Back to the Future Part 4: The Internet

Simon S. Lam
Department of Computer Sciences
The University of Texas at Austin

IP won the networking race

- Many competitors in the past
  - SNA, DECnet, XNS
  - X.25, ATM, CLNP
- IP provides end-to-end delivery of datagrams, best-effort service only
- IP can use any link-layer technology that delivers datagrams (packets)

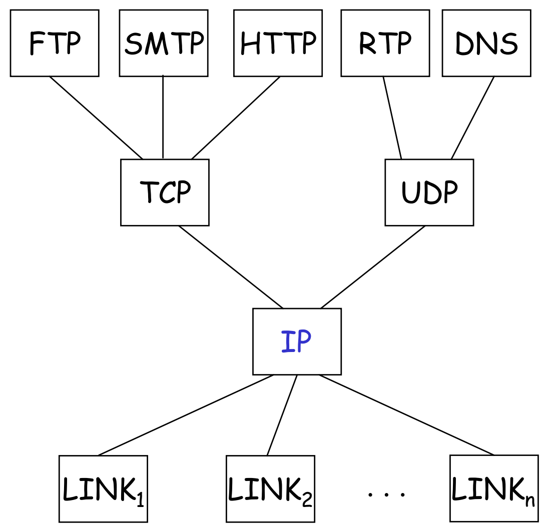

IP's underlying model is a network of queues

- Revolutionary change from the circuit model (Kleinrock 1961)
  - Each packet is routed independently using its destination IP address
  - No concept of a flow between source and destination, no flow state in routers
- Unreliable channels, limited buffer capacity
- Prone to congestion collapse

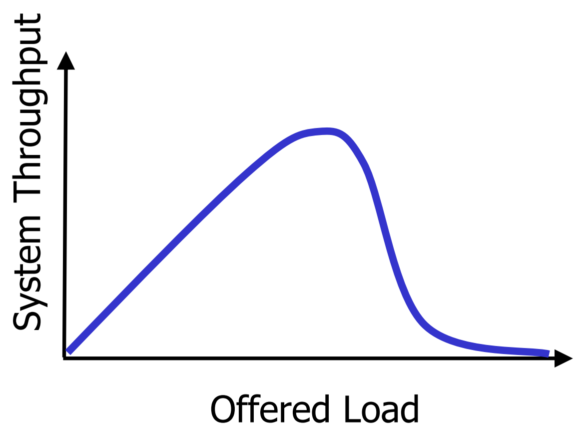

Congestion collapse—ALOHA channel

- The ALOHA System (Abramson 1970)
- Poisson process assumption
  - pure ALOHA throughput

$$S = G e^{-2G} \quad \text{(Abramson)}$$

  - slotted ALOHA throughput

$$S = G e^{-G} \quad \text{(Roberts)}$$

- ARPANET Satellite System (1972)

Congestion collapse—ALOHA channel (cont.)

- The early ALOHA System was vastly under-utilized

- Backoff algorithm for slotted ALOHA (Kleinrock-Lam 1973)
  - Retransmit a collided packet randomly into one of K future time slots
  - Poisson process assumption implies

    $K \to \infty$

Congestion collapse—ALOHA channel (cont.)

- ASS Note 48 (Lam 1973)

