2013 17th International Conference on Intelligence in Next Generation Networks
icin 2013
Unlocking Value from the Networks
Smarter Networks for Smarter Things 15-16 October 2013, Venice, Italy

S6: New Architectures for Service Delivery Moving to Content-Centric Networks

Barry Crabtree, Steve Appleby, Mike Nilsson,
Tim Stevens, Brahim Allan

British Telecommunications plc

Overview

- TV viewing behaviours
- The impact on network design
- Edge caching with pre-emptive filling
- Content Centric Networks
  - A new way to deliver content with caching

The evolution of TV

- All early TV transmissions were live
  - There was no capability to store programming for later transmission
- Over time, more and more TV content was filmed in advance and stored for later transmission
- Now most broadcast TV content is not live, and there are many ways to access TV content at times convenient to viewers
  - Personal video recorders
  - Online streaming and download services to various personal and portable devices
- *But most TV content is still viewed at the time of broadcast !!*

Images from Wikipedia and http://advanced-television.com.

Live TV is a popular as ever

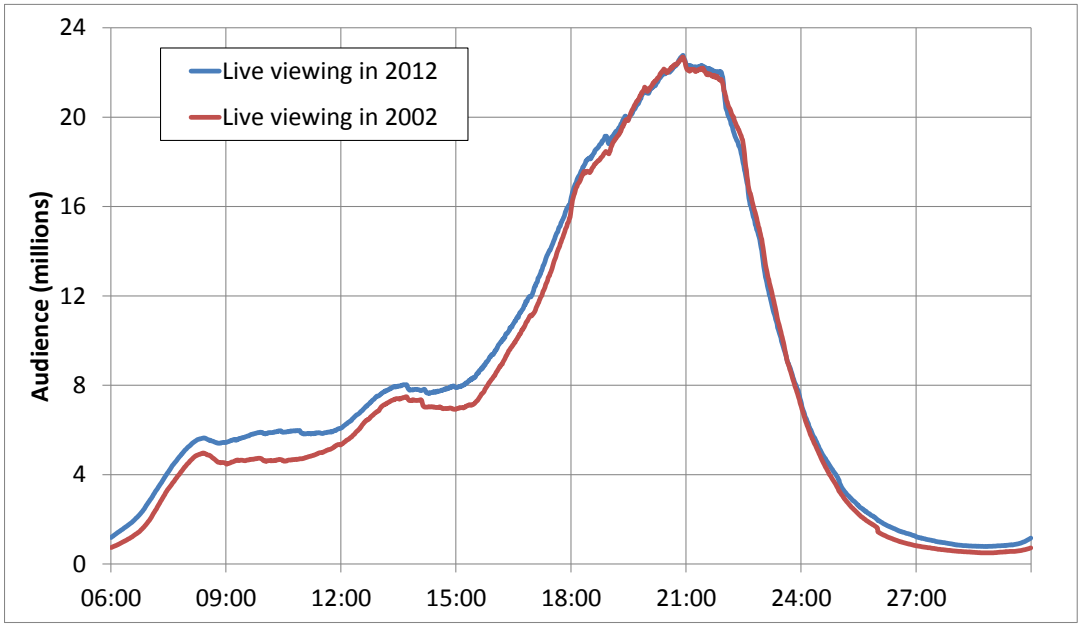

In the UK, in both 2002 and 2012, about 22 million people, 40% of the population, watched live TV every evening between 8PM and 10PM.

UK total live TV viewing has increased from 210 minutes per person in 2002 to 216 minutes per person in 2012 , primarily due to an increase in audience at non-prime-time.

- Live TV viewing levels maintained despite more personal video recorders and online video
- Instead, more viewing options has increased total consumption
- High levels of live viewing expected to continue over the next few years, driven by:
  - Sports and Interactive programming
  - Social media - at peak times, 40% of UK tweets are about live TV programmes
  - Inertia of current viewing behaviour

Data generated by accessing the BARB database, using Kantar media.

