The Video Tsunami: Internet Television, IPTV and the coming wave of Video on the Internet

Marshall Eubanks

AmericaFree.TV

tme@americafree.tv

March, 2008

Some Quotations

- In North America, Internet video has jumped from 10 percent of consumer Internet traffic in 2006 to 24 percent of traffic in 2007.
  - » Cisco Systems White Paper, ***The Exabyte Era***, January 14, 2008.
- Here's a prediction - by 2010, we will see several examples of cable operators cutting back on channel counts to turn more bandwidth over to broadband.
  - » Niel Weinstock, ***TVB / Television Broadcast***, February, 2008.
- In the years ahead, broadband on the computer will be the primary source of entertainment for kids
  - » Bob Iger, CEO, Disney Corporation, March 12, 2008.
- Video Road Hogs Stir Fear of Internet Traffic Jam
  - » Steve Lohr, ***New York Times***, March 13, 2008

March, 2008

The New York Times was thoughtful enough to motivate my talk…

HOME PAGE | MY TIMES | TODAY'S PAPER | VIDEO | MOST POPULAR | TIMES TOPICS

The New York Times **Technology**

WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HEALTH | SPORTS | OPINION

Search Tech News & 8,000+ Products Go

Browse Products -- Select a Product Category -- Go

Video Road Hogs Stir Fear of Internet Traffic Jam

By STEVE LOHR
Published: March 13, 2008

Caution: Heavy Internet traffic ahead. Delays possible.

E-MAIL
PRINT
SINGLE PAGE
REPRINTS
SAVE
SHARE

Multimedia

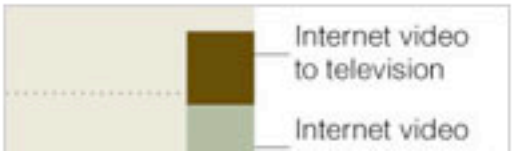

For months there has been a rising chorus of alarm about the surging growth in the amount of data flying across the Internet. The threat, according to some industry groups,

March, 2008

Agenda

- Definitions, History and recent events.
- Traffic Estimates Projections
  - Sources
  - Comparing Video, P2P, Web
- The “Long Tail”
  - Zipf’s Law, Pareto distributions and their implications
- Can Multicast Help ?
- Statistics from my Internet Television service.
- A look at current Video technology and the need for FEC.

The Video Tsunami

- What is going on ?
  - *Nobody really knows*.
    - But you can make some educated guesses.
- Video traffic (& bandwidth) is greatly increasing
  - Movies, TV shows, User Generated Content (UGC), Telepresence.
- Company and investor and content creator interest in greatly increasing.
- Unicast, Multicast, Peer to Peer (P2P), Push, Pull
  - All will be examined in turn.
  - I will mostly leave P2P to the next speaker…

March, 2008

IPTV versus Internet Television

- This nomenclature is still fluid…

  … so I am going to try and fix it.
- IPTV : Distribution of video over local networks by the local provider using the Internet Protocols (IP)
  - Typically a replacement of existing Hybrid Fiber Cable (HFC)
  - Over a network, but not generally over the Internet
  - Generally a Set Top Box (STB) is involved
  - First used MPEG-2 transport streams, then Ethernet, now IP Multicast is becoming common.
  - Also called “switched video” in the Cable / Telco world.
- Internet Television : Distribution of video channels to end users over the Internet
  - At present, generally a STB is **not** involved, playout is on a computer.
- Will they merge ? **Yes** (in my opinion).
  - I will present my reasoning in due course…

March, 2008

Not too long ago…

- Television was broadcast over the air or through co-axial cables and the Internet was something people did with phone lines or T1s.
  - Channels were limited, and you watched what was provided, when it was provided.
- The IETF has been concerned with video broadcast over the Internet for a long time.
  - The world's first "public" broadcast over the Internet was IETF-24, July 1992

    (I think.)
  - Many WG have played or are playing a role in the development of Internet video, including AVT, MMUSIC, MBONED, RMT, FECFRAME.

In this Century

Recent developments leading to the current video explosion have included

- The development of ever better codecs, from H.263 (Youtube Flash), MPEG-4 and others, leading to H.264 (2003) with a factor of 4 or improvement of video quality over MPEG-2 (1988)
- The common availability of broadband to the home and office.
- The restrictions placed on centralized social P2P systems such as Napster lead to the development of true P2P transport mechanisms, such as BitTorrent.
  - It is much easier to attack a server with a central server, such as Napster, compared to one without, such as BitTorrent.
- This has lead to an interesting conflation of a transport mechanism and a social movement.
- The availability of venture capital which lead to the rapid funding of almost any video venture imaginable.

Current Trends

- "Cable," DSL and Fiber systems are moving to IP Multicast deployments.
  - In walled gardens.
- P2P video has taken off in a big way, causing some providers to wage war against it.
- The "Long tail" has exploded.
  - To understand this it will be necessary to look a little into Zipf's Law and Pareto distributions, the mathematical basis for the Long Tail.
- But first, let's consider YouTube and general traffic patterns
  - YouTube simply provides a means for people to host video content they source themselves. YouTube built it and they did come…
  - The vast majority of the content is either amateur video straight from the camera or stuff "borrowed" from commercial TV.
  - They are making a lot of money from the Long Tail

Internet Video is Certainly In the News

Disruptive Telephony

Dan York on how Voice over IP is rewriting (almost) everything you thought you understood

Disruptive Conversations | Disruptive Telephony | Blog.DanYork.com | Blue Box: The VoIP Security Podcast | Voice o

« Today's Squawk Box podcast - A conversation with Jeff Pulver about VON.x | Main

March 07, 2008

The Oprah-tization of VoIP via Skype!

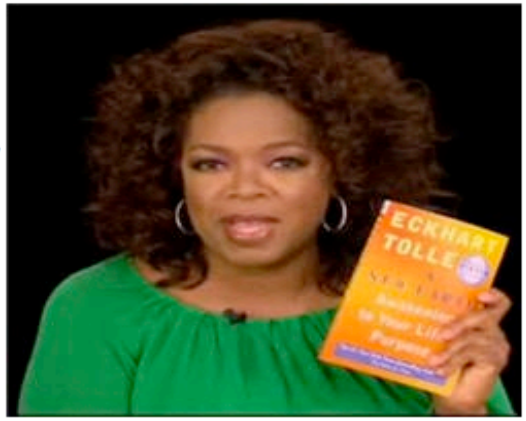

Did you ever think that we would utter "Oprah" and "VoIP" in the same sentence?

Well, we are... or more precisely "Oprah" and "Skype". If you have no idea what I am talking about, you need to head over to Skype Journal and read Jim Courtney's piece: "*Skype Sponsors Oprah's 'A New Earth' Web Event*"

Essentially, Oprah has taken her Book Club online to host a weekly web collaboration session for the next *ten weeks* with, oh, ***750,000*** members of her club! More information - and an audio introduction from Oprah - is available on the Skype campaign page. Monday night was the first first session and as Jim subsequently noted (as did Howard Wolinsky on the new Skype US blog), the session didn't go so well with regard to technical issues. As Oprah's company, Harpo Productions (Hint: spell "Harpo" backwards) said in their statement:

> *Monday night's webcast was one of the largest single online events in the history of the Internet.* ***More than 500,000 people simultaneously logged on to watch Oprah Winfrey and Eckhart Tolle live, resulting in 242 Gbps of information moving through the Internet.*** *Unfortunately, some of our users experienced delays in viewing the webcast. We are working to identify the specific causes for the problems experienced and will work diligently to rectify them.*

The Growth of YouTube.com (from Alexa)

