Encouraging bandwidth efficiency for peer-to-peer applications

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, NY

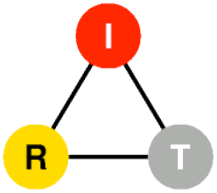

Overview

- Video bandwidth consumption
- Cost of providing video content
- Economics
- Mechanisms
  - network topology indication
  - scavenger service
  - indication of charge
- Problem mainly of economics

Bandwidth consumption

- 4 hours/day of TV @ 18 Mb/s HDTV ➠ 972 GB/month
- Columbia University caps at 350 MB/hour ≈ 252 GB/ month

Economics of the eco system

- Long term, minimize **overall** cost of content delivery
  - across end user, provider, ISP
  - thus, focusing only on efficiency of HTTP misses the complete story
- Components
  - media storage
  - media server bandwidth (can't serve whole ISP from one disk)
  - delivery bandwidth (upstream & downstream)
- Re-use of existing components vs. new components
  - e.g., end user DVR storage vs. dedicated cache servers
  - local bandwidth vs. wide-area bandwidth vs. content provider bandwidth
- Allow cost allocation
  - e.g., rentable caches --> both content provider and ISP benefit

Economics of bandwidth

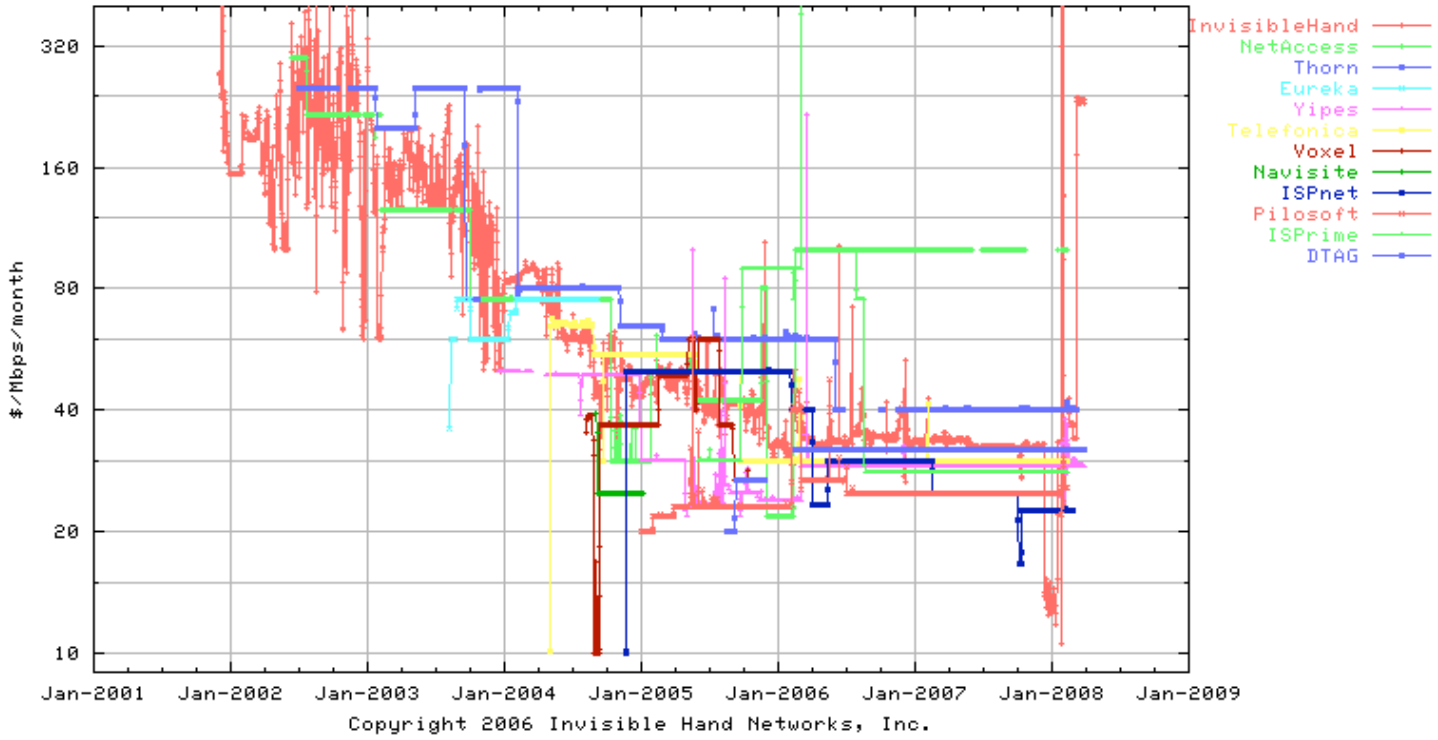

- Transit bandwidth $40/Mb/s/month ~ $0.125/GB
- US colocation providers charge $0.30/GB to $1.75/GB
  - CDNs: $0.08 to $0.19/GB

Cost of bandwidth

- Thus, 7 GB DVD → $1.05
- HDTV viewing ~ $120/month for WAN bandwidth
- Netflix postage cost: $0.70 round-trip
- Typical PPV charges: $4/movie (7 GB)
- Local bandwidth cost is amortization of infrastructure
  - driven by peak load, not average
- Asymmetric vs. symmetric networks