Viewer distribution across TV channels

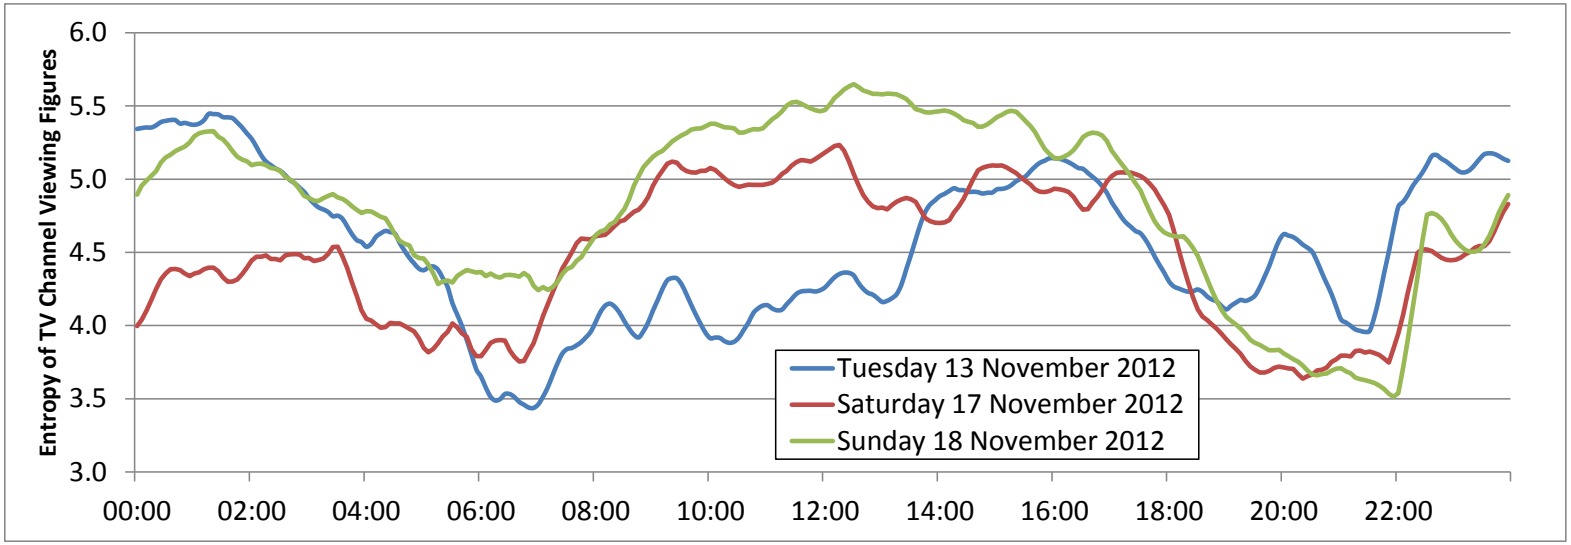

where p_c = proportion of viewers on channel c

- The number of people watching each of 114 channels is recorded every 5 minutes
- The entropy of the channel choice is calculated
  - Lower entropy indicates viewer concentration on a smaller number of channels
- Entropy varies with time of day and day of week:
  - Low entropy at breakfast and peak time
  - Higher entropy at night, in the afternoon and on weekend mornings

Data generated by accessing the BARB database, using Kantar media.

Viewer distribution across TV channels

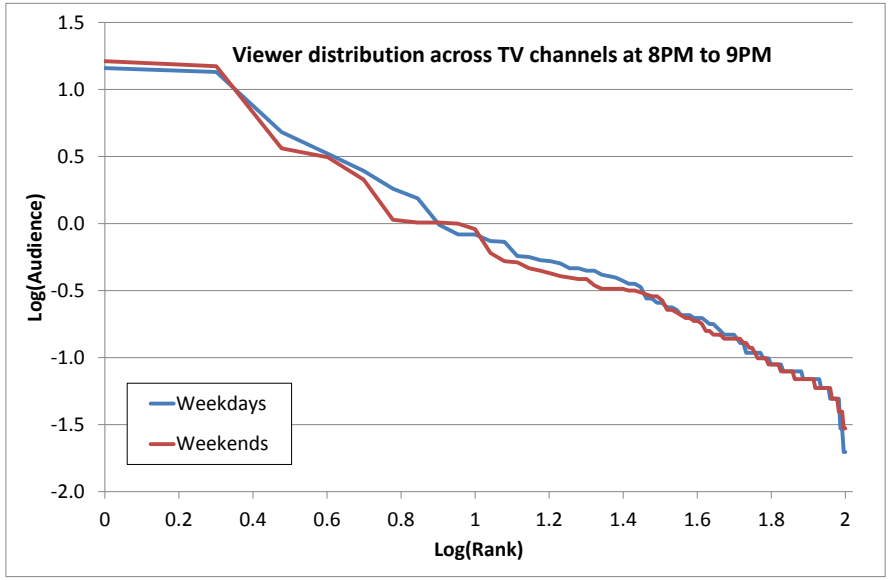

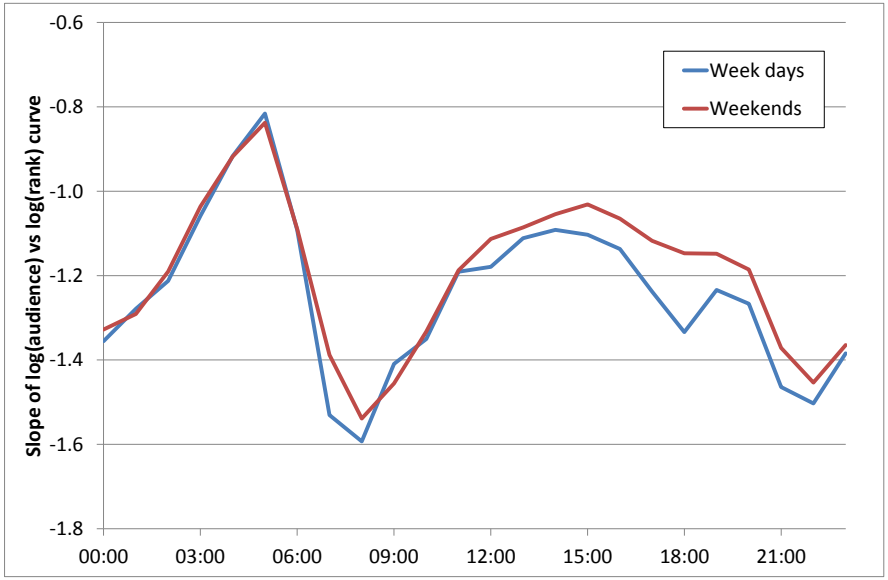

- The number of people watching each of 114 channels is recorded every hour over all of 2012
- The channels are sorted for each hour by decreasing popularity
- The logarithm of the audience against the logarithm of the rank (position) approximates a straight line
  - This behaviour is known as Zipf's law
- The slope of this line varies with time of day, and to a lesser extent, day of week:
  - The slope is steeper at breakfast and peak time
  - A steeper slope indicates viewer concentration on a smaller number of channels

Data generated by accessing the BARB database, using Kantar media.

Varying popularity of programs

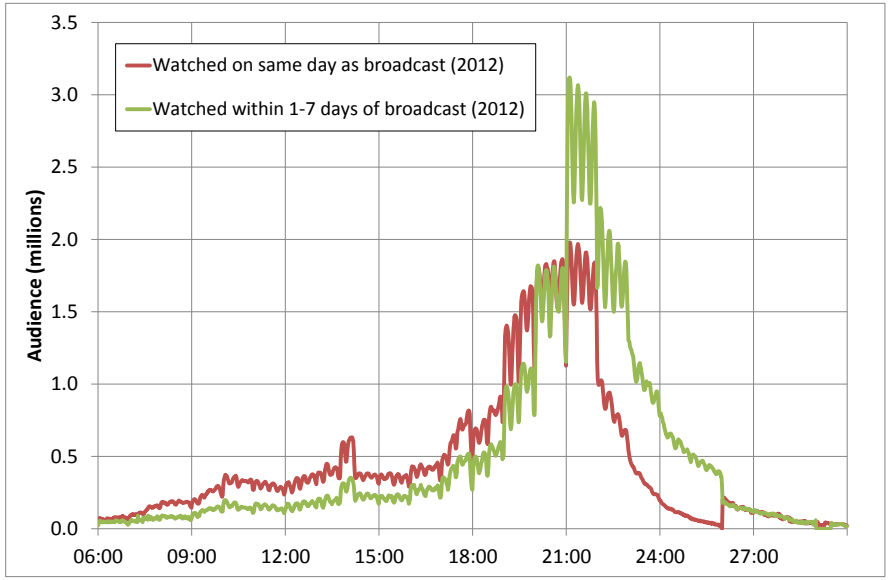

Average daily PVR recordings in the UK in 2012.