~ Monthly Audience

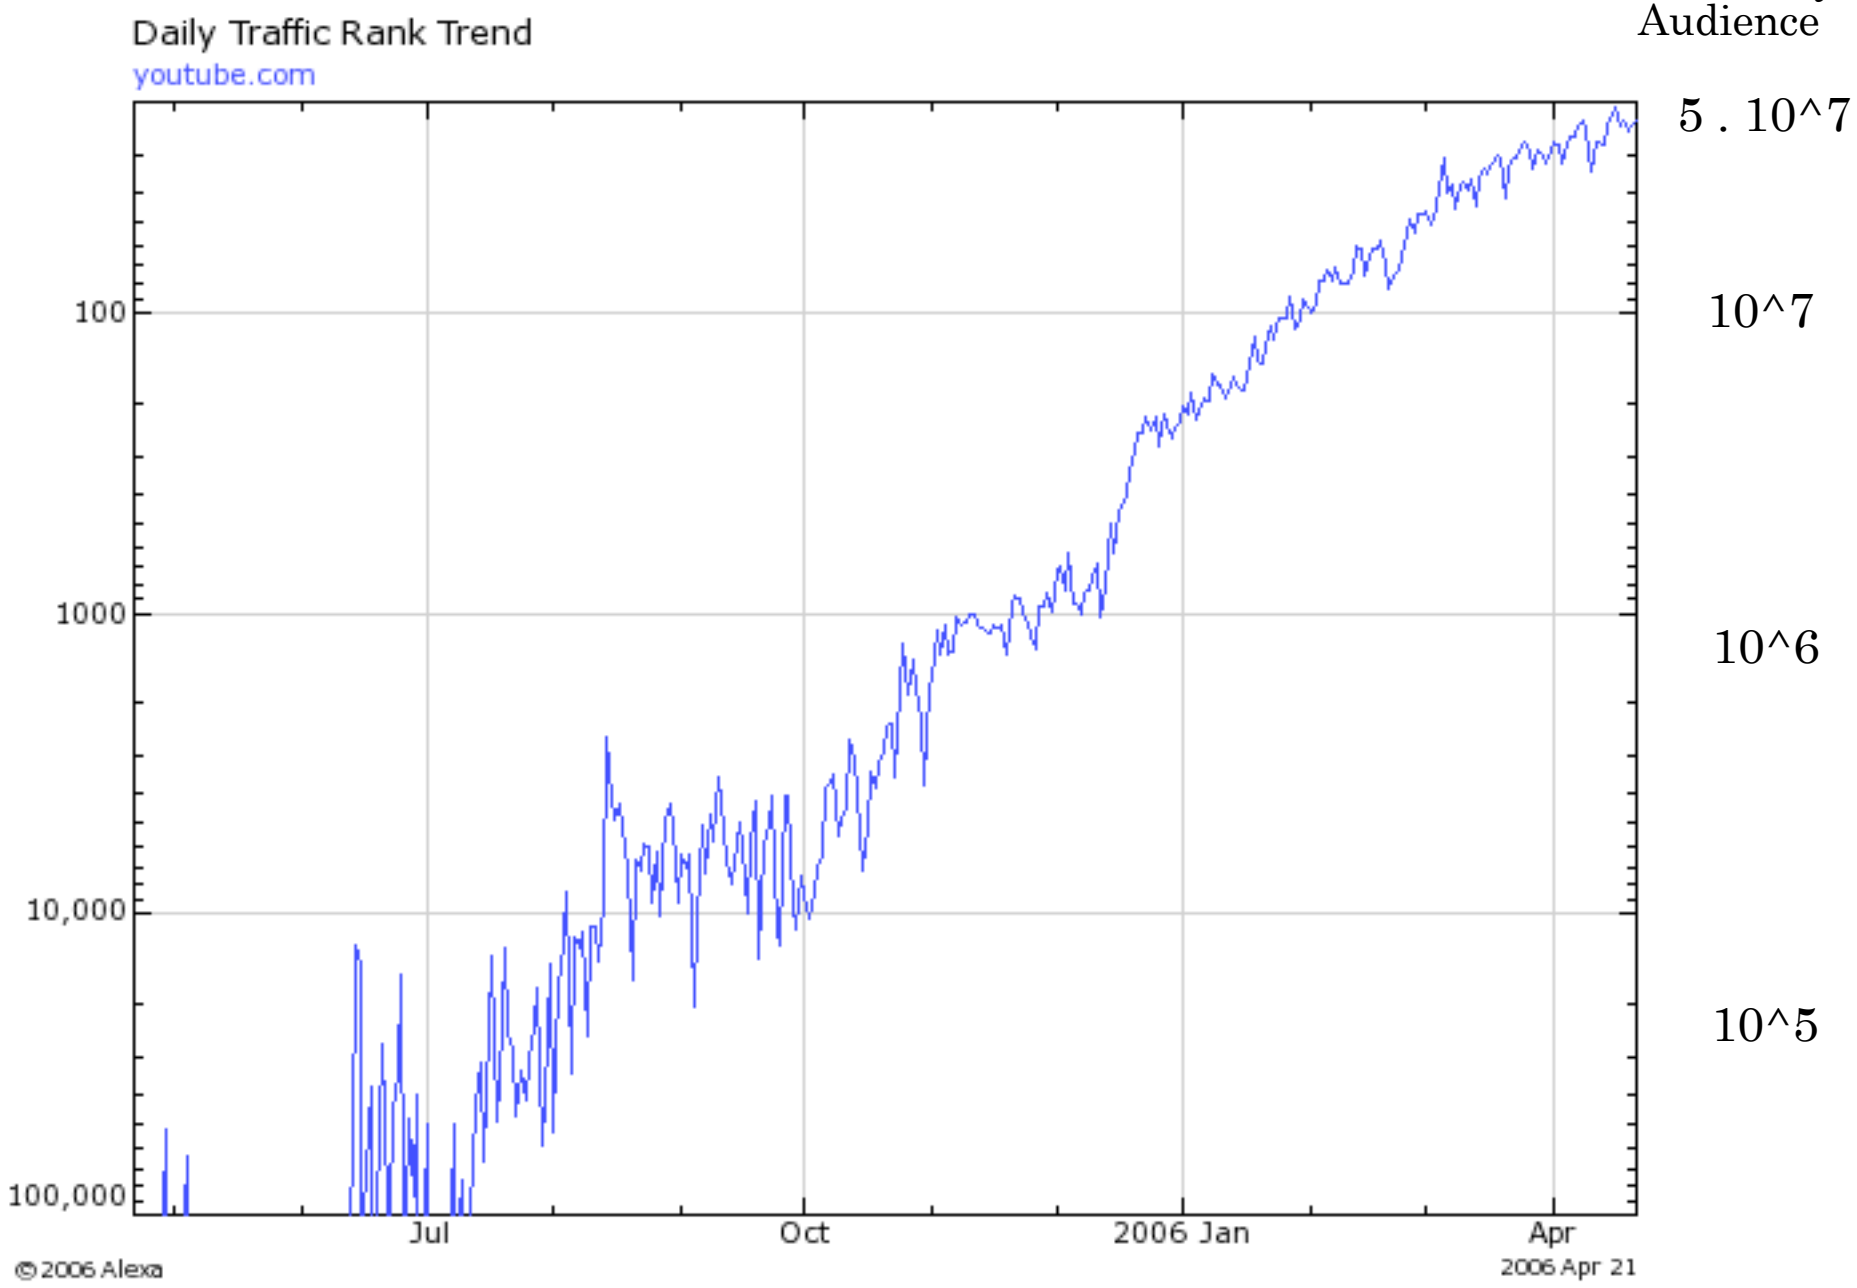

5 . 10^7

10^7

10^6

10^5

But how much video is there, really ?

- It is frequently hard to find definitive numbers for the past, and the future is always cloudy.
  - You have to use what you can get.
- Cisco put together a white paper, which seems fairly sound, and which is publicly available.
  - Global IP Traffic Forecast and Methodology, 2006–2011
- The Minnesota Internet Traffic Studies (MINTS) estimates both traffic for year end 2007 and growth rate per annum, with error estimates.
  - http://www.dtc.umn.edu/mints/home.html
- So let's look at these traffic projections, and then at a way of mathematically dealing with what's going on.