Cost for providing content

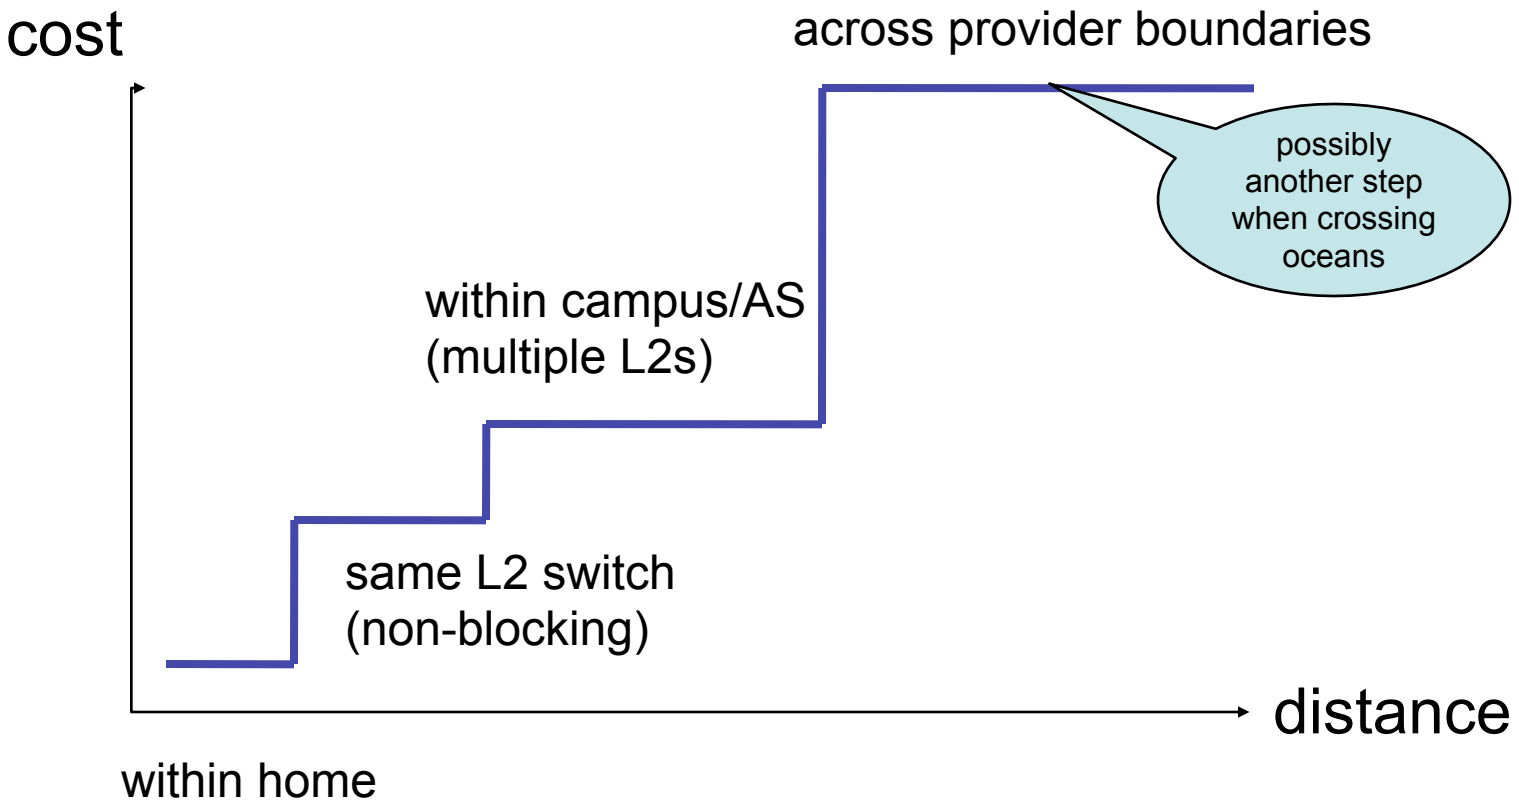

Example: FiOS TV architecture

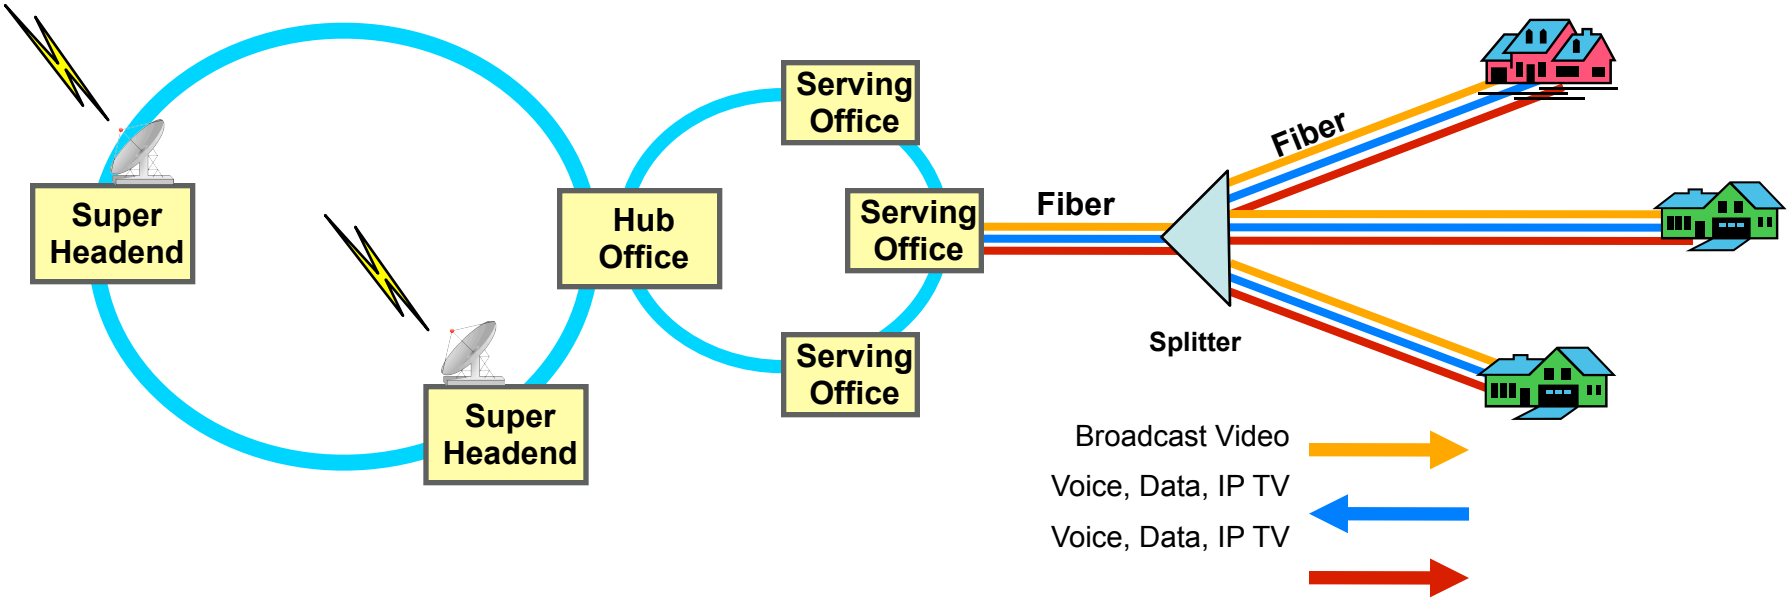

J. Savage (Telecom ThinkTank), Nov. 2006

- 2 national super headends