K = 15

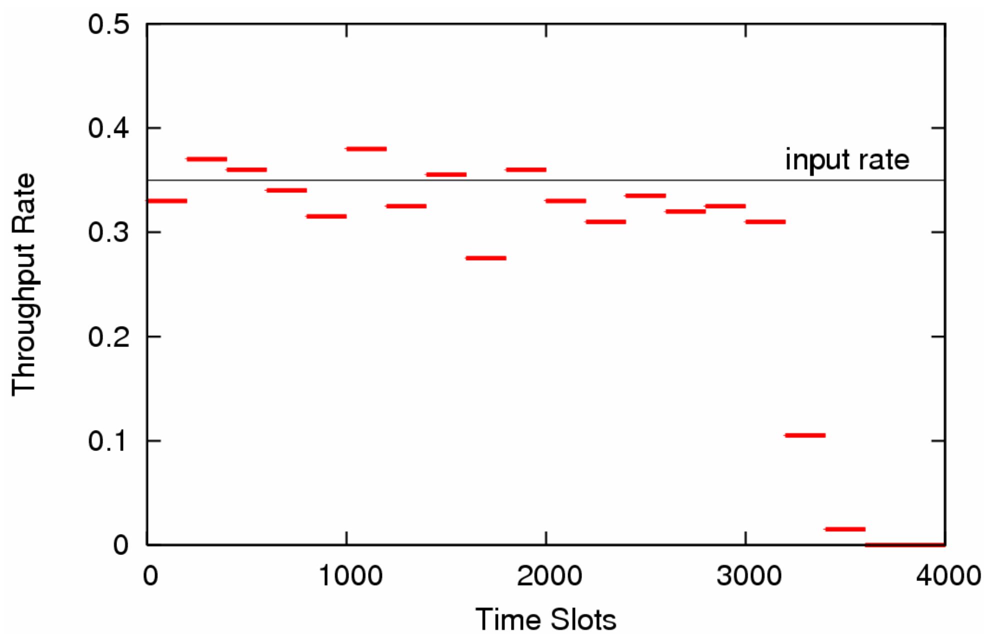

Congestion collapse—ALOHA channel (cont.)

- Adaptive backoff algorithm (Lam 1974): Retransmit a packet with m previous collisions into K(m) slots, where K(m) is monotonically nondecreasing in m

K(1) = 10

K(m) = 150, $m \geq 2$

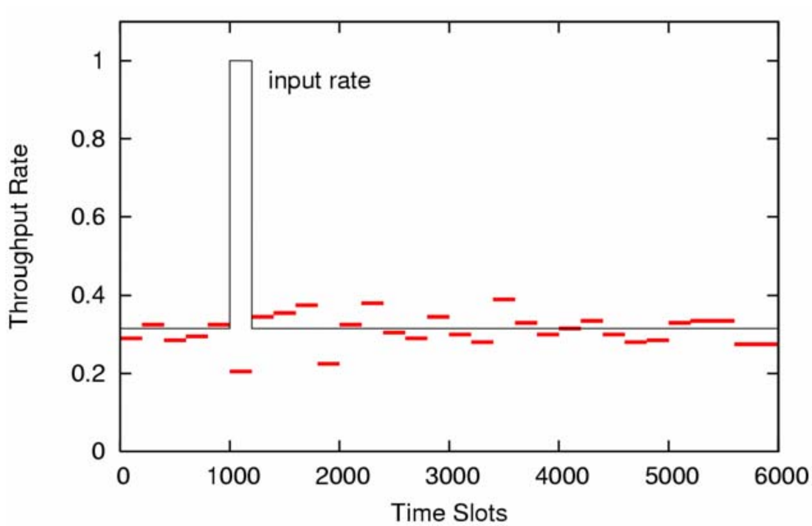

Congestion collapse—packet networks

- Network of queues with limited buffers
  - Static window flow control in TCP not sufficient (Lam, late 1970s)

