SWINBURNE UNIVERSITY OF TECHNOLOGY

Driving the demand for high speed consumer access

24 June 2007

Grenville Armitage

garmitage@swin.edu.au

Will bandwidth demand ever stop?

Today I want to discuss why bandwidth to the home is not driven by unbounded demand

This will involves some arm-waving, and back-of-the-envelope guesstimates

Hopefully provoking some thought on what unknown service will drive further demand for bandwidth...

(A peek at my conclusion....)

Conclusion: A *typical home* will not need more than 60Mbit/sec inbound and 15 Mbit/sec outbound

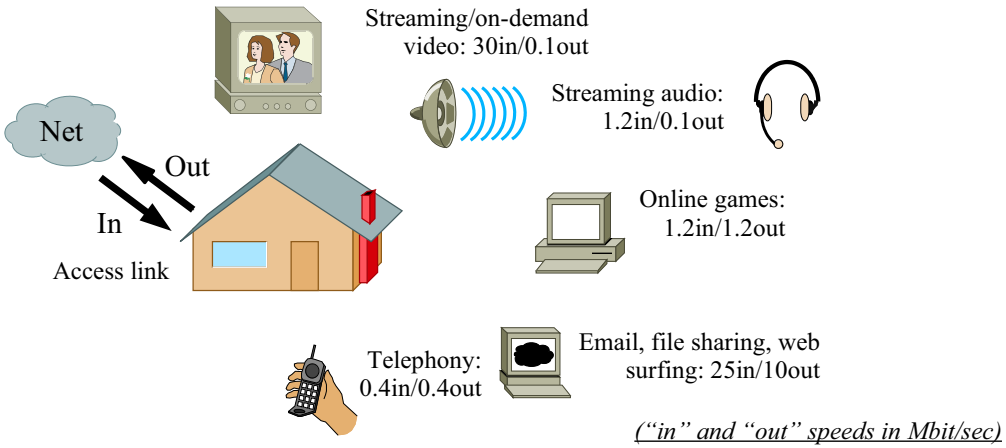

("in" and "out" speeds in Mbit/sec)

(A peek at my conclusion....)

Conclusion: A *typical home* will not need more than 60Mbit/sec inbound and 15 Mbit/sec outbound

We will need optical technologies that deliver 10s of Mbit/sec per home while cheap to manufacture, deploy and operate at metropolitan scales

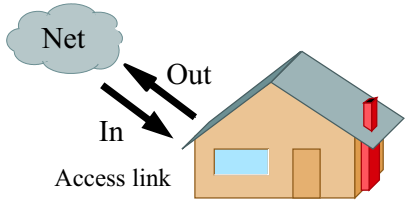

Also: W. Harrop G. Armitage, "Quantifying the broadband access bandwidth demands of typical home users," Australian Telecommunication Networks and Application Conference 2006, Melbourne, Australia 4-6 December 2006.

Comfort from previous predictions

"I think there is a world market for maybe five computers." (Thomas Watson, chairman of IBM, 1943)

"There is no reason for any individual to have a computer in their home." (Ken Olsen, president, chairman and founder of Digital Equipment Corp., 1977)

"640K ought to be enough for anybody." (We'd like to think it was Bill Gates, 1981)

(numerous sources, including "31 Of The Worst Predictions", http://www.singlegrain.com/blog/31-of-the-worst-predictions/, 23May2007)

With that in mind, here's my analysis....

What are people doing in their homes?