- 9 video hub offices
- 292 video serving offices

Verizon's FTTP Architecture

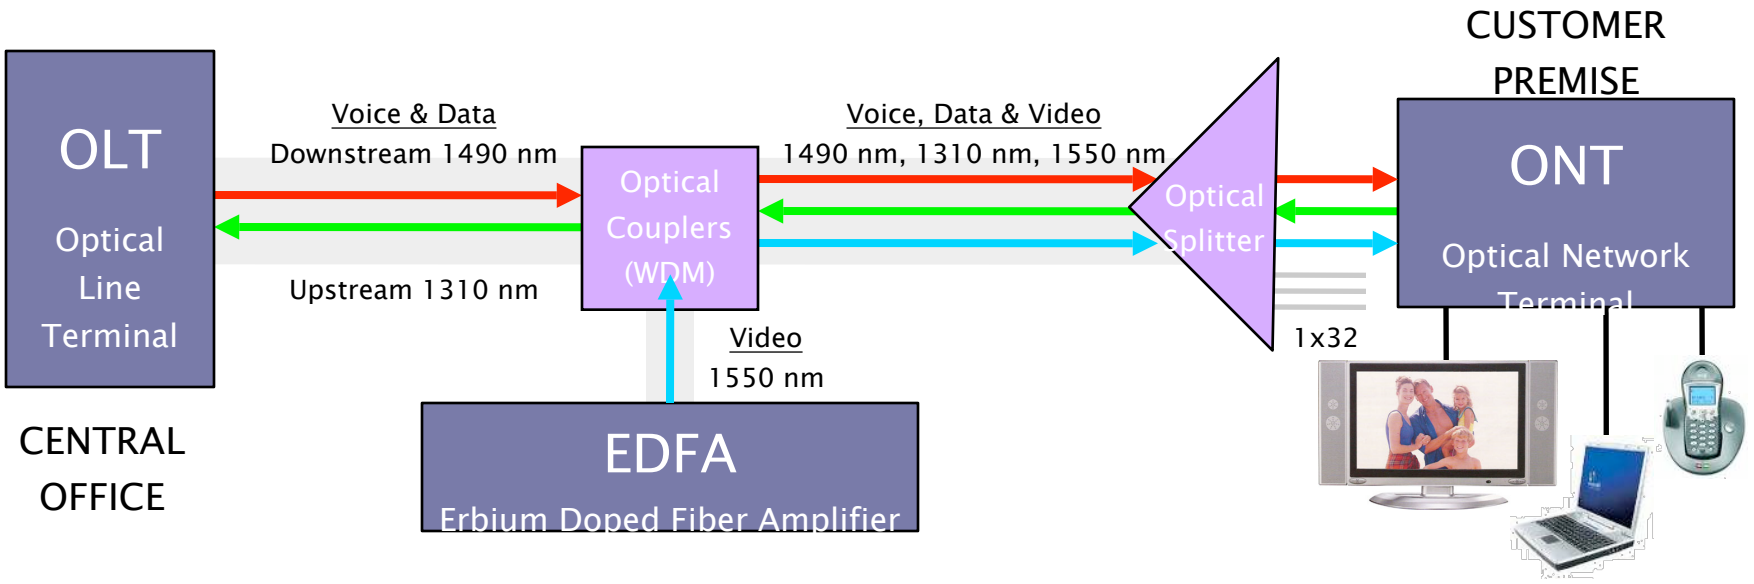

Bandwidth & Services

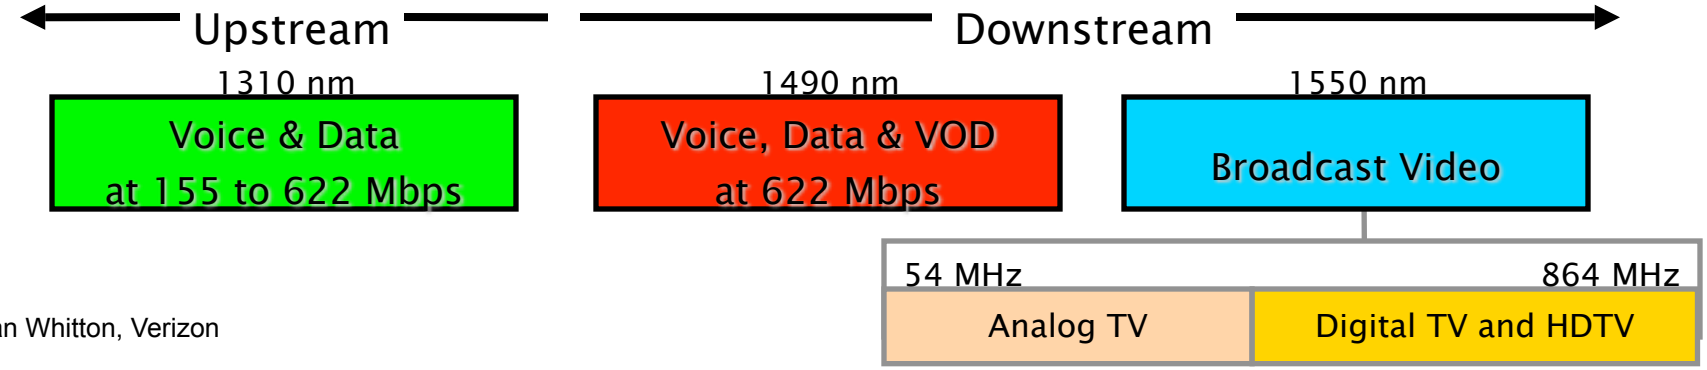

VoD requirements

short clips < 10’ (long tail)

feature-length

- avoid Netflix queue
- avoid stocking 20,000 DVDs

- Example: Superbad grossed $33M during August 17 weekend (in US)
- = roughly 3M viewers
- = roughly 1% of US population
- ⇒ if VoD, each neighborhood has likely one copy
- 2 problems:
  - get initial copy to neighborhood
    - multicast, OTA
  - distribute in neighborhood
- only viable for top 1000 content
- need data on popularity distribution