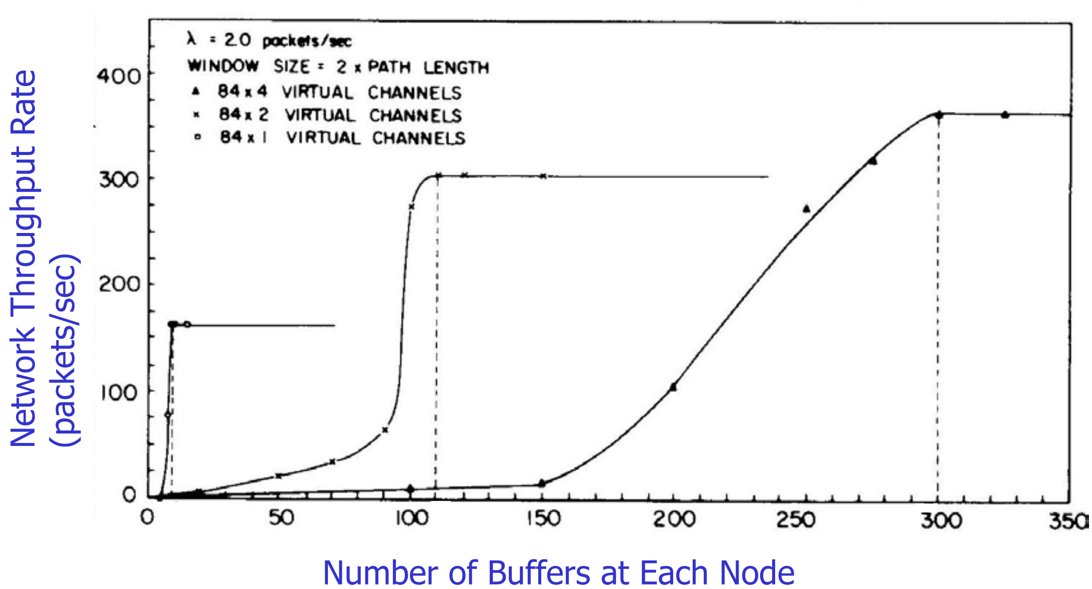

Internet congestion control

- Van Jacobson's algorithms for TCP congestion control (late 1980s)
  - main reason for stability of the current Internet
- UDP does not perform congestion control
  - preferred by voice and video applications
- More and more voice and video traffic will impact Internet stability
  - I believe in differentiating voice and video flows as well as flow admission control

Efforts to extend/replace IPv4 (past 15 years)

- IP multicast
- QoS support - IntServ, RSVP, DiffServ
- Active Networks research program of DARPA
- IPsec – retrofitting IP with security
- IPv6 – replacing IPv4
  - 128 bit IP address
  - flow concept to support QoS
- Mobile IP
- . . .

Don't mess with IP?

- In recent years, the research community has moved on to other areas
  - P2P overlays
  - Wireless (ad hoc networks, sensor networks, satellite networks)
  - Measurements
- But the IP foundation, currently relying on over-provisioning, still needs work

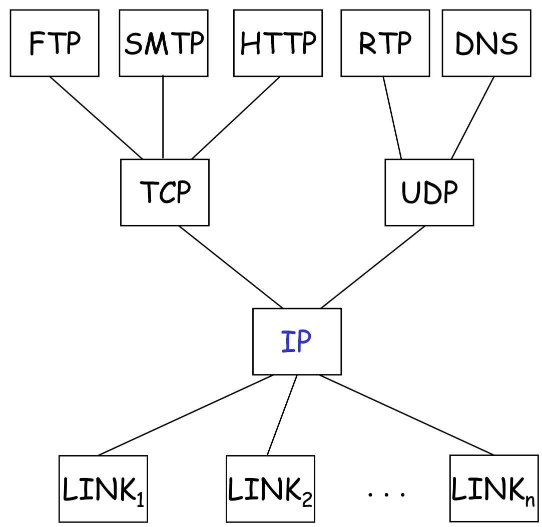

Is P2P overlay a panacea?

- P2P overlay supports many new distributed applications
- P2P overlay is inefficient in its use of underlying Internet resources
- P2P overlay does not directly address IP's foundational issues (stability, QoS)

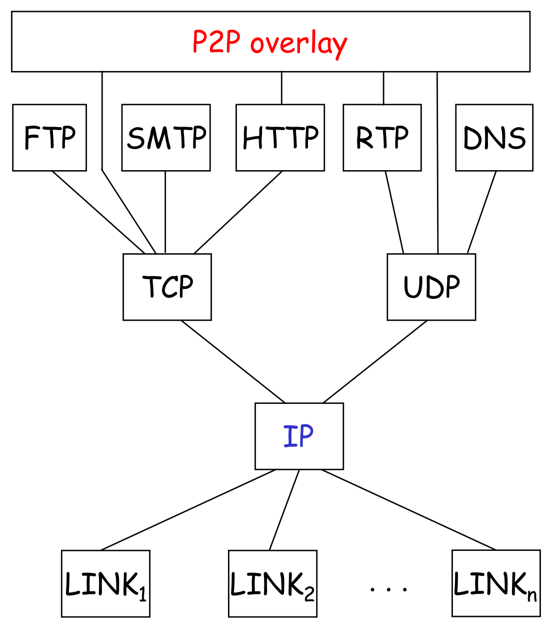

IP as the future universal interface

- Payload in the form of IP packets to enable new applications across different telecom networks

  analog → digital → packet

- Recent news
  - Microsoft unveiled plans for developing IPTV (October 2003)
  - SBC to invest $6 billion on fiber to home for TV services (June 2004)

Changing network traffic mix

- Much more voice and video traffic
- Current traffic of a major telecommunications carrier
  - Circuit switched voice 1.2 petabytes/day
  - Internet traffic 1.5 petabytes/day
- Television services over IP
  - Back of the envelope calculation:

    4-8 Mbps, 10,000 seconds/day, 10^8 TV sets

    → 500-1000 petabytes/day to end users

A pragmatic approach

1. Learn from the evolution of Ethernet

Ethernet technology today is very different from Ethernet technology 20 years ago.