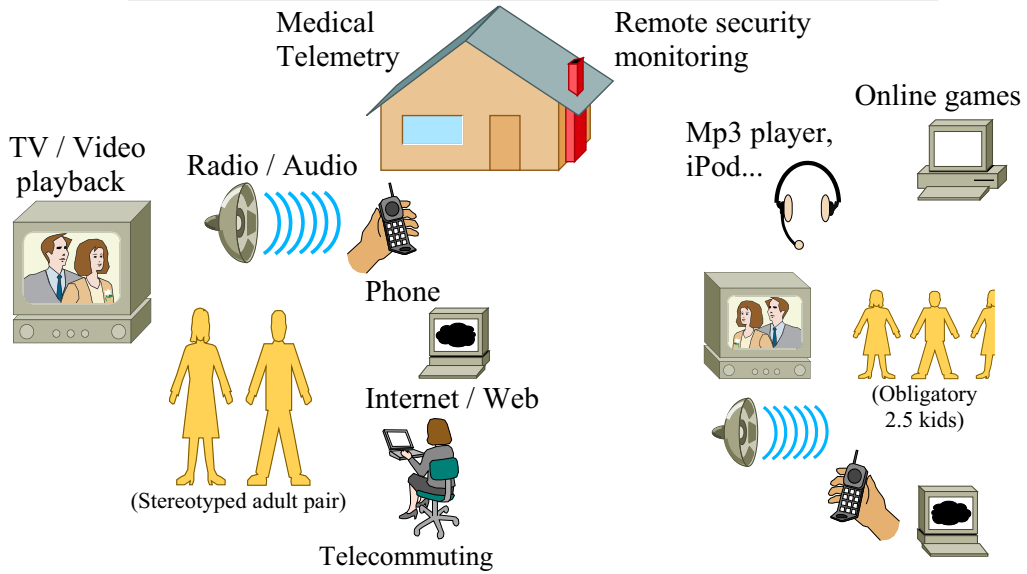

Driving bandwidth at the edges

Activity at the edge drives cross-core traffic

Businesses

VPNs, data sharing, replication, backups, B2B, remote access...

Consumers

Entertainment, remote VPN access to work, online transactions, phone/conferencing with friends/family, games...

Symmetric traffic (e.g. Telecommuters pull/push)

Asymmetric traffic (content providers to consumers)

Transit (e.g. businesses supplying VoIP/conferencing connectivity, Instant messenger services, etc)

Pressures on consumer demand

- Upward pressure
  - Desire
  - Convenience
  - Capability
- Downward pressure
  - Financial costs
  - Availability / (in)convenience
  - Ability to consume

What tends to drive big estimates?

- "Deliver all media over one digital link"
  - Replace existing RF, coax, HFC, POTS loop with a common physical/access link (e.g. Optical loop)

Service	Today	Future
TV (free / cable)	50 – 450MHz (air) 50 – 750MHz (HFC)	Packet switched delivery (all channels)
Radio (free)	0.5 - 1.6MHz (AM) 88 – 108MHz (FM)	Packet switched delivery (all channels)
DVD / movies	Cable TV Sneaker-net (rentals)	Packet switched delivery (on demand)
CD / songs	Radio / Internet	Packet switched delivery
Interactive games	Buy / Rental / Internet	Packet switched delivery

Some changing realities

- Today's TV and radio are intrinsically broadcast or multicast
  - RF 'access link' carries all channels all the time
  - Channel selection: filtering at the receiver
  - Apparent equivalent digital bandwidth is massive
- Packet switched delivery of video and audio can be unicast or multicast
  - Channel selection implemented at (or near) the source
  - Customer link only carries content while channel selected
  - Substantially reduced bandwidth required at last hop (proportional to number of active channels at one time)

Changing realities #2

- Previously
  - Broadcast TV / radio requires no upstream signaling
  - The public (consumers) were not content creators (no upstream distribution channel)
- Now
  - Upstream channels allow distributed control of content delivery (e.g. 'channel selection' applied at source)
  - The public can be (and are) creators – peer2peer file sharing, webcasts, multiplayer games, telecommuting, ...
  - Yet asymmetric link technologies are still a reasonable match to typical usage patterns

Driving the need for delivery speed

- The last 20+ years have seen a dramatic evolution in cheap, large-scale data storage and processing capabilities
- Content creation, delivery and presentation may be (and often are) decoupled
- Compression techniques getting better, driving down the *necessary* bandwidth for given experience

Creation, delivery and presentation

- Creation:
  - Audio/video from live action, Studio-based TV, retrieval from DVD / HD storage, game servers
- Delivery:
  - Packet switched transmission
- Presentation:
  - Media players, TV, Hi-Fi Stereo systems, iPods, mp3 players, online games, on-screen chat room interfaces, news websites, etc

- Creation and presentation occur at the network 'edges'
- Delivery is where network resources are consumed
- Bandwidth requirements during delivery depend on the temporal relationship between creation and presentation

Further decoupling of each step

- Creation ► immediate or delayed delivery
  - Real-time streaming (immediate delivery)
  - Participatory virtual game worlds (immediate delivery)
  - Retrieve-on-demand, such as VoD or common web surfing (delayed delivery)
- Delivery ► immediate or delayed presentation
  - Real-time streaming or watching 'on demand' content (immediate presentation)
  - Participatory virtual game worlds (immediate presentation)
  - Download into (optionally portable) device for later use (delayed presentation)

Or more succinctly...