Overall traffic projections

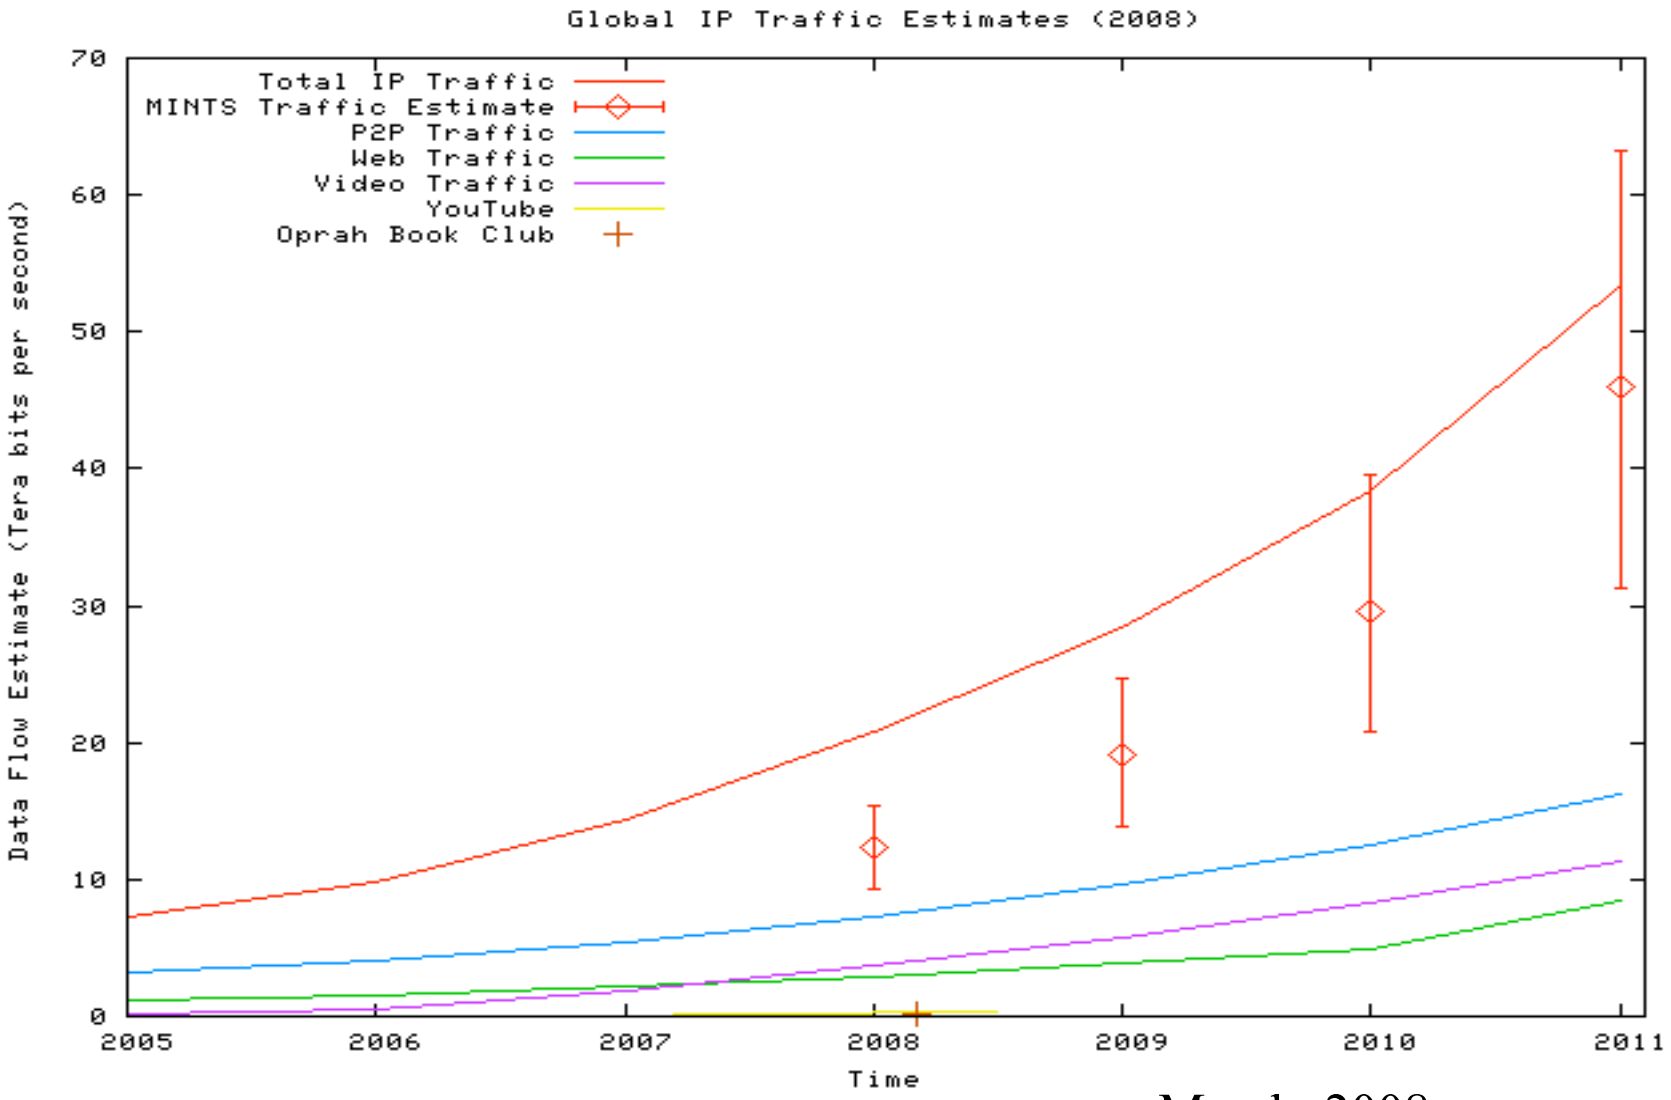

March, 2008

Relative traffic projections

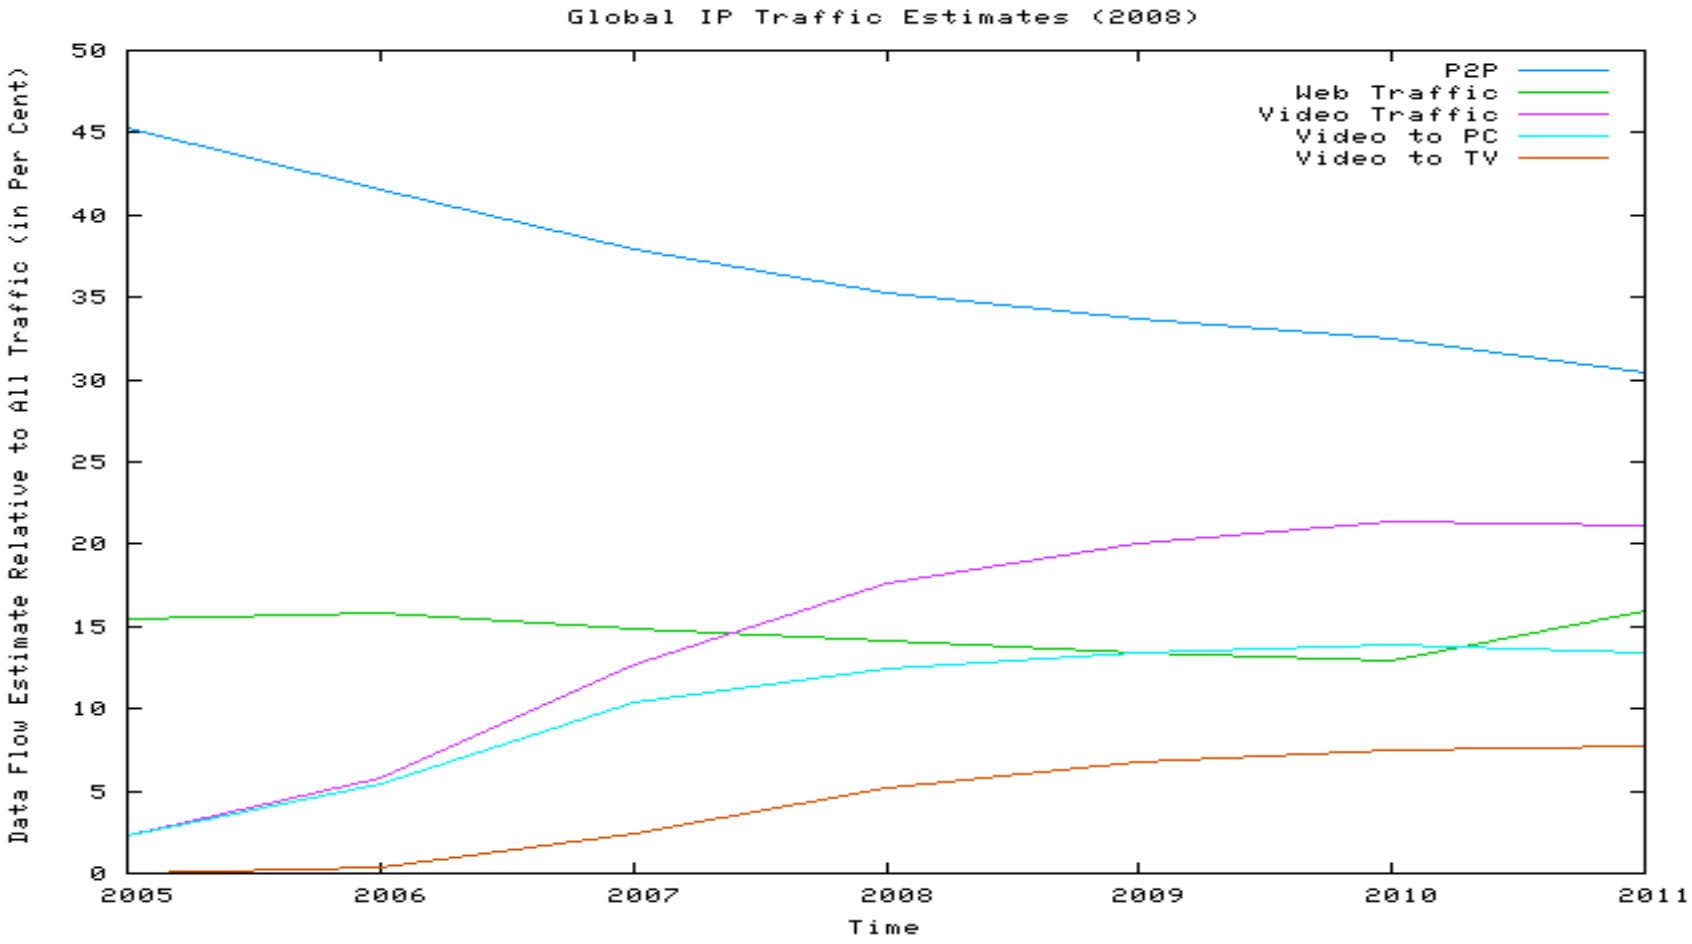

- From Cisco's ***Global IP Traffic Forecast and Methodology, 2006–2011***

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/net_implementation_white_paper0900aecd806a81aa.pdf

March, 2008

P2P Dominates
But what is P2P transporting ?

- P2P is a *transport* mechanism.
  - It is also a social movement, but that is artificial, and is due to various social and legal restrictions.
- In the past, P2P traffic was mostly audio
- Now it is reasonable to assume that it is largely video
  - Adding Peer to Peer and direct IP Video, we get ….