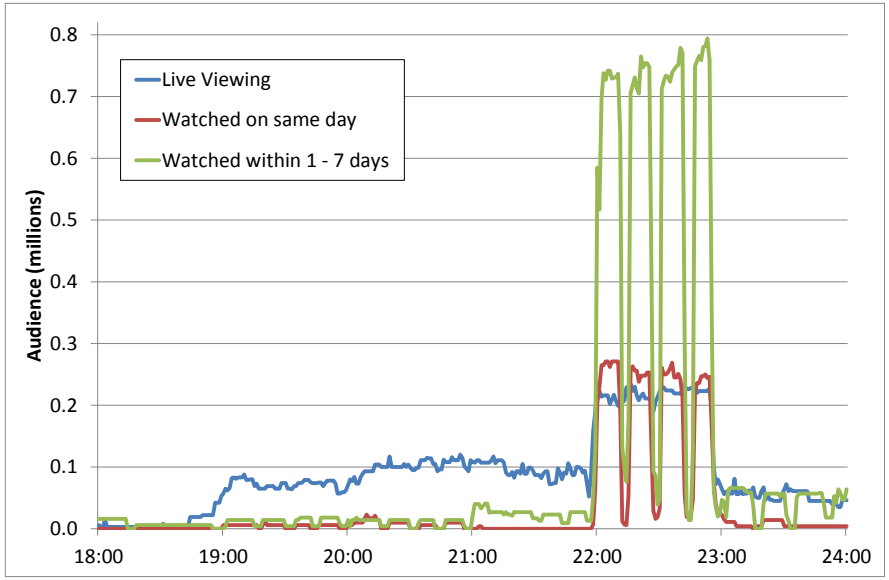

Viewing figures for Fox, 9 November 2012.

- All programs are not equally popular
- This is not just related to time of viewing
- Peak time programs are the most recorded
- On some channels, the audience increases significantly for specific programs

Data generated by accessing the BARB database, using Kantar media.

Varying popularity of TV content

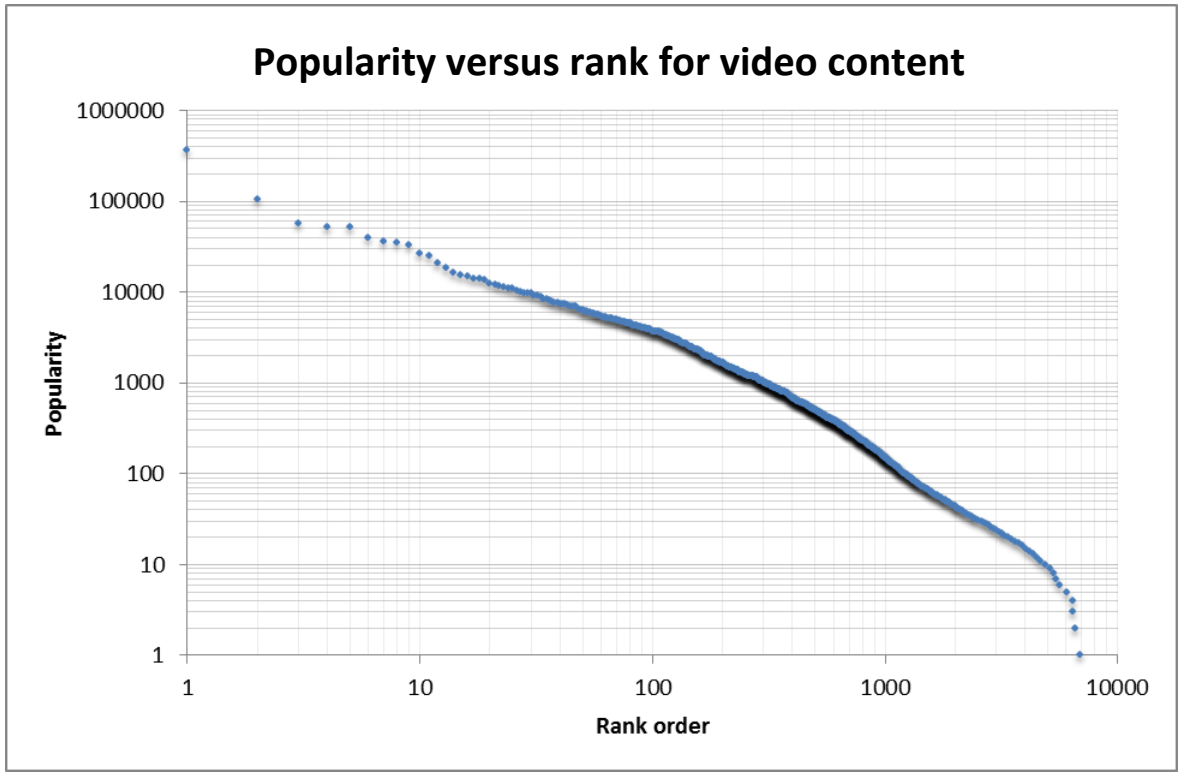

George Kingsley Zipf
1902–1950
Source: Wikipedia

- All TV content assets are not equally popular
- The popularity follows Zipf's Law where the logarithm of popularity is proportional to the logarithm of rank (the position in a list of content sorted by popularity
- In this example, the most popular asset is viewed 10 times as often as the asset that is 10th in the list of popularity, and is viewed 100 times as often as the asset that is 100th in the list of popularity.