Example: Columbia University

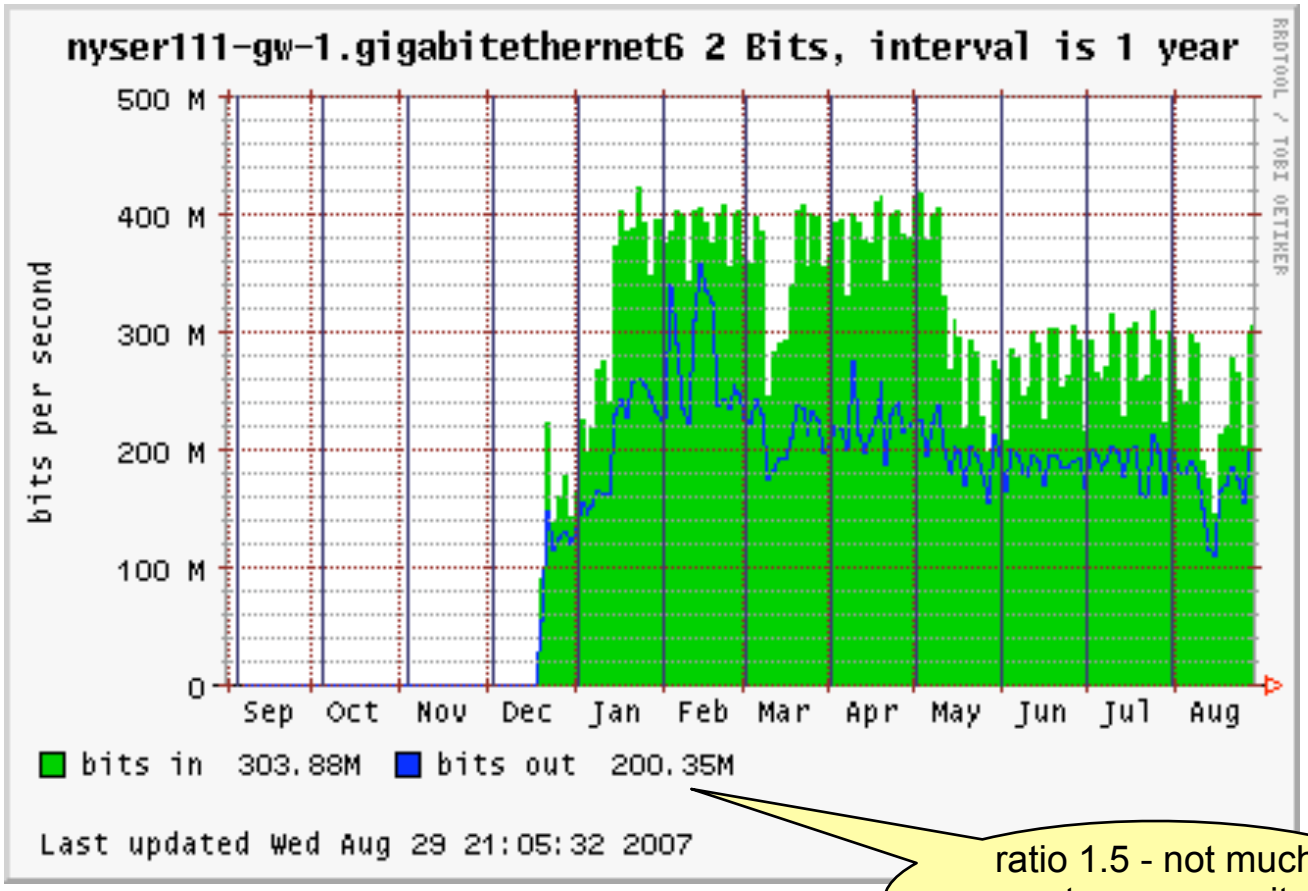

ratio 1.5 - not much upstream capacity left

Diurnal variation

- Use off-hours to download content?