Relative Traffic Projections

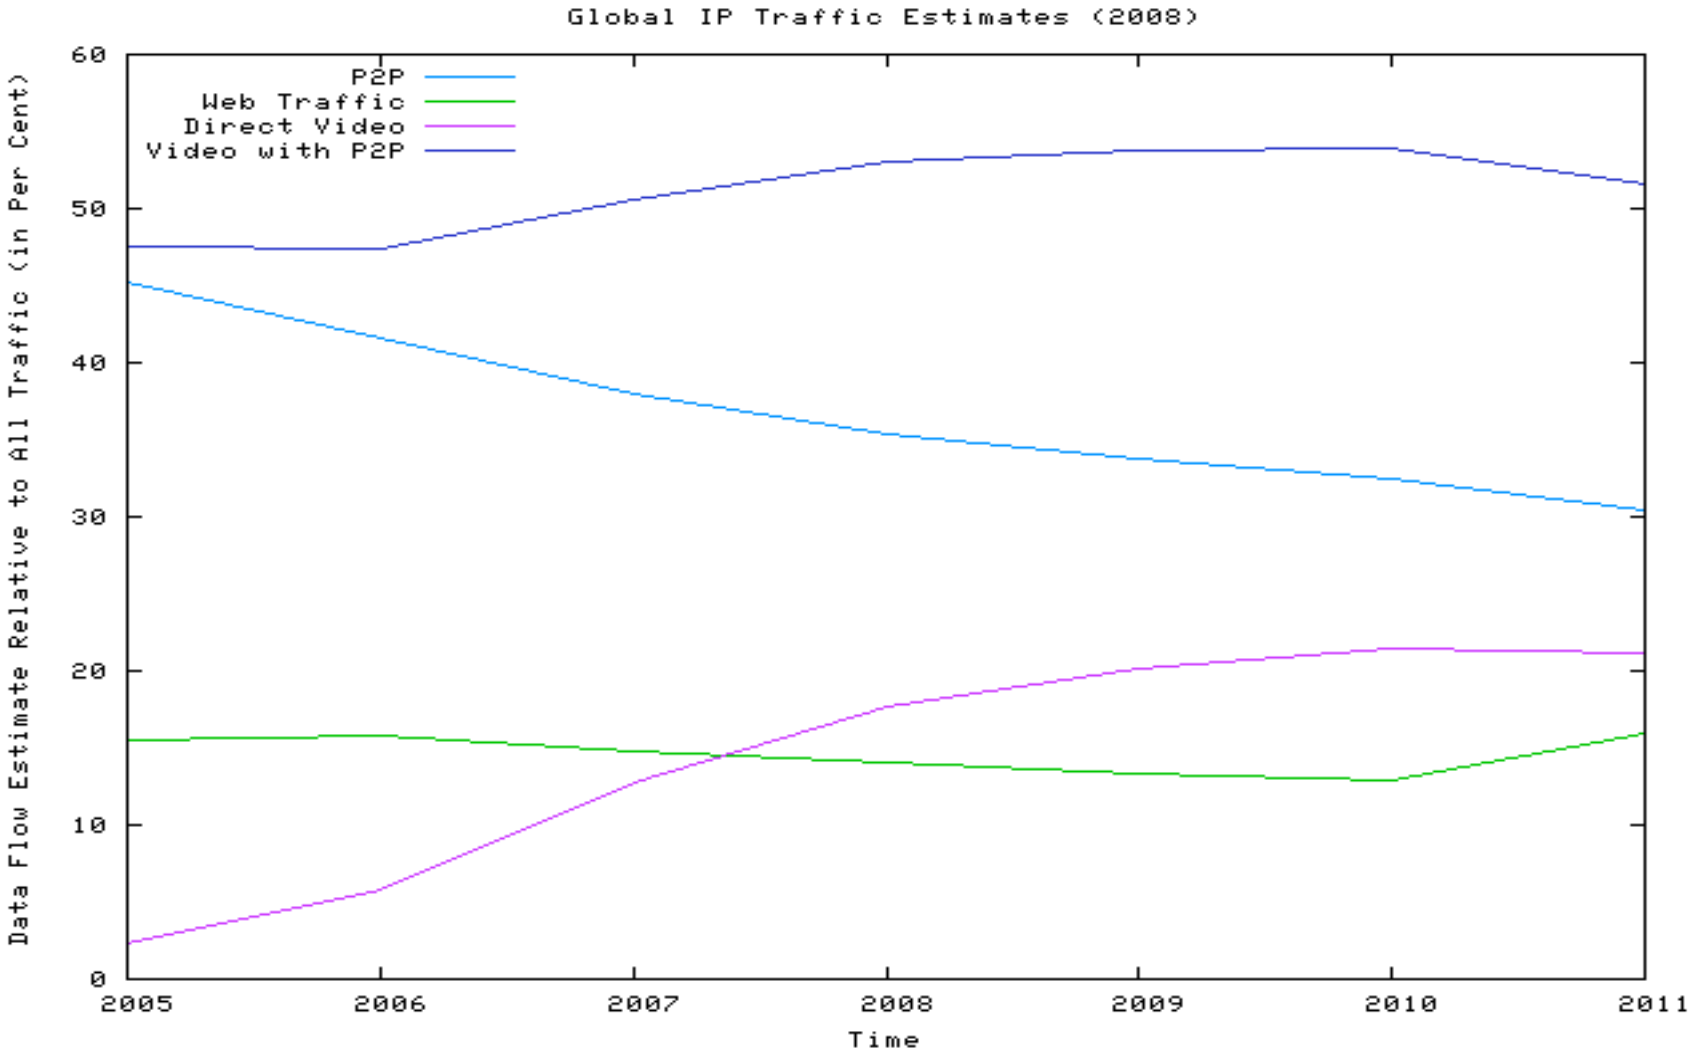

March, 2008

Traffic Projections (Conclusions)

- The Cisco White Paper agrees reasonably well with other estimates for overall traffic.
  - They also include VPN traffic, which muddles the video picture somewhat.
  - The MINTS 55 ± 5 % per annum growth estimate is larger than but roughly consistent with the Cisco White Paper.
  - It is reasonable to assume that video will continue to be ~ 50% of this traffic.
  - Will other transport mechanisms supplant P2P, as the White Paper indicates ?
    - Maybe.
  - I think that that is likely to some degree, as the Long Tail develops
    - So, let's look at the Long Tail

Zipf's Law and the Long Tail.

- Zipf's Law has been found in virtually every case of (not artificially limited) content selection.
  - Technically, this is a Pareto distribution.
- Zipf's law postulates a power law relation between the frequency of selection and the rank order of the option, expressed mathematically by Zipf's equation :

$$P = K \; R^{Z-1}$$

- where R is the rank order, Z is the Zipf exponent, K is a constant and P is the frequency of selection.
  - The Higher Z, the bigger the "Long Tail" is
- In video rentals, Chervenak found Z ~ 0.27.
- For Amazon.com Book sales, Brynjolfsson *et al.* found Z ~ 0.13
- For YouTube, Gill *et al*. found Z ~ 0.44
- For Web Site Usage, I find that Z ~ 0.15

March, 2008

The "80-20" rule

- A common, heuristic, version of Zipf's Law
  - 80 percent of the business comes from the top 20% of the content.
  - The observed Zipf's Law in Video Rentals, for example, implies that the top 20% of the titles generates 62% of the rentals.
    - Not bad for a heuristic.
- Zipf's law probably arises from similar distributions in social networks
  - I think that "6 degrees of separation" and the Pareto distributions have the same root cause.