Internet traffic levels by time of day

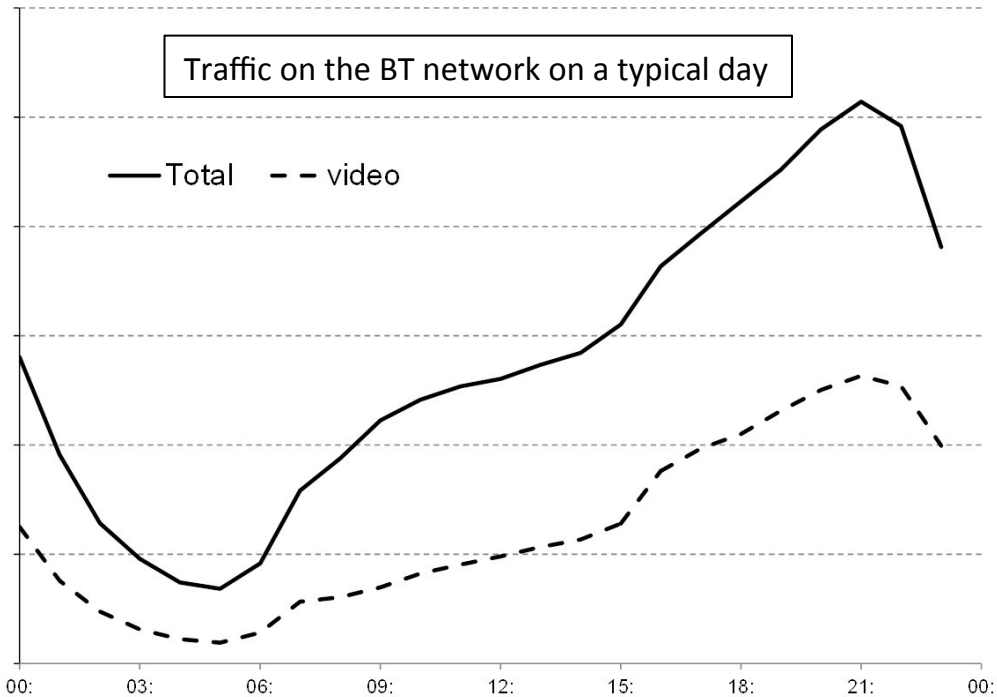

Over 50% of the total internet traffic is video traffic

This is consistent throughout the day.

The average traffic throughout the day is about 50% of the peak

The peak is predicted to increase quicker than the mean: by 2017, peak time traffic is predicted to increase four times and the average traffic three times

- Service providers need to design for the peak level of traffic
- Caching at the edge of the network, near end users, reduces core network bit rate demand
- Delivering content off peak could reduce core network bit rate demand further

Could content be delivered to edge caches pre-emptively?

Separating content delivery from content rights

- Traditionally content delivery and content rights were combined
  - Broadcast TV using conditional access and smart cards
  - Physical Media, e.g. DVD, with content protection
- Modern DRM systems allow content delivery to be separated from content rights
  - Encrypted content can be freely distributed, and "side-loaded" onto devices
  - Content is associated with rights that limit play back:
    - to specific people or specific devices
    - to specific time periods
  - Content could be distributed before play back is allowed
- The Digital Entertainment Content Ecosystem has created a digital rights locker ecosystem named UltraVioletthat operates on these principles

Can live-broadcast content be cached pre-emptively?