- *Delivery* may be required at:

  Faster than, at or slower than real-time

- 'Faster than' can imply *any* attainable speed...
  - Pragmatically it will be bounded by available technology and a consumer's personal cost-benefit trade off
- The last two are bounded by content encoding and a person's ability to consume the presentation
  - Let's consider these in more detail....

What can a typical home *consume?*

- 'Slower than' is upper-bounded to 'at real-time'
- Just how much content can a home consume at real-time?
- Constraints include:
  - Number of occupants
  - Number of rooms
  - Concurrent visual media consumption
  - Concurrent audio media consumption
  - Concurrent telemetry, immersive 'game' participation, etc

A person consuming visual content

- Not an unbounded system
  - We tolerate various sizes and quality (vs cost or convenience)
- Visual resolution limits
  - Ability to resolve pixels
  - Screen sizes limit the required resolution
  - Home theater screens in *typical* homes are bounded by possible seating arrangements

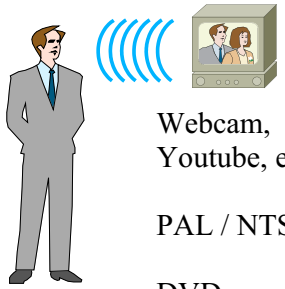

Webcam, Youtube, etc

PAL / NTSC

DVD

High Def

Home Theater

Some visual limits

- Bandwidth of a human eye-brain connection ... ~10Mbit/sec

  Koch et al., "How Much the Eye Tells the Brain", Current Biology 16, pp.1428-1434, July 25 2006

- Optimal viewing distance for current HD content (at 1888x1080) is ~3 to 4 times the height of the screen
  - "HDTV displays: How good do they need to be?" 24 May 2004. http://broadcastengineering.com/hdtv-displays/

Video on Demand / Streaming

- Consumed in real-time, and no need to deliver faster than it is consumed
  - If you need it now then download in real time, if you need it later then we download at ≤ real time
- Compression is common and improving:
  - E.g. Apple's H.264 literature claims: 1-2Mbps for 640x480@24fps, 5-6Mbps for 1280x720@24fps and 7-8Mbps for 1920x1080@24fps
    - http://www.apple.com/quicktime/technologies/h264/ (24jun07)
- Perhaps 3 concurrent HD channels? < 30 Mbit/sec
  - Or perhaps 3 concurrent regular TV channels and some 'youtube'-style video? < 10 Mbit/sec

A person consuming audio content

- Not an unbounded system
  - We tolerate various levels of quality and 3D immersion (vs cost or convenience)
- Spatial and frequency resolution limits of hearing
  - How many channels?
  - In how many rooms?

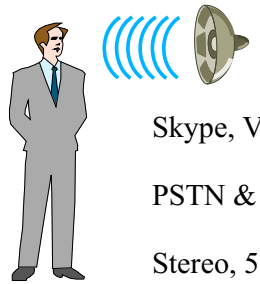

Skype, VoIP, etc

PSTN & mobiles

Stereo, 5.1, 7.1, ...

Home Theater

Streaming audio

- Consumed in real-time, and no need to delivery faster than it is consumed
  - If you need it now then download in real time, if you need it later then we download slower than real time
- Under 0.5 Mbit/sec for ridiculously good quality
  - Large market happy with CD-quality @ 128Kbit/sec
    - (smaller markets of audiophiles notwithstanding)
- Perhaps 5 concurrent channels? -> 2.5 Mbit/sec

Voice and Games

- Voice over IP / Telephony
  - Primary requirement is low loss/latency
  - Worst case ~100kbit/sec per active call
  - perhaps 4 concurrent calls? -> 400kbit/sec
- Online games
  - During periods of activity the multiplayer games typically aim for well under 100kbit/sec per player
  - During game update (non-playing) periods perhaps one or two megabit/sec is desirable, but 100-300kbit/sec tolerable
  - Perhaps 4 concurrent players? -> 400 to 1200kbit/sec