- Transmission rates: 10 Mbps, 100 Mbps, 1 Gbps, 10 Gbps
- Switching protocols:
  - CSMA/CD on a cable,
  - CSMA/CD on a hub,
  - collision-free switching,
  - full-duplex point-to-point, both WAN and LAN
- A variety of coding techniques and media

Only the Ethernet frame interface remains the same.

A pragmatic approach (cont.)

2. Accept in some form a competing idea that has been vanquished again and again by IP but refuses to die and go away:

virtual circuit packet switching

- X.25 1970s – 1990s [L. Roberts]
- Frame Relay 1980s – present
- ATM early 1990s – present
- Label switching, MPLS late 1990s - present

Real circuit switching also under IP

- IP over SONET
- GMPLS forwarding based on TDM time slot, wavelength, or optical port
  - Nesting of label switched paths (LSPs) reminiscent of multiplexing/demultiplexing hierarchy in telephone networks

IP should be a "big tent"

- To "rule" the world of communications, IP has to attend to the needs of new constituents (voice, video)
  - multiple services to support diverse applications
- For the research community, over-provisioning should be considered a temporary fix, not a permanent solution
- While the core is over-provisioned, access paths to the core are not

The hourglass shape reconsidered

- Although link-layer technologies are diverse, including virtual and real circuits, only best effort service is available to Internet applications

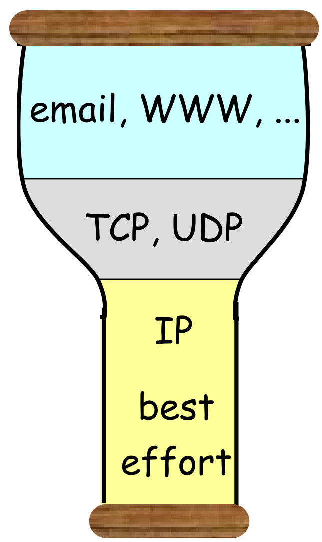

IP needs a broader base

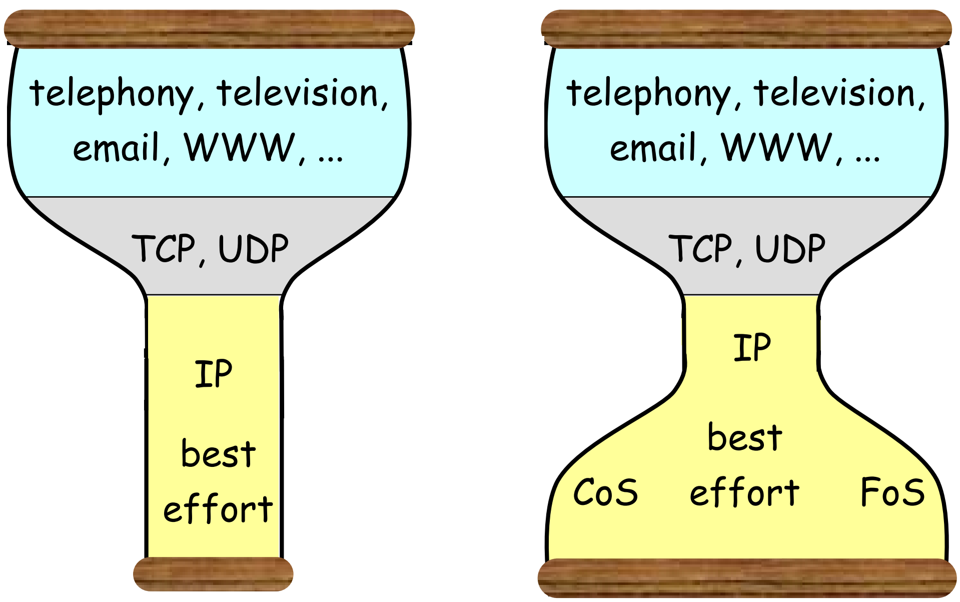

CoS Class-oriented Service

FoS Flow-oriented Service

What is flow-oriented service?

- Dynamic signaling, flow admission control, flow state, quantitative QoS metric
  - Flows subject to admission control rather than random packet drop
  - Make use of virtual or real circuits in link layer

- Am I reviving IntServ?