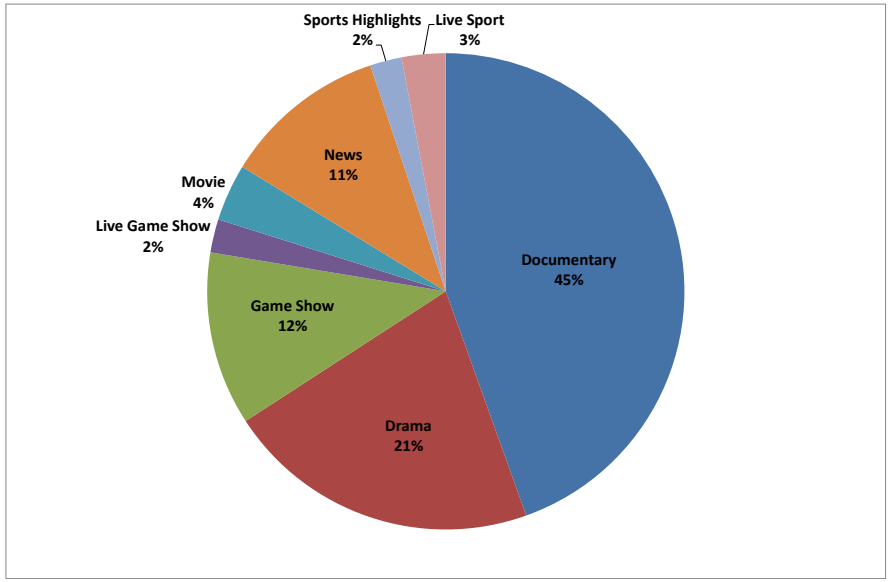

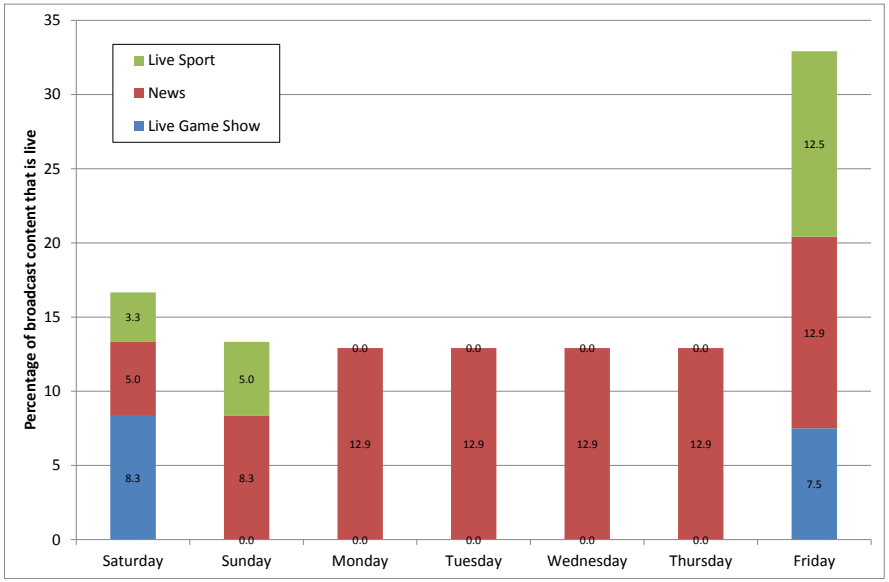

- Analysis of TV content on the five most popular channels in the UK during the week 20-26 July 2013, during the peak viewing period of 7PM to 11PM
- 16% is broadcast live, consisting of 3% Sport, 2% Live Game Shows and 11% News

A large proportion of broadcast content is not actually live but filmed much earlier
Content could be pre-emptively cached ahead of 'broadcast' or 'availability' time
But there are content rights and content availability issues

Can live-broadcast content be cached pre-emptively?

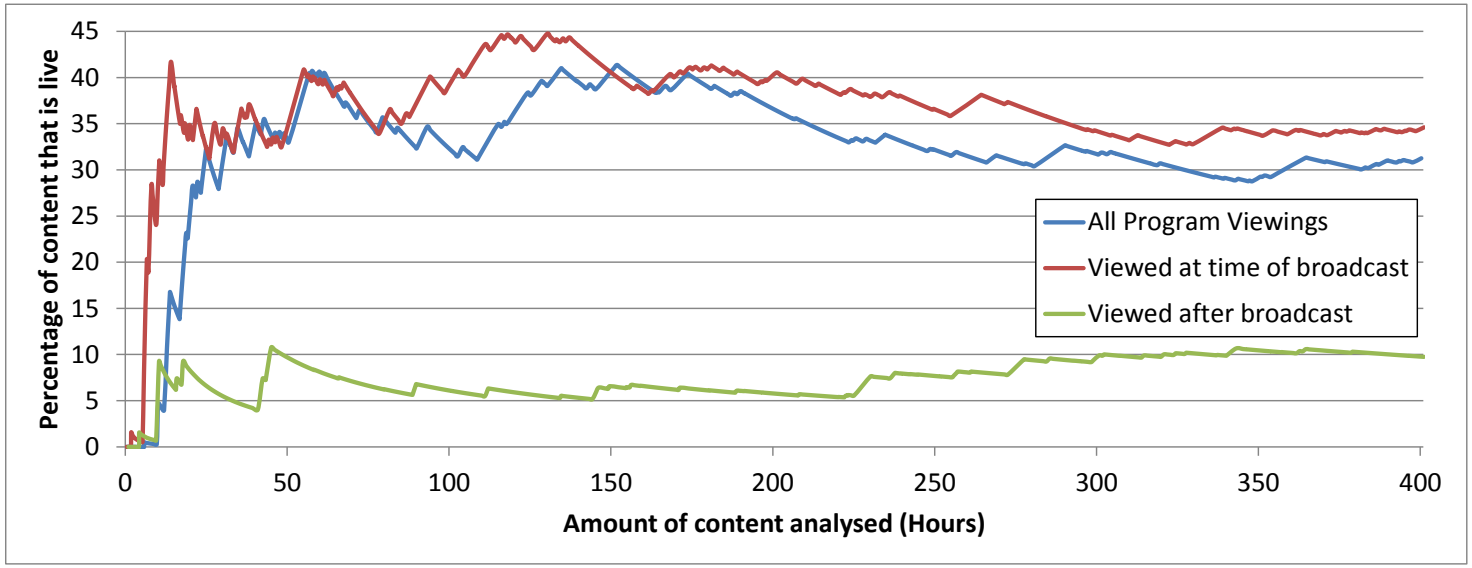

- Analysis of over 500 of the most viewed TV programs in the UK during the week 20-26 July 2013
- This represents 400 hours of program content, and over 50% of the total viewer minutes
- About 35-40% of content that is viewed at the time of broadcast is actually live content
- But only about 5-10% of content viewed later, using a PVR, was originally broadcast live

There is an opportunity to pre-emptively cache content to satisfy later demands for online 'catch-up', thus moving traffic from peak time to off peak
But there are content rights and content availability issues

Data generated by accessing the BARB database, using Kantar media.

Which content assets should be cached?