Zipf's Law in Web Site Usage

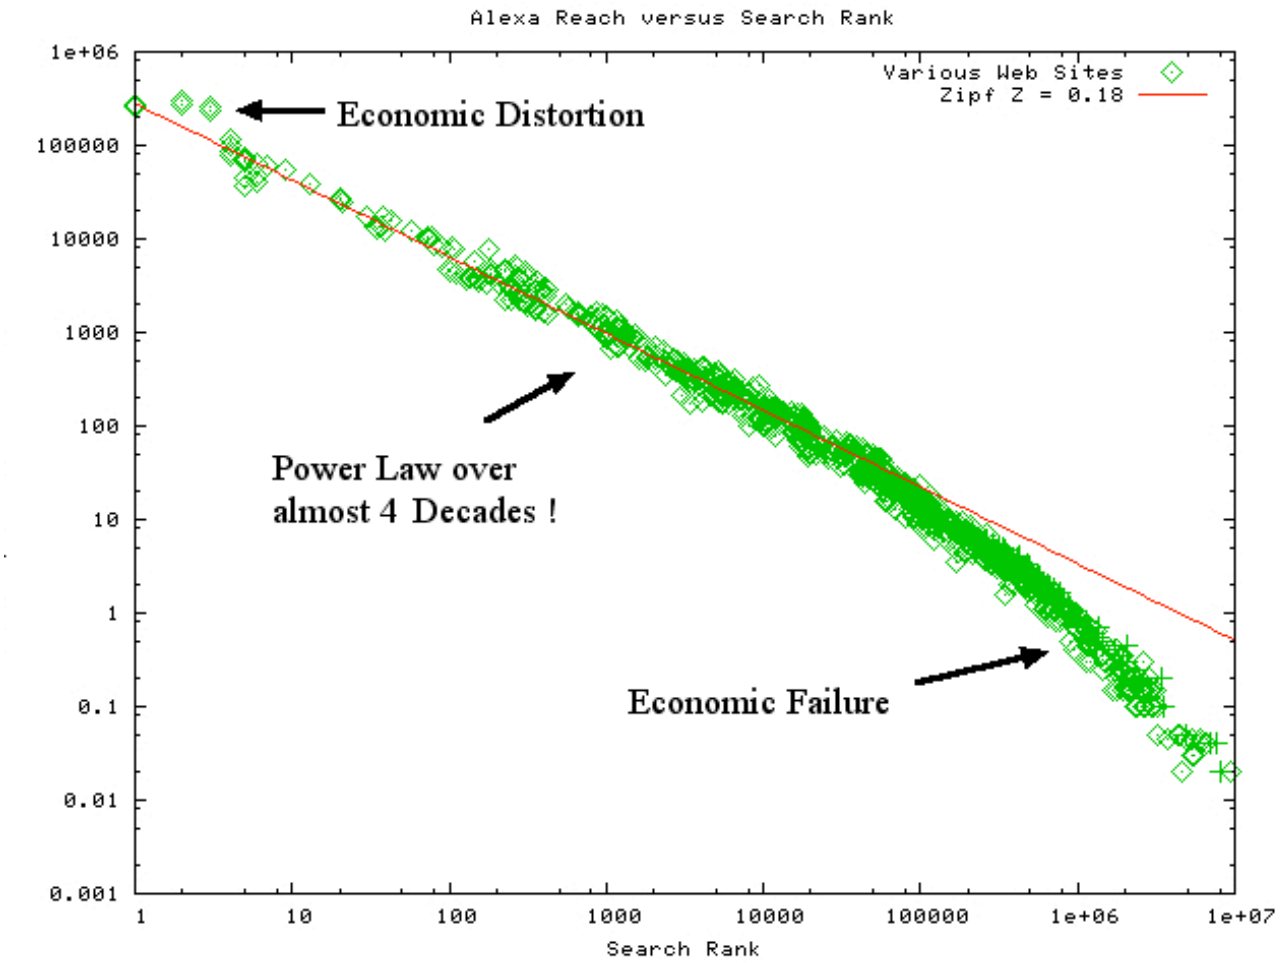

At some point, the audience becomes too small to make a profit.

Reach is the proportion of all Internet users who visit a given site, expressed per million users.

It is thus a measure of audience size.

Data from Alexa : http://www.alexa.com/data/details

March, 2008

Well, what can you do with this ?

- You can model the total video audience distribution.
- Suppose that video content (movies, shows, channels, etc.) has Z ~ 0.15, and the power law distribution holds for 100,000 video channels
  - 1 ≤ rank < 10 has a relative usage of 10.3 %
  - 10 ≤ rank < 100 has a relative usage of 12.6 %
  - 100 ≤ rank < 1000 has a relative usage of 17.5 %
  - 1000 ≤ rank < 10000 has a relative usage of 24.7 %
  - 10000 ≤ rank < 100000 has a relative usage of 34.9 %
- This is a fairly low value for Z, and thus for the importance of the Long Tail.
- The Long Tail thus cannot be ignored. For almost any reasonable value of Z, there will be a substantial audience in the aggregate of the “niche” content.

Zipf's Law and the Video Universe

- With our assumptions, the world can support a lot of video channels !
  - 100's of thousands of profitable channels does not seem outrageous world-wide.
  - The long tail in video content will be long indeed.
- Although the appropriate Z exponent for video channels is unknown, and although there may be other limits to the expansion of the video universe (e.g., cost and availability of content), it seems clear that the video universe will continue its rapid expansion.

Commercial Multicast Video Distribution

- Multicast is becoming the preferred means of distribution for video (TV) to Set Top Boxes (STB) over IP Networks (i.e. IPTV).
- Why ?

  It saves money.

  It uses the IP Infrastructure
- The Buzzword of the day is “Triple Play” - Data, VOIP, and Video on the same network
  - And this requires Multicast Video

Multicast Streaming : Walled Garden or Global Utility

- Most current plans for multicast streaming is entirely behind the scenes.
  - The "walled garden" approach.
    - Video packets and user packets never touch
- I don't think that Zipf's Law will allow the walls to stand.
  - As you will see, statistical models predict 10,000s to 100,000s of channels in the USA alone.
  - It's hard to see how the walled garden can be extended to encompass this
  - The model will switch from content selection to service provision, with protected "major" content and maybe best effort for everything else.
- What better way to get this additional video content except by IP multicast from the source.

Can Multicast relieve some of the video pressure on the network ?

- Maybe.
- If most IPTV providers convert to IP Multicast, then they could carry much of the Video Long Tail as best effort traffic, or best effort with FEC.
- There is no ***technical*** reason why this can't happen.
  - IPTV walled gardens have to open
  - There needs to be a standard for Electronic Program guides
  - Internet TV needs to adopt to Multicast.
    - The IETF is doing its part…

AMT

- Automatic Multicast without explicit Tunneling
  - draft-ietf-mboned-auto-multicast-08.txt
  - The idea is to provide a shim with automatic multicast failover
    - If you have native multicast, data arrives natively.
    - If you do not, a request is anycasted to a AMT relay, which unicasts the data to you (encapsulated)
    - No host / application modification is required.
    - Multicast is frequently available in the core, not at the edges. This protocol is intended to fix that.
  - Draft is nearing submission to the IESG.
  - If anyone is interested in supporting trials, please let me know.

Some Details of Current Internet Television

- For the time being, Internet Television will continue to be (mostly) unicast.
- I happen to run an Internet Television broadcaster, AmericaFree.TV
  - 20 channels, 19 in English, 1 in Spanish
    - Mostly long form content (movies)
  - Advertising Supported (no subscription)
  - 2007 Cumulative Audience of 4.2 million unique viewers
  - UDP and TCP streaming enabled for unicast
- Our audience mostly comes from “Push” not “Pull”
  - I think that will be typical in the Long Tail
- Tonight will be a stress test of the new configuration to support large numbers of joins in a short time.
  - Thanks for participating !

