seem fairly low compared to 1Gbit/sec even though our upper band numbers are worst case in a number of ways: they assume that compression will not gain efficiencies (due to CPU speed increases) over time. They assume a household of 5 and they assume no prior download of content. Finally they assume no sharing of content between users (eg two or more users watching the same content).

B. Managed bandwidth versus Statistical multiplexing

Our estimates in the previous section make a key simplifying assumption – the home broadband service provides ideal isolation between traffic belonging to different classes of application. In other words, we have assumed that some form of quality of service mechanisms are deployed to ensure bursts of packets from one application do not temporarily 'starve' other applications (such as HD channels being rendered in real-time). This can only really be achieved if the ISP deploys, and manages, moderately sophisticated queuing and traffic classification in their routers.

A more likely situation is that an ISP will utilise statistical multiplexing to 'manage' the interactions between packets belonging to different application flows. Statistical multiplexing manages through benign neglect and over-provisioning. Enterprise IP network operators often use a rule-of-thumb that the link speed should be 2 to 5 times the average capacity. In this case our pessimistic estimate would require something in the order of one Gbit/sec at the link layer to support the proposed mix of application traffic.

V. CONCLUSION AND FUTURE CONSIDERATIONS

We have considered the probable demand for bandwidth that is presumed to motivate the deployment of new broadband access technologies. We constructed plausible estimates of what a 'typical home' might need given two different service provision scenarios – managed bandwidth

and best-effort statistical multiplexing. With a family of five, all consuming high quality HD content at the same time, our base bandwidth requirement ranges from 58Mbit/sec to 113Mbit/sec. This is the bandwidth that would be required from a physical layer technology if IP layer bandwidth management techniques are deployed at either end of the access links. If 'best effort' statistical multiplexing is used to share capacity across each home's applications, the physical layer requirements tend one gigabit/sec under the most pessimistic estimations.

The current trends in silicon processing speeds allows for increasingly improved codecs to be used at end nodes. Combined with our stated bandwidth consumption numbers for various services, it seems that the roll out of very high bandwidth services will need to be considered carefully.

In the future we intend to consider the level of asymmetry in access link bandwidths that may be required or tolerated by consumer applications.

REFERENCES

- [1] Tezuka H., Suda Y., Ishikawa H., Shibusaki M., "A study on large capacity file transfer with gigabit to the home (GTTH) system", Fifth Asia-Pacific Conference on Communications and Fourth Optoelectronics and Communications Conference (APCC/OECC '99), Volume 1, 18-22 Oct. 1999
- [2] Shioiri, S. et al, "A 2.5-Gb/s high-speed PLC module for Gigabit-To-The-Home (GTTH) system", Optical Fiber Communication Conference and the International Conference on Integrated Optics and Optical Fiber Communication, Volume 3, 21-26 Feb. 1999
- [3] "IEEE 802.3 10Gb/s PHY for EPON (Ethernet Passive Optical Network) Study Group" http://grouper.ieee.org/groups/802/3/10GEPON_study/index.html, August 2006
- [4] "Apple - iPod + iTunes", <http://www.apple.com/itunes/>, August 2006
- [5] "iTunes Music Store Downloads Top One Billion Songs", <http://www.apple.com/pr/library/2006/feb/23itms.html>, August 2006
- [6] "Really Simple Syndication", <http://www.rssboard.org/>, August 2006
- [7] "DVB - Digital Video Broadcasting", <http://www.dvb.org/>, August 2006
- [8] "Digital Broadcasting Australia", <http://www.dba.org.au/index.asp?sectionID=120>, August 2006
- [9] "Moving Picture Experts Group", <http://www.chiariglione.org/mpeg/>, August 2006
- [10] "Apple - Movie Trailers", <http://www.apple.com/trailers/>, August 2006
- [11] "H.264 : Advanced video coding for generic audiovisual services", <http://www.itu.int/rec/T-REC-H.264>, August 2006
- [12] "Dolby Laboratories Inc.", <http://www.dolby.com/>, August 2006
- [13] "DTS", <http://www.dtsonline.com/>, August 2006
- [14] G.Armitage, M.Claypool, P.Branch, "Networking and Online Games - Understanding and Engineering Multiplayer Internet Games," John Wiley & Sons, UK, April 2006 (ISBN: 0470018577)
- [15] Y. Ohtaki, K. Toraichi, Y. Ishiyama, "On compressing method of EEG data for their digital database", IEEE International Conference on Acoustics, Speech, and Signal Processing, 23-26 Mar 1992
- [16] P. R. Rijnbeek, J. A. Kors, M. Witsenburg, "Minimum bandwidth requirements for recording of pediatric electrocardiograms", *Circulation*. 2001;104:3087.
- [17] "Japan demonstrates next-gen TV broadcast", <http://www.eetimes.com/showArticle.jhtml?articleID=173402762>, August 2006
- [18] Koch et al., "How Much the Eye Tells the Brain", *Current Biology* 16, pp.1428-1434, July 25 2006
- [19] J. Hamalainen, "HDTV displays: How good do they need to be?", <http://broadcastengineering.com/hdtv-displays/>, August 2006
- [20] Lee, S, et al, "A mobile pet wearable computer and mixed reality system for human-poultry interaction through the internet", *Personal and Ubiquitous Computing*, November 2005