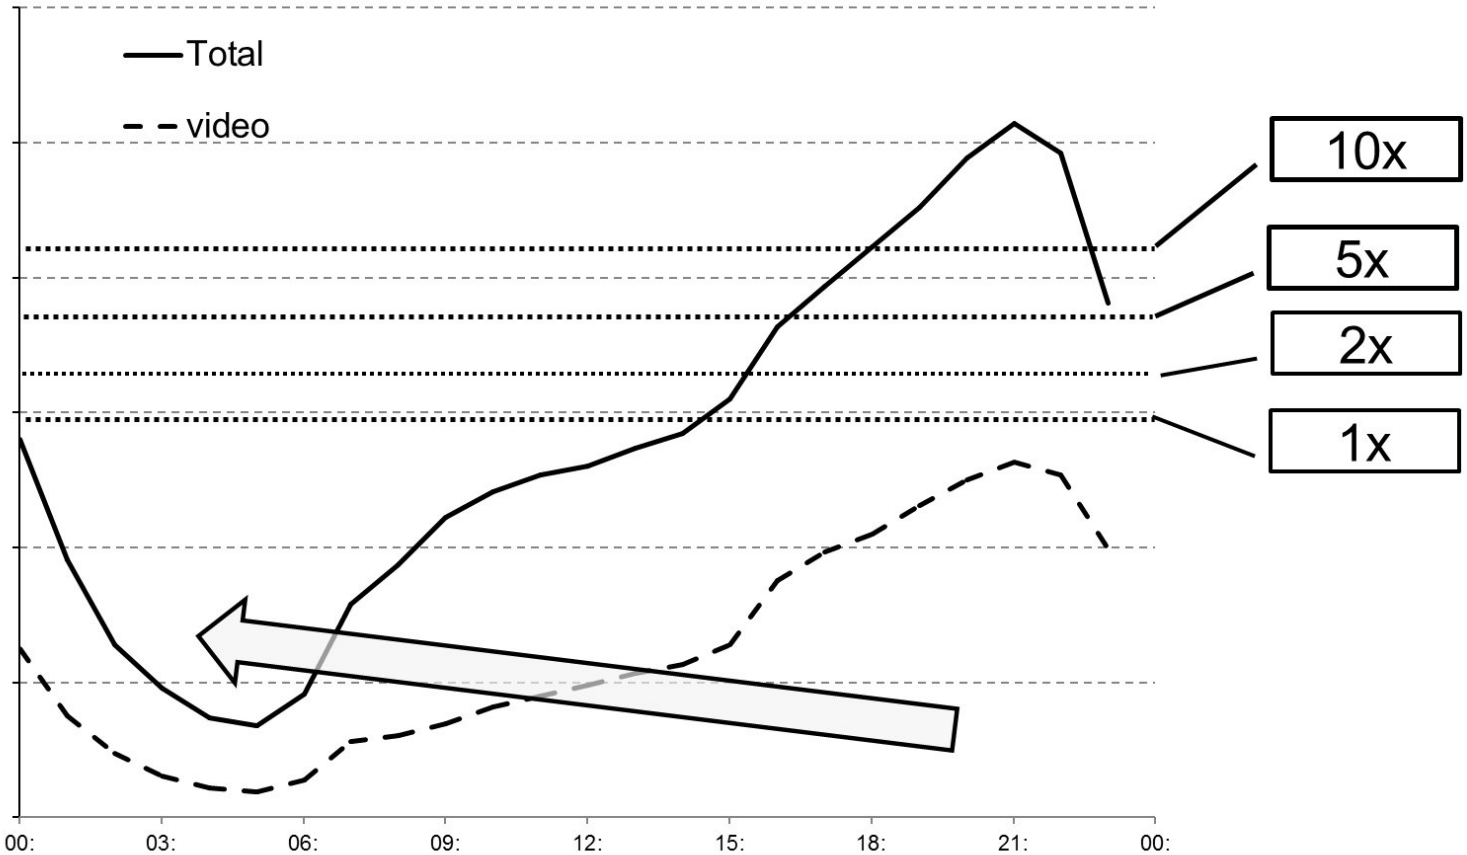

Given the Zipf distribution of content popularity, caching can be very effective even if the selection of content assets to cache is not perfect

How long does content need to be cached?

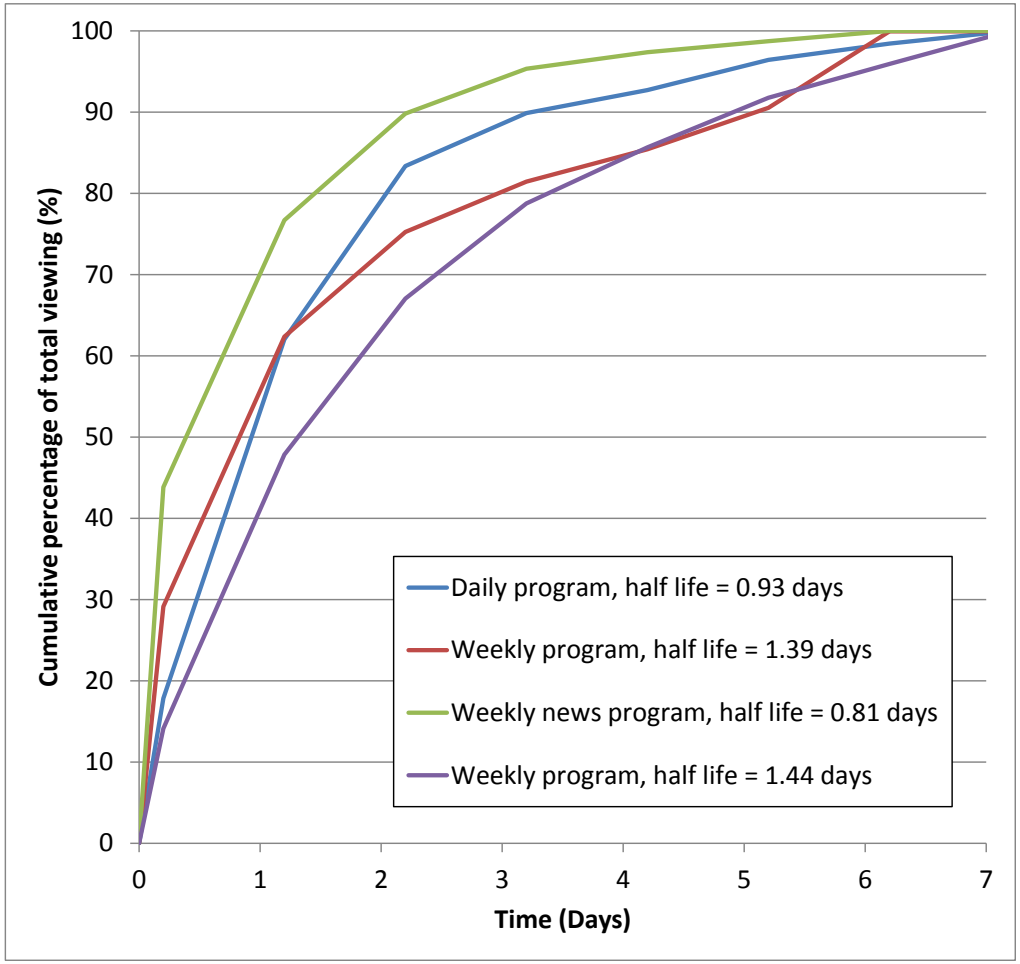

where = half life and is time

Chart showing cumulative viewing of selected "catch-up" TV content assets

Curves matched to exponential distribution

"Half Life" of content calculated: this is the time to get 50% of viewings, or half the time to get 75%, etc.