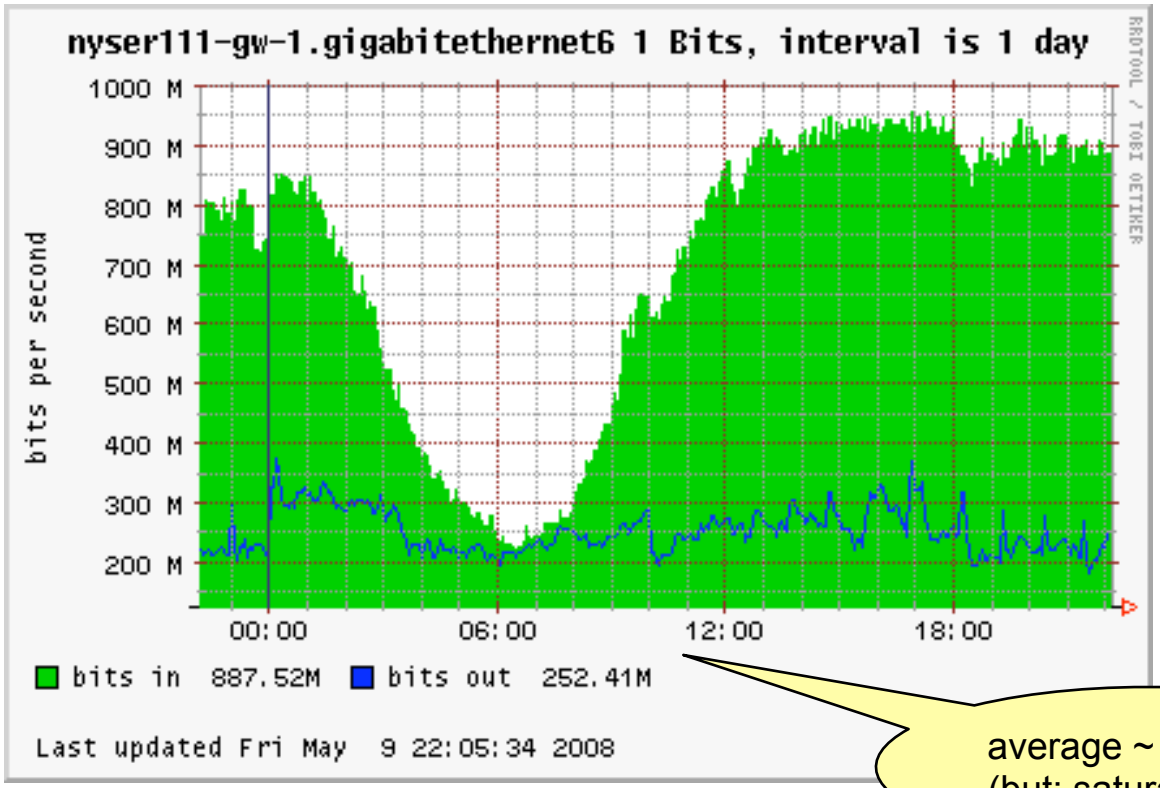

average ~ 0.8 peak
(but: saturated link)

Mechanisms

- Goals and requirements
  - application neutral: not just for BitTorrent or VoD
  - no lawyers
  - no saints ➠ assume economically-rational actors
- Mechanisms
  - network proximity
  - real-time cost and charging information
  - common DiffServ code points
    - Scavenger service

Network topology discovery

- Incentive:
  - lower cost (later)
  - better performance - throughput and delay (e.g., VoIP relay node)
- Indications
  - AS number, ...
  - symmetric vs. asymmetric bandwidth
    - symmetric: local cache
    - asymmetric: cache in ISP network
  - see p4pnet.org
- Mechanisms
  - separate protocol (e.g., web service)
  - STUN
  - DHCP (requires NAT upgrade)

Discovery

- Similar problems
  - discover network topology information server
  - STUN server
  - HELD server
  - LoST server
  - SIP local network configuration
- All likely provided by ISP
- Develop common set of discovery mechanisms
  - DHCP
  - DNS (SRV, NAPTR, ...)
  - anycast
  - ...

Scavenger service

- Explored by Internet2 QoS working group
- Less-than-best effort
  - lower scheduling priority than regular BE traffic
- Avoids self-interference
- Requires no admission control
- Improve RT service performance, but does not address wide-area cost issue
- Requires well-known (or discoverable) DS code point

Indication of charging

- If volume-based, need application-visible charging indication
  - "current cost of 1 GB to 128.59.16.1 is $0.15"
  - "predicted cost in 3 hours is $0.05"
  - "you have 47.5 GB of free local traffic left"
  - "you are currently in penalty box"
- May differ upstream vs. downstream
- Applications can then prefer local content
- or defer to later
  - "Do you want to watch the movie now ($4) or wait until 10 pm ($2.52)?"

DiffServ & Bandwidth charging

- Only two options:
  - limit supply of (high-priority) bandwidth ("1000 minutes of VoIP/ month") OR
  - charge for bandwidth
- Probably need to differentiate "local" and "long-distance" traffic
  - see "free local calls"
- Charging exposes user to risk
  - mis-behaving application or malware
    - need SE-Linux-like capability limitation
  - DoS attacks
    - need permission-based sending

Conclusion

- Simple network mechanisms needed
  - allow applications “to do the right thing”
  - prevent self-interference
  - work for both symmetric and asymmetric networks
  - incentive: better performance or lower cost
- Local network retrieval only works for short-tail content
  - what is the fraction of bandwidth for top-1000 content?