Observations on “Neutrality” and Priority

Charles L. Jackson

clj@jacksons.net

1 October 2010

Roadmap

- Wireless Priority
- Wireless and H-ARQ
- Congestion Control on the Honor System
- Priority Routing More Generally
- Concluding Thoughts

In Wireless, Priority Creates Capacity

- Voice vs. Data
  - Modern wireless voice systems require some capacity to be kept in reserve to protect against signal fading on multiple connections.
  - This reserve can be used for a best-efforts data service.

Example

The system analyzed below can support only about 40 voice users.

With 40 voice users, 40% of the data capacity can still be used.

Uplink tradeoff is not as favorable.

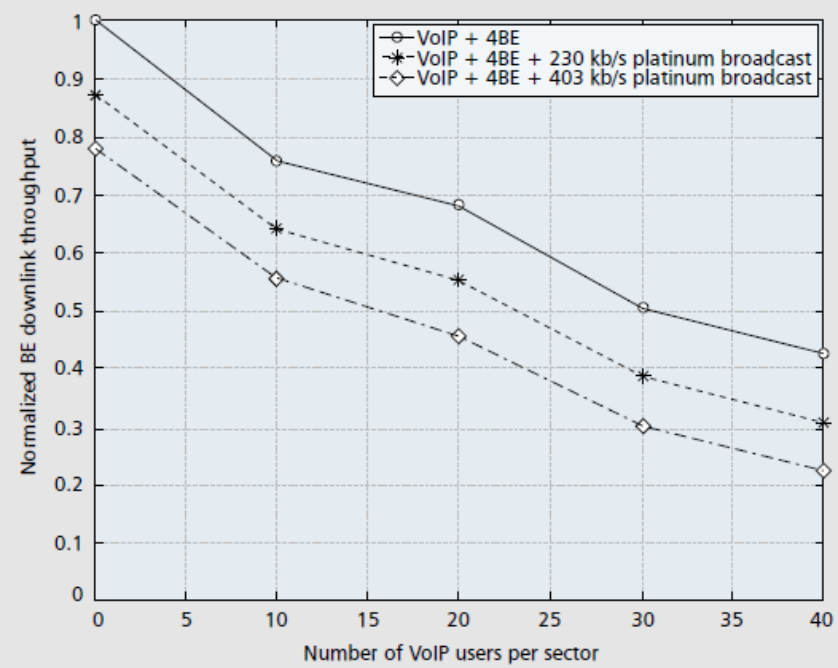

Source: *VoIP over cdma2000 1xEV-DO Revision A,* M. Yavuz et al, IEEE Communications Magazine, February 2006.

Voice Priority

- If one required all bits to be treated equally (no priority for voice over data) then either
  - this background capacity cannot be exploited, or
  - voice performance declines markedly.
- There will larger consumer benefits in a world with two wireless services—voice and best-efforts data than in a world with one service—voice.

ARQ

- Data links can detect blocks with errors and retransmit them. Called Automatic Repeat reQuest (ARQ).
- Ethernet LANs don't bother with ARQ.
  - Relatively few errors on wired LANs.
- The rate of block errors is much higher (thousands of times higher) in wireless than in a LAN; consequently, wireless links often use ARQ.

Wireless and H-ARQ

- Modern wireless links can also use a technique called hybrid ARQ (H-ARQ) that allows correcting a frame that arrived with errors without incurring the cost of a full retransmission.
- Basic idea:
  - Base transmits to mobile a seven-bit plus parity message (e.g., “1011110P”).
    - P: a parity bit that makes the number of 1s even.
  - Mobile examines number of 1s in received block, finds it is not even, concludes error occurred.
  - Requests partial retransmission (some help).
  - Base transmits, “The message starts with ‘101’.”

H-ARQ and Wireless Internet Traffic

- Should a wireless system use H-ARQ on a TCP datagram? How about a UDP datagram carrying voice?
  - TCP: If a datagram is not retransmitted at the link level, will be fully retransmitted later by the TCP process—consuming more wireless capacity than would H-ARQ retransmission.
    - Pay me now or pay me more later.
  - UDP carrying voice: A voice system can tolerate occasional losses of packets.
    - Pay me now or don't pay me at all.
- Large literature on "cross-layer design"

Internet Congestion Control

- Works on the honor system.
- Violating the honor system can improve performance for the application that violates.
- “In the current Internet architecture, congestion control depends on parties acting against their own interests. It is not in a receiver’s interest to honestly return feedback about congestion on the path, effectively requesting a slower transfer. It is not in the sender's interest to reduce its rate in response to congestion if it can rely on others to do so.
  - Welzl et al., http://tools.ietf.org/html/draft-irtf-iccrg-welzl-congestion-control-open-research-05.txt , Aug 31, 2009. Emphasis added.

Congestion Control II

- *Peterson and Davie,* 4th Ed., p. 470

  > It is possible for an ill-behaved source (flow) to capture an arbitrarily large fraction of the network capacity. . . . Such an application is able to flood the Internet's routers with its own packets, thereby causing other applications' packets to be discarded.

- An ethical but unhappy developer on the Google Chrome web browser project:

  > There's not much difference between 6 connections per server and 8 total connections per server. . . But, Firefox bumped their total limit of connections per server way up, and now there is a big difference between 6 and 15 **:-(**

- Should ISPs do anything about applications that violate this vital honor system?

Glass's Denial of Service Attack

- Brett Glass operates Lariat, a small ISP in Wyoming.
- DS-3 (45 Mbps) connection to larger Internet.
- In May, 2009 Microsoft made a large security update available for Windows.
- Many users' machines started downloading.
- The DS-3 link saturated and service quality plummeted.
- Glass restored service by throttling back the Microsoft downloads.

Massive Network Failures

- Earthquake in December 2006 took out 12 of the 18 submarine cables lying on the ocean bottom between Taiwan and the Philippines.
  - One ISP restored service by blocking video downloads and gaming traffic.
- Netgear hard coded a router to query the University of Wisconsin's network time protocol (NTP) server.
  - If a user's firewall blocked incoming UDP packets, the router would send one query per second.
  - 700,000 devices
  - If all active, 426 Mbps towards U Wisconsin.
- Would it be reasonable for an ISP to drop an NTP query headed to U Wisconsin?

Priority Routing More Generally

- Free Press claimed:

  Priority routing is a “zero sum game” because speeding one packet slows another.

- But, delaying a VoIP packet imposes greater costs than delaying a file-download packet.
- A zero-sum ambulance analogy:
  - Pulling over to let an ambulance pass speeds up the ambulance by 10 minutes but slows down 100 commuters by 6 seconds.

Questions to Consider

- If there were to be a catastrophic failure of transmission systems, would you prefer that your ISP as well as other ISPs blocked BitTorrent and gaming?
- Do you agree, that in a world of network neutrality, it is likely that

  Aggressive but delay-tolerant applications will thrive and latency-sensitive applications will stumble along. Regulation and the physics of networks rather than consumer preferences will determine which firms and applications succeed in the market.