Daily programs generally have shorter half lives than weekly programs

However, weekly news programs have shorter half lives

Content with a longer half life needs to be retained in the cache for a longer time.

Data generated from BT Wholesale CDN content request log files.

Content Centric Networking

- A new transport protocol built on ***Named Data***
  - A packet address is a content name rather than a location
  - Content is cached at routers rather than simply queued for forwarding
- Created by Van Jacobson et al at the Palo Alto Research Center (PARC)
- Emerging Networks Consortium
  - An initiative to advance the development of Content-Centric Networking
  - Founded by PARC in 2012
  - Founding members include Alcatel-Lucent, BT, France Telecom-Orange, HCL, Huawei, MACH, Panasonic, and Samsung Electronics Co., Ltd.

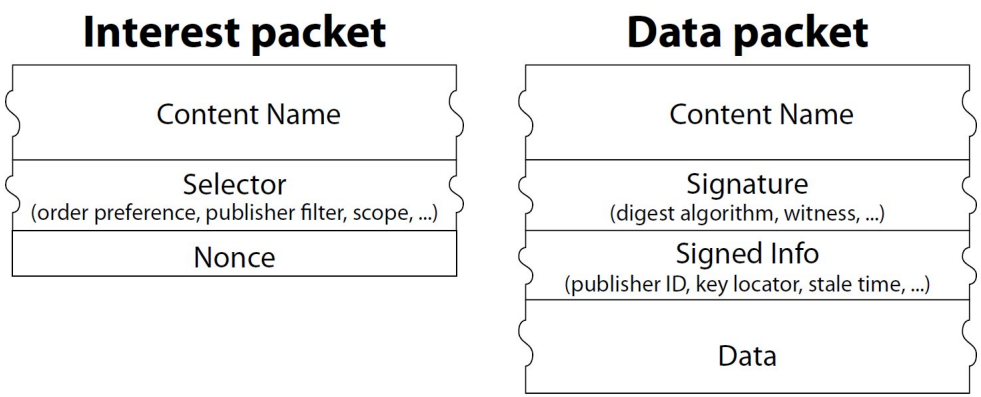

Figure taken from V. Jacobson et al., "Networking Named Content," ACM CoNEXT Dec 2009, pp. 1-12.

CCN Packet Forwarding

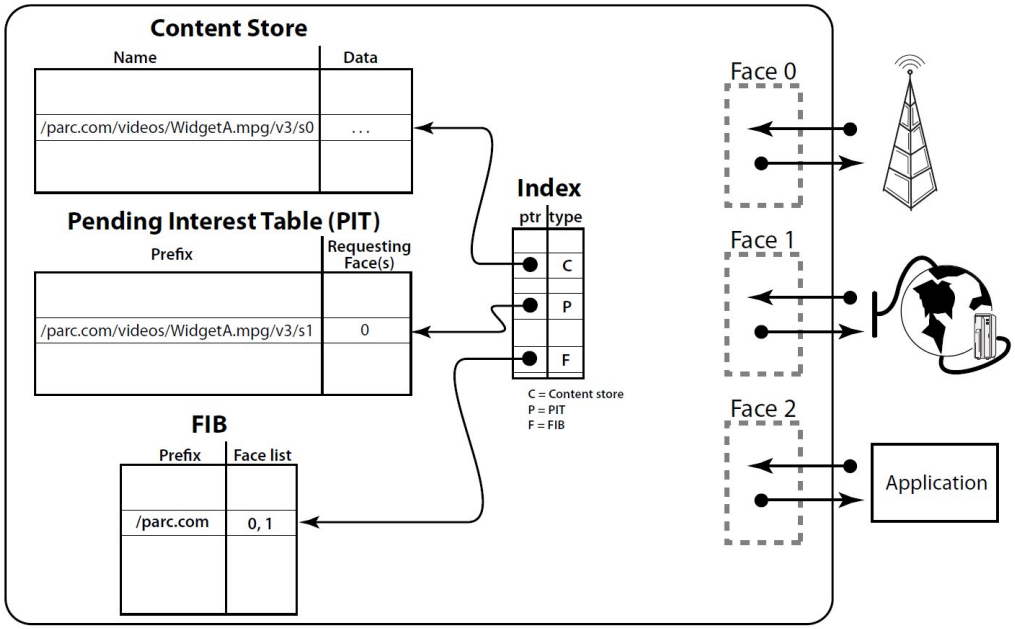

A CCN node operates like an IP node:

- a packet arrives on a face,
- a lookup is done on the name,
- and an action is taken.

Figure taken from V. Jacobson et al., "Networking Named Content," ACM CoNEXT Dec 2009, pp. 1-12.

- Content Store

  Same as buffer memory on an IP router but configured to cache rather then queue Data, thus saving upstream demand and downstream latency

- Pending Interest Table

  'Bread crumb' trail of upstream *Interests* to direct downstream *Data*

- Forwarding Information Base

  A table of outbound *Faces* for *Interests*, where each entry may list multiple *Faces*

CCN Message Processing

Processing *Interest* Packets

1. Lookup the *Content Store*
2. If found, transmit the *Data* and delete the *Interest*
3. Otherwise look up the *Pending Interest Table* (PIT)
4. If found, add the arrival *Face* to the *PIT entry* and delete the *Interest*
5. Otherwise look up the *Forwarding Information Base* (FIB)
6. If found, create a *PIT entry* and forward the *Interest* on one or more *Faces* according to *Strategy*
7. Otherwise delete the *Interest*

Processing *Data* Packets

1. Lookup the *Content Store*
2. If found, discard the *Data* as *Interest* is already satisfied
3. Otherwise look up the *Pending Interest Table* (PIT)
4. If found, store the *Data* in the *Content Store*, transmit the *Data* on all *Faces* listed in the *PIT entry* and delete the *PIT entry*
5. Otherwise the *Data* is unsolicited and should be discarded, but instead may be stored in the *Content Store*