  Not exactly

- Per-flow state and dynamic signaling are not scalable!

  I know

Two good engineering ideas for voice and video

1. Flow aggregation

- Current examples
  Virtual paths in ATM, Label stacks in MPLS, RSVP aggregation
- Each flow is routed along a sequence of "virtual channels" each of which carries a flow aggregate
- Flow aggregation reduces state information and signaling overhead thus improving scalability
  - In the extreme case, a router has just two flow aggregates (voice and video) for each outgoing channel

Two good engineering ideas (cont.)

2. Statistical guarantee

- A natural service guarantee for a voice or video flow is the flow's loss probability for a given packet delay bound

  Prob[packet delay > x] < ε

  Very hard problem! Substantial research in the past but needs much more work to be applicable.
- Statistical multiplexing gain from flow aggregation

Major research issues

- How to derive service guarantee of a flow from the service guarantee provided to a flow aggregate?
- Dynamic configuration and provisioning of virtual channels for flow aggregates.
- How to efficiently compute the end-to-end statistical guarantee to a flow, under practical assumptions?
  - design, modeling, analysis
  - approximation methods
  - measurement-based techniques

QoS research is not done

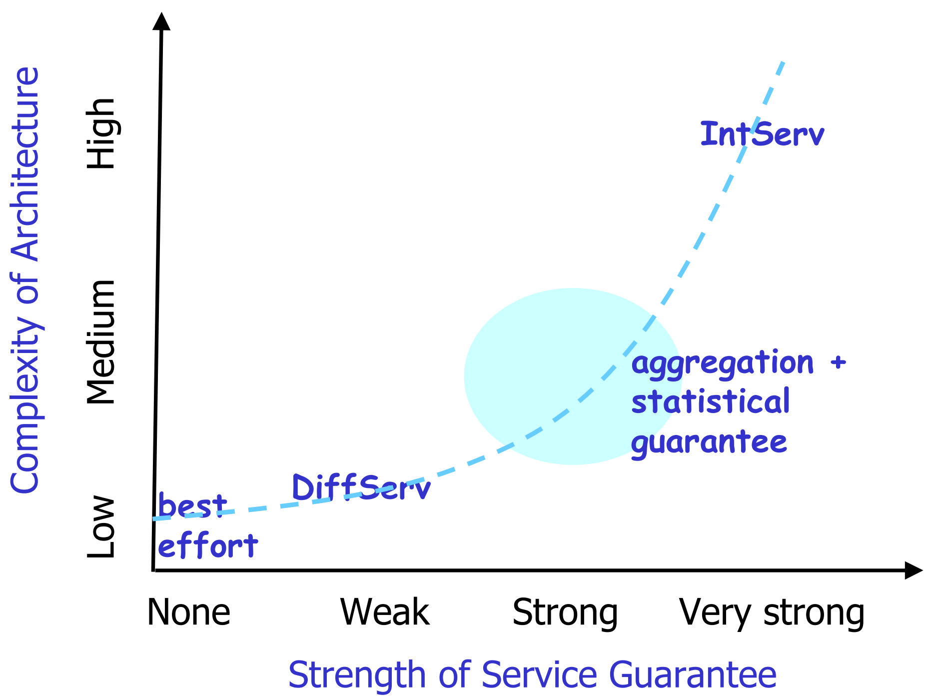

Inter-provider QoS is a major challenge

- Business and legal issues
- Framework for competitive ISPs to cooperate
  - A quantitative QoS metric for inter-provider agreement
  - A small set of standardized traffic specs for voice and video
  - . . .
- It would be nice to have a de facto standard!

How to get one?

- The SSL model
  - Application-driven—the need to secure web transactions
  - Now SSL used for other applications as well
- The FedEx model
  - Profit-driven—someone takes risk

Possible scenario: Some large ISP takes risk and provides QoS services to a large part of the Internet. Success leads to universal global coverage and a de facto standard.

Conclusions

- For the research community, over-provisioning should not be considered a solution
- To become the universal telecom interface, IP needs to be a "big tent"
  - A flow-oriented service needed to support television and telephony services
- QoS research is not done
  - Flow aggregation and statistical guarantee merit further investigation

Conclusions (cont.)

- From history
  - Internet—almost 30 years from initial research to commercial deployment
  - Packet radio—about 25 years
  - QoS research began in late 1980s

  Widespread commercial deployment of QoS within 10 years!