March, 2008

AmericaFree.TV Audience
2008-03-11 @ 1644 EDT

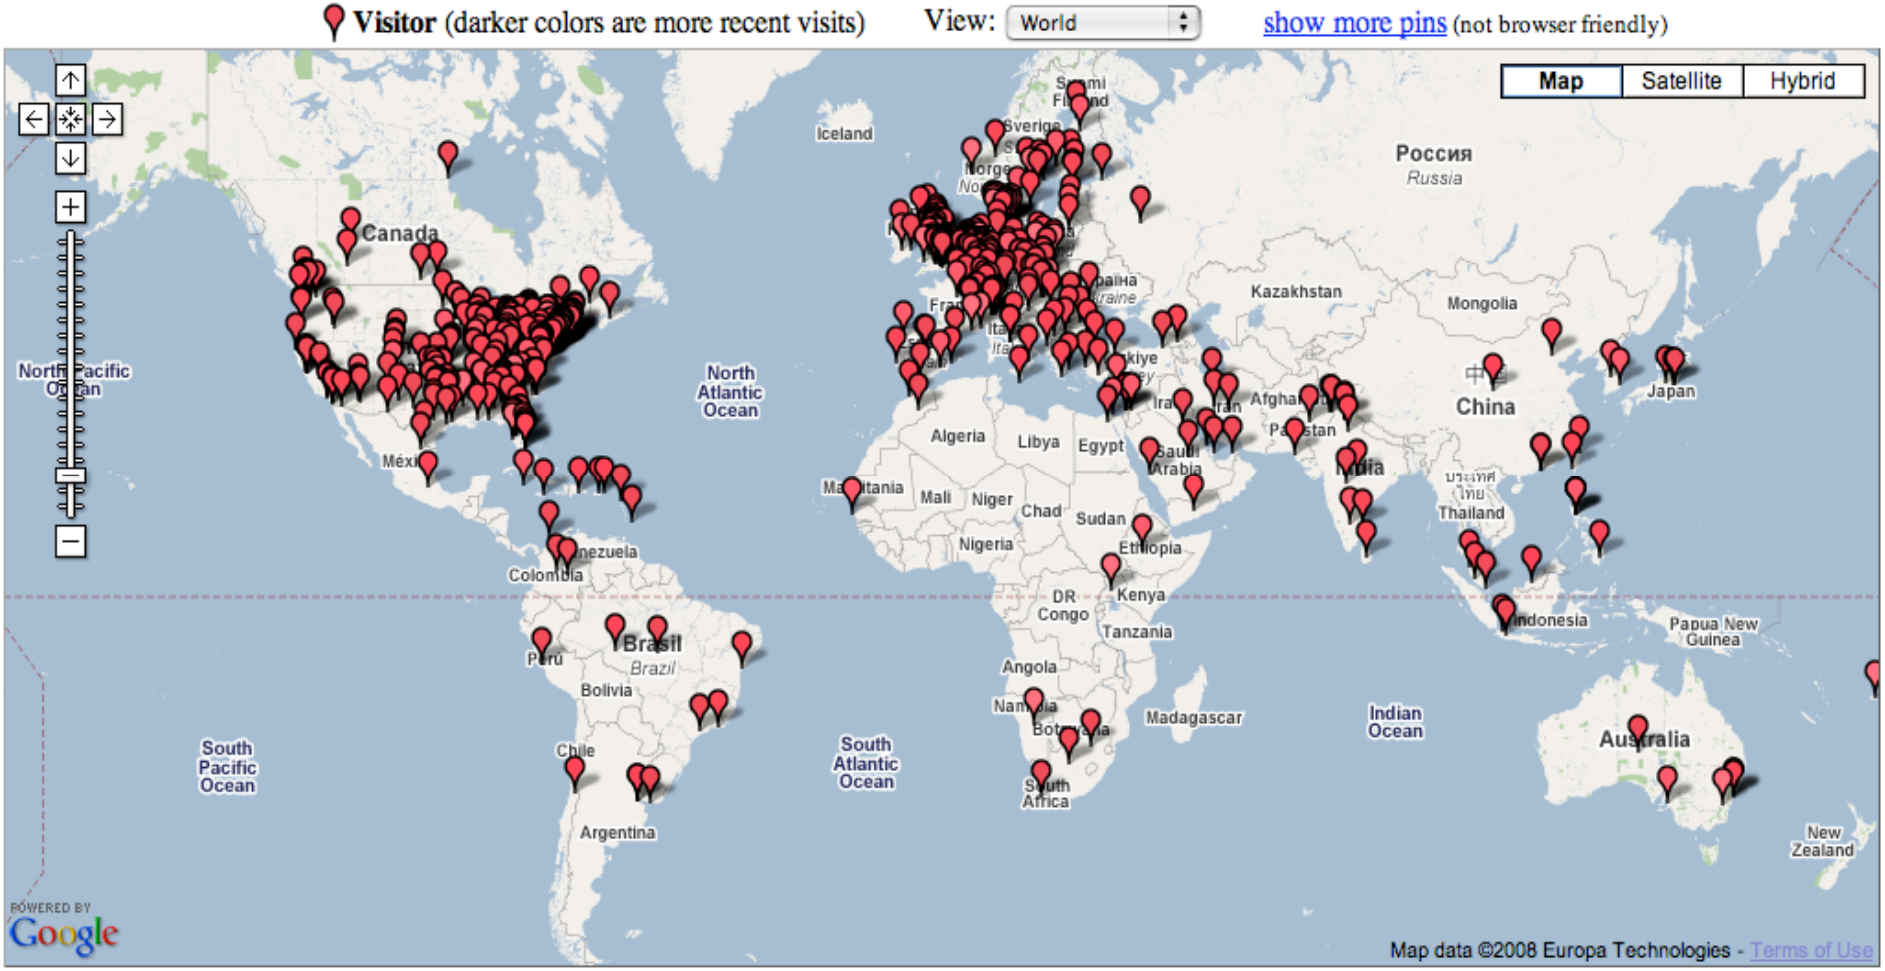

The first conclusion is that audiences are global

March, 2008

The Audience Likes Higher Bit Rates

- Channels are simulcast at a variety of bit rates, up to 2 Mbps, and including 3GPPx for cell phones.
- For the week February 23 - March 1, 2008 :

Sub-Channel	Fraction of Users	% of total Duration
1 Mbps (SD)	17 %	13 %
500 Kbps (SD)	54 %	44 %
250 Kbps (1/4 SD)	37 %	37 %
96 Kbps (1/8 SD)	1 %	3 %
Cell Phones (3GPPx)	3 %	3 %

- ~ 57% of the viewing time is spent watching our SD channels at 500 Kbps or higher.

(Note : Some users watch more than 1 sub-channel, so the user totals add up to > 100 %)

March, 2008

There is an audience for streaming HD Video

- For the week February 23 - March 1, 2008 :
  - For the HD content only (3GPPx is not offered for this content).

Sub-Channel	Fraction of Users	% of total Duration
2 Mbps (HD)	38 %	20 %
1 Mbps (SD)	42 %	34 %
500 Kbps (SD)	35 %	30 %
250 Kbps (1/4 SD)	21 %	14 %
96 Kbps (1/8 SD)	5 %	2 %

(Note : Some users watch more than 1 sub-channel, so the user totals add up to > 100 %)

March, 2008

Bandwidth Choice is surprisingly consistent

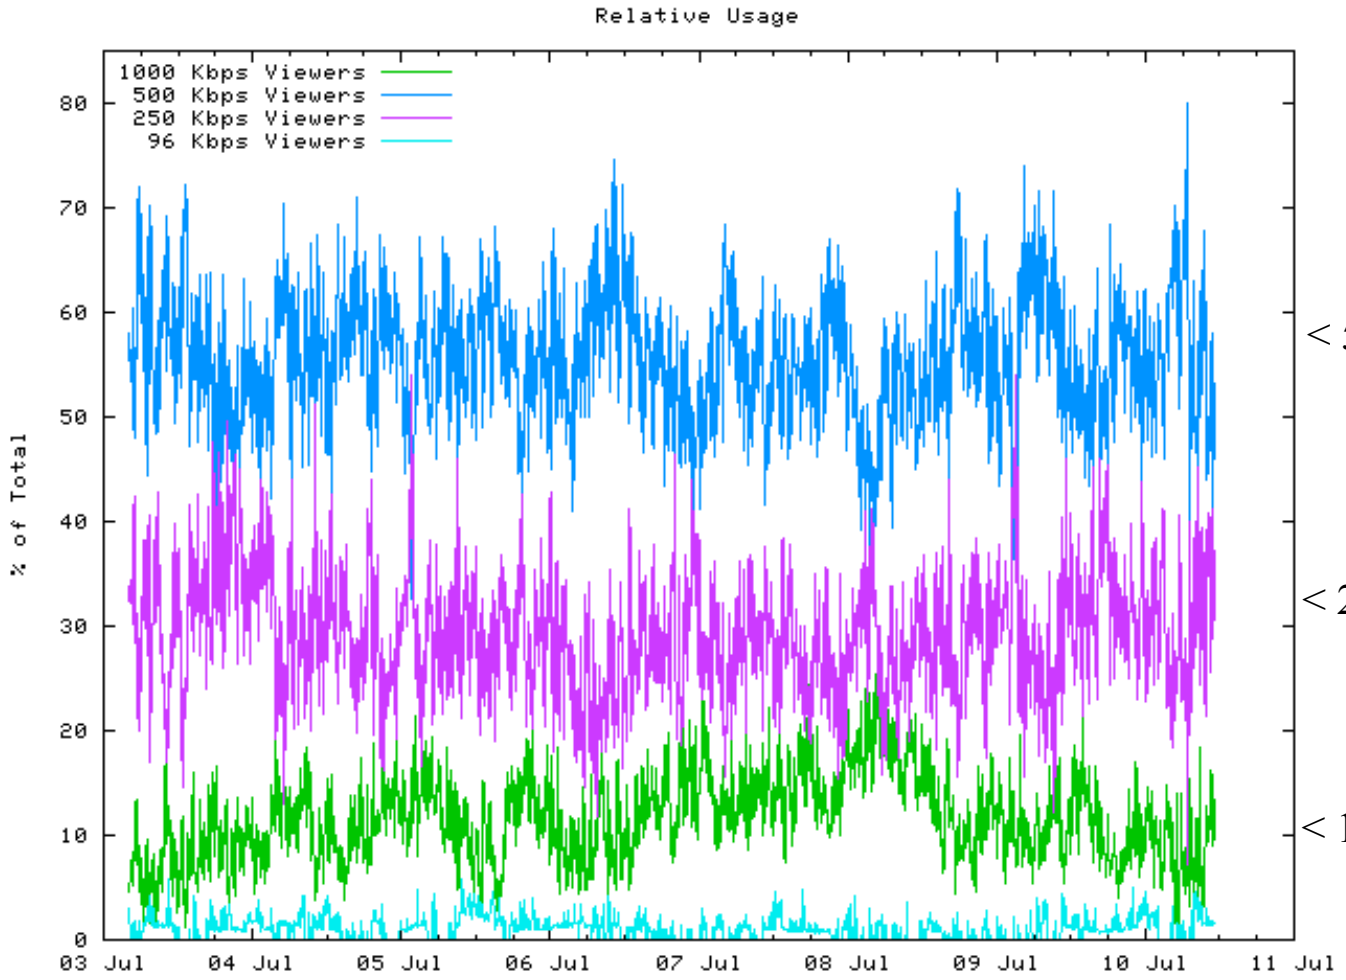

< 500 Kbps SD

< 250 Kbps Q SD

< 1 Mbps SD

March, 2008

Transport type

- Something like 5% of the audience is Multicast
  - Based on HTTP logs - direct multicast audience feedback would be useful.
- Over time, the amount of UDP traffic has been slowly decreasing.
  - For the week February 23 - March 1, 2008 :
    - 45.2 % of the sessions are UDP
    - 54.8 % of the sessions are TCP
- Why not use P2P for transport ?
  - It is not quite ready for streaming.
  - We make a profit from the existing transport mechanisms, so I don't feel a strong driver to change.