The *Data* 'consumes' the *Interest*

The CCN Strategy Layer

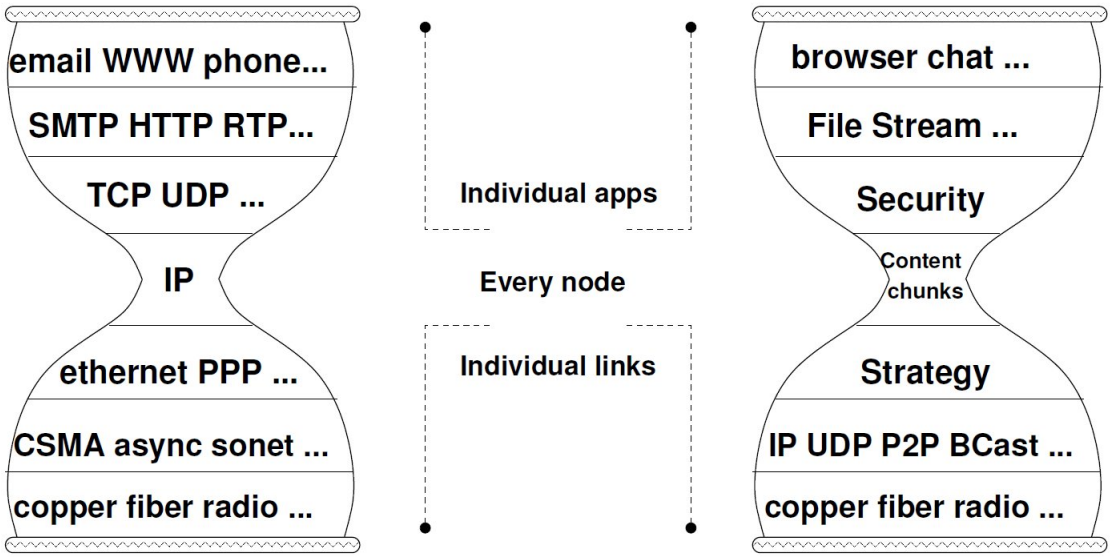

The strategy layer makes the fine-grained, dynamic optimization choices needed to best exploit multiple connectivities under changing conditions.

Figure taken from V. Jacobson et al., "Networking Named Content," ACM CoNEXT Dec 2009, pp. 1-12.

- The FIB may contain multiple Faces for any Content Prefix
  - Data can not loop in CCN
- The Strategy Layer determines where to send Interests
- The Strategy Layer can (occasionally) run experiments:
  - The same interest can be sent on multiple Faces
  - Unlike IP, in CCN, the Data will return on the same path as the Interest
  - Therefore the best Face can be determined
- This provides automatic failover protection

CCN Performance

Some hard to design services 'come for free' with the CCN architecture

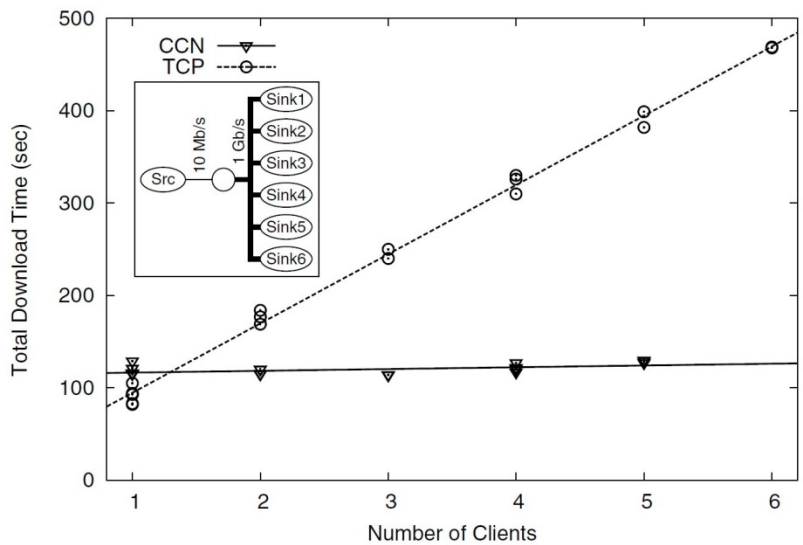

Comparison of HTTP/TCP and CCN

- Each client pulls the same 6MByte file
- Total transfer time is measured for an increasing number of clients
- HTTP/TCP outperforms for one client due to higher, signature, overhead of CCN
- CCN outperforms for more clients due to intermediate node caching

Figure taken from V. Jacobson et al., "Networking Named Content," ACM CoNEXT Dec 2009, pp. 1-12.

- CCN node caching can deliver live TV as efficiently as using multicast
  - The efficiency of multicast can be achieved for all client devices, even those without multicast support
- Retransmission of lost information is more efficient than with multicast
  - Data is retransmitted only over the link where loss occurred, reducing latency
- 'Re-start' services can be implemented easily
  - Data for a live TV program may be cached near to a client wishing to view a program from the start

Summary

- Over 50% of current internet traffic is video
  - This is expected to grow in the future
  - It peaks during the evening and the peak is expected to grow faster than the average
- Content popularity approximately follows Zipf's laws
  - Popular items are viewed orders of magnitude more often than lower ranked items
- Much TV viewing is at time of broadcast
  - Viewing preferences are driven by the broadcast schedule
  - But much broadcast content is not actually live
- Content could be cached close to clients ahead of viewing time
  - Caching reduces the average core network bandwidth
  - Pre-emptive caching reduces peak core network bandwidth
- Content-Centric Networking
  - Content is requested by name rather than address
  - Content caches are an integrated component of the network
  - Some hard to design services 'come for free' with the CCN architecture
  - Could initially be an overlay on existing IP networks, as adopted by the CCNx consortium

The End

Mike Nilsson