Email and File Sharing

- Email
  - Fast as possible, but we tolerate 10s of seconds delay
  - e.g. 5 Mbyte email in 30 secs (1.3 Mbit/sec) is okay
  - Allow extra, so perhaps 4 Mbit/sec per user and 2 concurrent emails at any given time -> 8 Mbit/sec
- File sharing
  - Fast is good, but tolerate 1000s of seconds delay
  - e.g. 400 Mbyte file in 30min (1.8 Mbit/sec) is okay
  - So perhaps 4 Mbit/sec per user and 4 concurrent full-rate sharing at any one time -> 12 Mbit/sec

Web surfing and ???-casting

- Web surfing
  - Goal is to render pages 'fast', but diminishing returns
    - Does the user care about 0.1sec vs 0.3 sec render time?
    - For pages under 80Kbyte in 0.3 sec -> 2.1 Mbit/sec
    - Allowing for multiplexing of users, perhaps 5 Mbit/sec ?
      - Linked content falls under previous categories
- ???-casting (podcasting, etc....)
  - Download content for later playback/consumption
    - “Need it now, am going mobile in XX minutes”
    - Genuine case for download speeds faster than rate at which the content can be consumed by human
  - Difficult to characterise realistic bandwidth need...

A modest house, perfect b/w sharing

• Voice over IP	0.4 Mbit/sec
• Online games	1.2
• Video on Demand	30
• Streaming Audio	2.5
• Email	8
• File sharing	12
• Web surfing	5

A grand total of ~60 Mbit/sec

(assuming worst-case high quality video, no guesses for pod-cast style downloads, all services in use at once)

(and yes, I'm waving my arms vigorously – approximations abound....)

Is perfect bandwidth sharing realistic?

- The preceding estimate assumes each class of application has a protected slice of bandwidth
  - Protect UDP-based apps that use only what they need
  - Isolate TCP-based apps that otherwise opportunistically consume whatever bandwidth is 'free'
    - Web surfing, email, peer2peer, burst download for later playback,....
- And that implies limited QoS on the access link
  - .... is IP QoS on consumer access links realistic?
    - Certainly not over wireless -> need isolated (copper or optical) infrastructure

What about stat muxing?

- Statistical multiplexing, aka Benign Neglect
  - Over-provisioning is the usual solution, spare / unused bandwidth to cope with random, correlated traffic bursts
  - Previous estimate up by factor of 5-6? -> 360+Mbit/sec
- Yet opportunistic traffic flows may still cause transient starvation of capacity to others
  - Impacting concurrent video, audio, interactive games, etc
- We probably still need basic IP rate-caps per application class
  - Which implies based IP QoS on access links anyway

How much for Faster than real-time?

- A difficult category to predict
- Download content into portable device, quickly, so delayed presentation (playback) can occur while disconnected from the network
  - e.g. Podcasts, personal CD collection, ...
  - Download size bounded by presentation medium
- What is this worth to a customer?
  - ~1.5Mbit/sec to load 6Hrs of 128bit/sec audio in 30min
  - ~15Mbit/sec to load 6Hrs of 128bit/sec audio in 3min
  - What would customer pay for higher-speed physical layer just so they can burst for 3 – 30min once or twice a day?

What *could* I do with 20:1 Mbit/sec ?

- 20 down and 1 up is plausible with ADSL2+
- With bandwidth protection between classes of IP traffic, we could provide:
  - 4 concurrent TV channels (H.264) (10 Mbit/sec down)
  - 2 concurrent 'radio' stations and 2 high quality streaming audio stations (1.5Mbit/sec down)
  - 2 concurrent VoIP calls (0.2Mbit/sec up/down)
  - 2 concurrent online game players (0.2 Mbit/sec up/down)
  - ~8Mbit/sec down and 0.6Mbit/sec up for other elastic, non-realtime applications (email, web, etc...)
- So how do we generate more demand for higher bandwidths than ADSL2+ on the link to the home?

The concluding challenge ?

- More raw speed isn't the issue
  - Compression algorithms, smart endpoints, distributed control of delivery have seen to that
  - There's no killer-app waiting to push home access needs over 100Mbit/sec (except lack of per-class QoS control)
- The challenge is delivering cheap, controlled packet -switched bandwidth over the consumer's last mile
  - Optical networking needs to keep pushing costs down (to manufacture, deploy and own/operate) for modest (10s of Mbit/sec) infrastructure at metro scales

Thank you for listening....