Audience Duration Histogram

- For 54 days in 2008, average viewing duration is 11.9 minutes. Some, however, view a lot more… (up to 24 days total in 54 days)

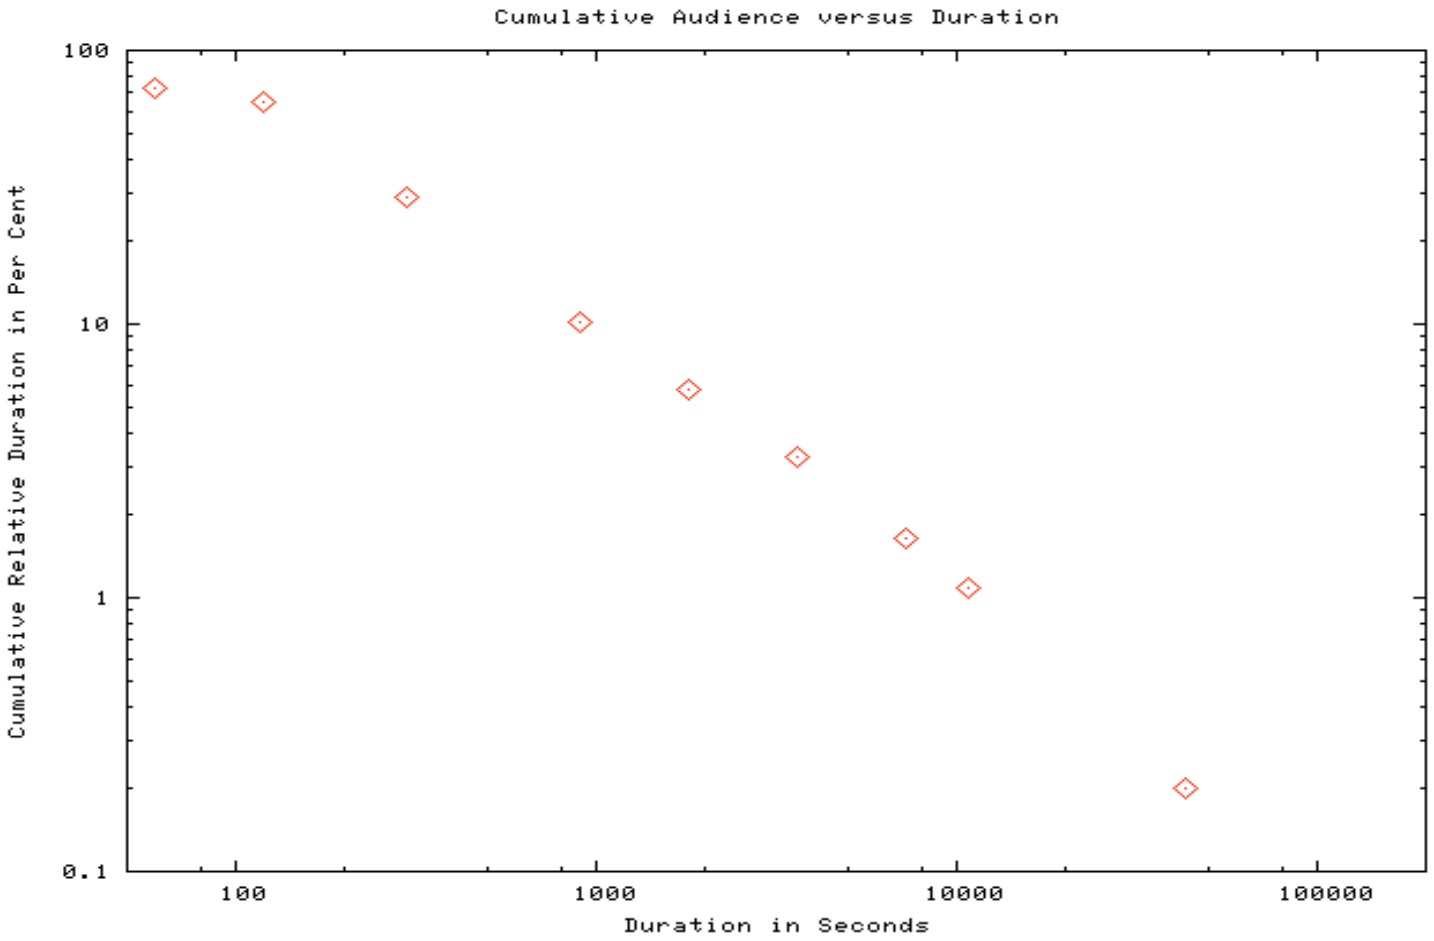

March, 2008

Implications

- Suppose in 2010 that 500 million people worldwide watch 12 minutes per day at 2 Mbps. (Or 100 million for 1 hour per day.)
  - That is ~ 8.3 Tbps

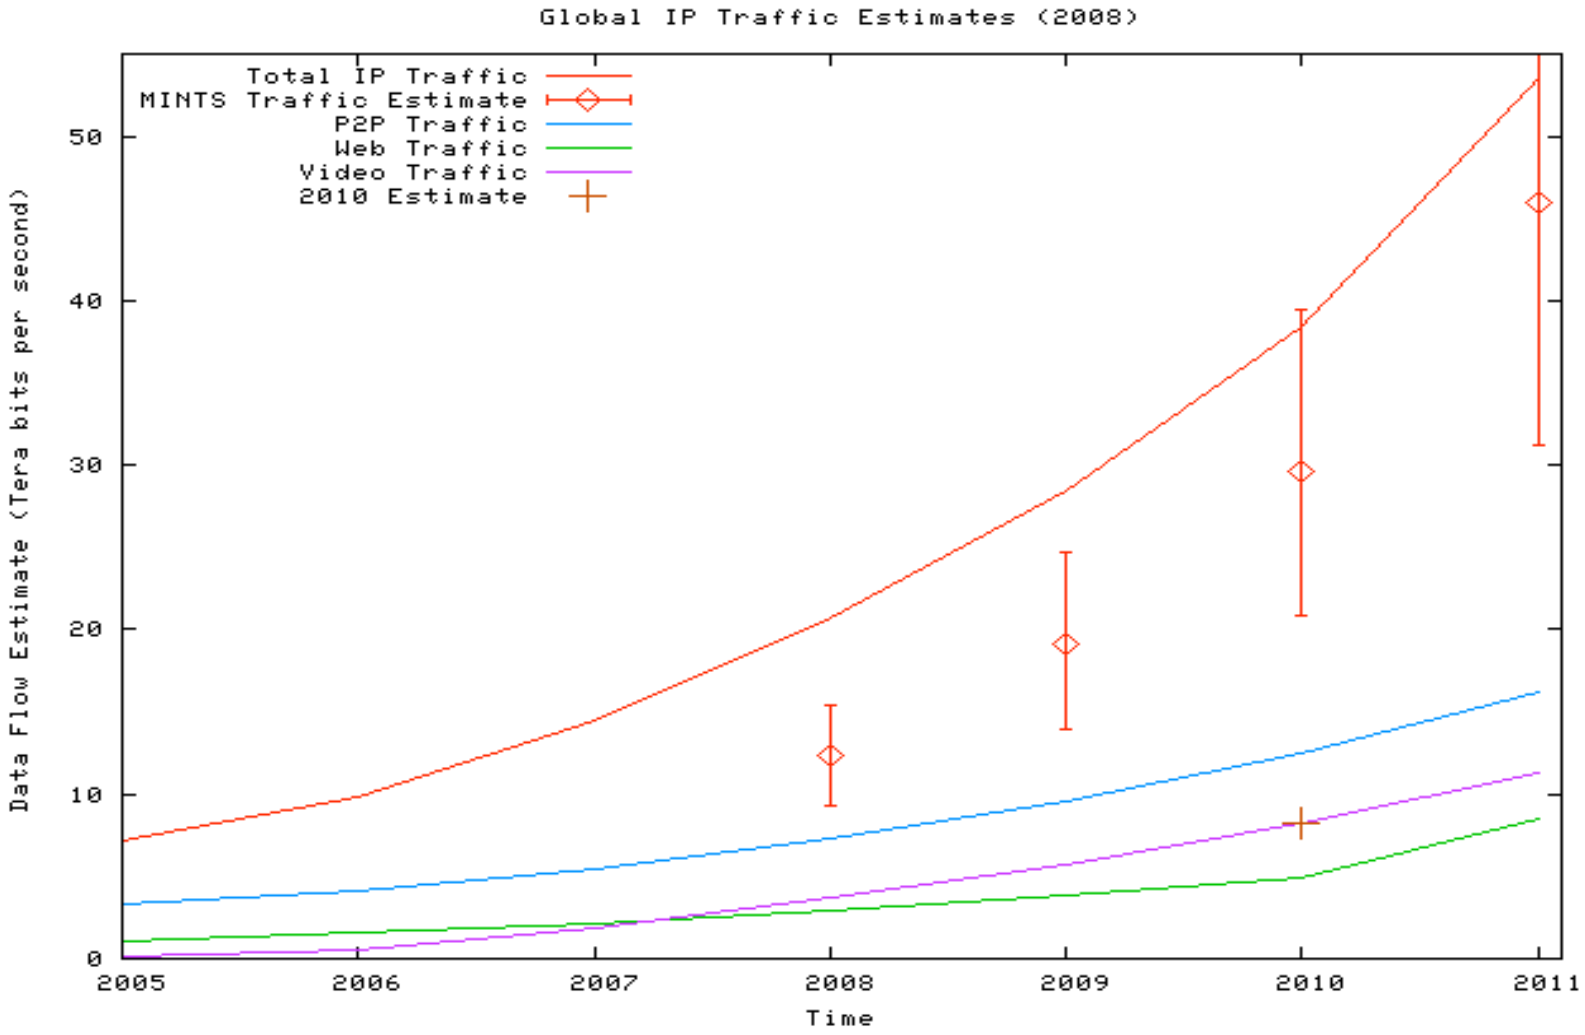

How can Best Effort Video Provide a Good User Experience ?

- I think that a major missing piece of the puzzle, and one that is being worked on in the IETF in RMT and FECFRAME, is in Forward Erasure Protection.
- I want to show why this work is so important.

MPEG-X & H.26X

- All of these standards have similar frameworks
  - The fundamental basis for compression is the macroblock (16 x 16 luma pixels or 8x8 chroma pixels), arranged into Slices, and then into frames.
  - All allow the use of previous (or future!) frames to predict the current frame (or macroblock)
  - Encoding is thus the compression of a prediction residual.
  - All allow for motion compensation to improve interframe prediction.
  - All use block based transforms and quantization to low pass filter the residual visual information

The Group of Pictures (GOP)

- There are three kinds of MPEG frames:
  - I (intra-coded)
  - P (predictive-coded
  - B (bidirectionally predictive-coded)
- There is one and only one I frame per GOP
  - It is encoded by itself, with no information from other frames
- P frames are encoded using the difference from the last I or P frame.
- B frames are based on the difference between the previous and next I or P frames.

Not all MPEG packets are created equal

- An I frame encoding is less efficient, so it might be 10 times as big as a P frame for the same video quality.
  - And the quality of the I frame determines the quality of the entire GOP.
  - Typically in HD an I frame is >> 1 packet. A P frame may not be. A B frame may be only one packet.
- Unless there is a repair mechanism, packet losses from I and P frames cause video errors that persist until the next I frame.
  - If any one of the I frame packets are lost, there will be errors persisting for the entire GOP.
  - If the GOP is (say) 20 video frames, 25% of the packets might be I frames for a 1 Mbps stream.
  - There would thus be a 25 % chance that a random single packet drop would corrupt 20 frames of data.
  - There is a 54% chance that a random packet drop would corrupt 10 frames of data.

NTT quality test for Raptor

- Two month trial: December 2003 - January 2004
- 300 subscribers in the Tokyo area
- FTTH (48%), ADSL (52%)
- Service: Video on Demand
  - Commercial service also uses multicast for scheduled/live delivery
- Format: MPEG2 (6 Mbps / 3 Mbps)
- 100 titles: movies, music, animation (30 - 120 minutes)
- Blind test
  - 50% Raptor
  - 50% No-FEC
- User access definition
  - Watched for at least 3 minutes

Courtesy Digital Fountain

March, 2008

NTT quality test for Raptor: Minutes per access

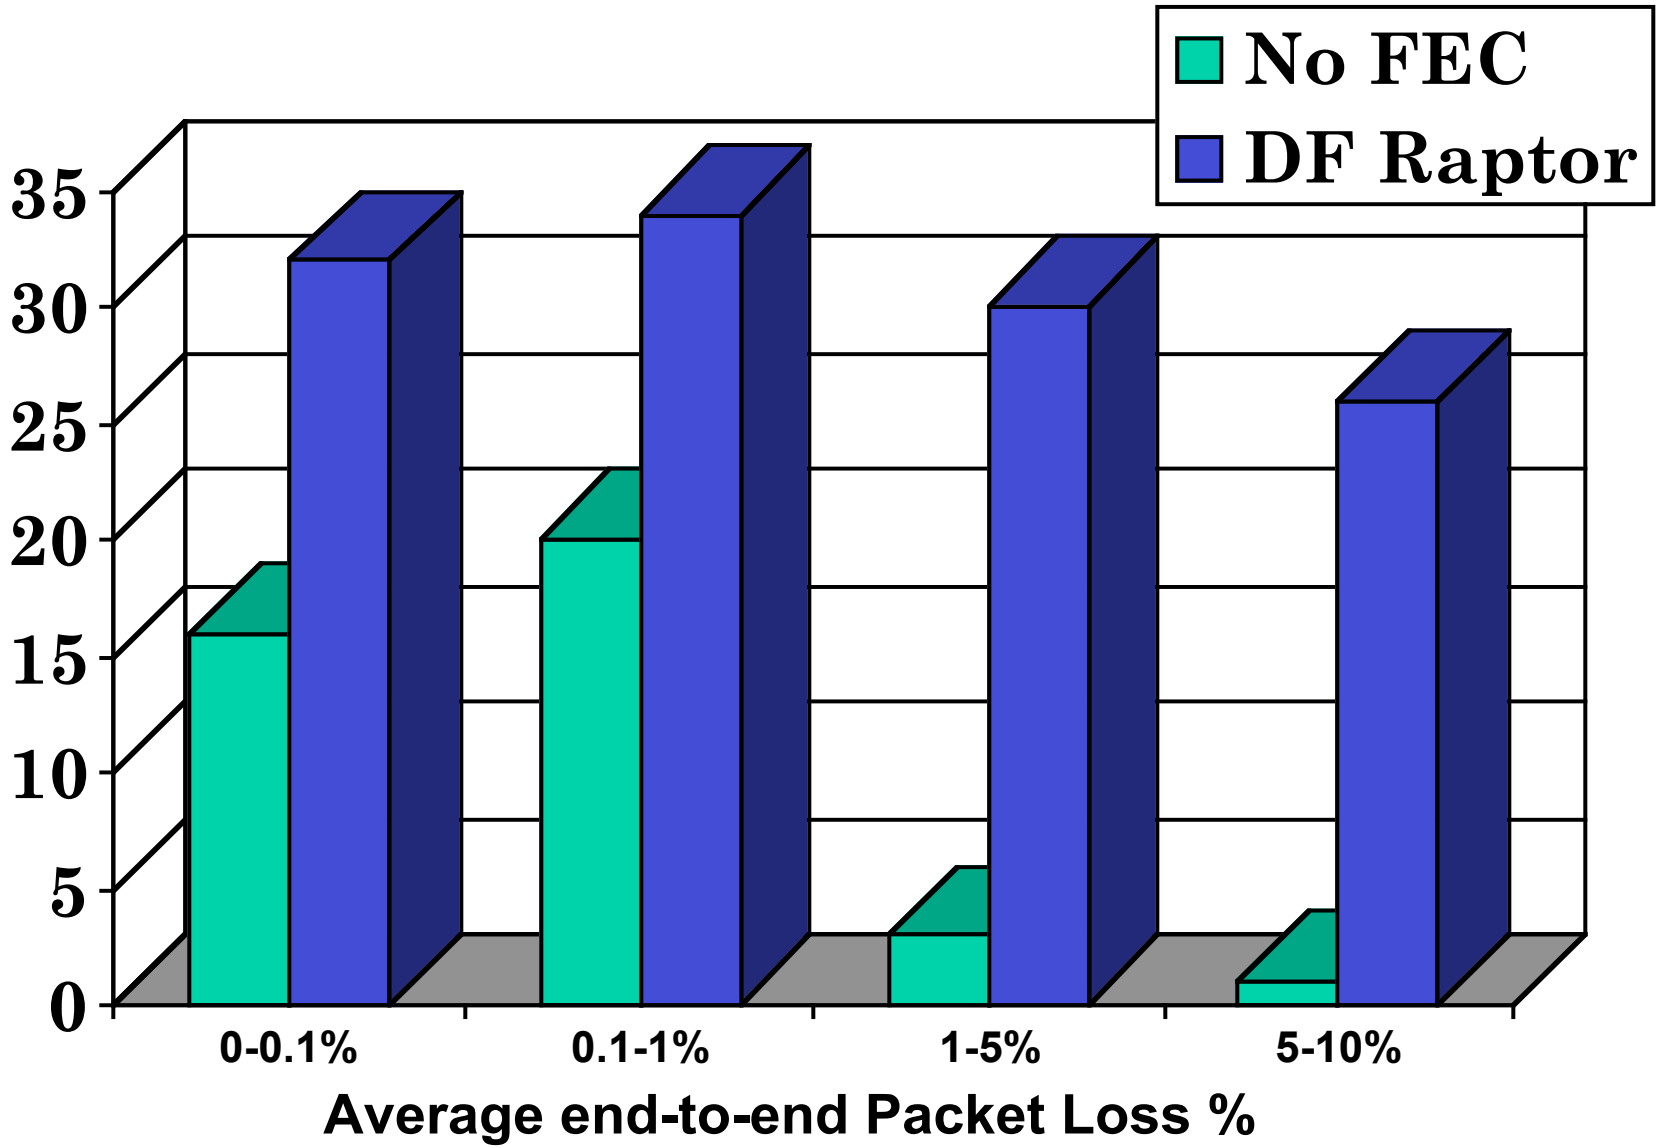

Source: NTT Trials. Blind test over Internet infrastructure.
User accesses of more than 3 minutes only

March, 2008

Courtesy Digital Fountain

Conclusions

- As far as we can tell, video is likely ***already*** ~ 50% of Internet usage, and that is not likely to decrease.
- Video usage is going to driven more and more down into the Long Tail.
- The existing business models for video are going to be under stress.
  - These are disruptive changes.
- Network models based on bursty web usage and large amounts of overprovisioning may be in trouble.
  - People watch for long periods of time.
- Overall, though, I don't see any reason why video will "break the net"
  - Growth will continue, but it doesn't seem to be disproportionate.

March, 2008