Time-Driven Priority Scheduling
operating principles, benefits, and implementation

Mario Baldi
Turin Polytechnic
mario.baldi@polito.it
staff.polito.it/mario.baldi

The context

- Quality of service
- Packet networks
  - IP
  - ATM
  - …

Am I wasting your time?

- Are you happy with the service provided by today's Internet?
- Do you think that network resources are (going to be) plenty and (almost) free of charge?
  - Transmission capacity
  - Switching capacity

Talking about network resources

- Transmission capacity:
  - Can intercontinental links be (almost) free?
- Routing and switching capacity: technological advances
  - Cisco CRS-1 (announced in 2004)
    - 92 Tb/s →72 line-card shelves + 8 switching fabric shelves = 80 shelves
    - 1.28 Tb/s (or 640 Gb/s) per shelf

Talking about network resources

- Size
  - Power consumption
  - Cooling
  - Etc.
- Technological advances
  - Cisco 12000 (announced in ?)
    - 1.28 Tb/s (or 640 Gb/s) in one shelf
- What happened to Moore's law?

Quality of service today

- Overprovisioning
- Differentiated Services (DiffServ)
  - Class of service
  - Preferential overprovisioning
- Other (un-deployed) solutions
  - Integrated Services (IntServ)
  - Sophisticated queuing algorithms
- Why are they not being deployed?
  - In my view, complexity and low cost/benefit ratio

What can be done instead?

- Guaranteed quality
  - Enable new services
- High network utilization
  - Keep costs down

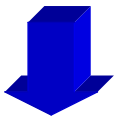

- New affordable services for users
- New revenue for service providers
  - Users are willing to pay for service
  - Providers have an easy way of billing

How can this be done?

Time-Driven Priority

- Packet scheduling algorithm
- Low complexity
  - Computational (none)
  - Implementation
- High scalability
  - No per flow processing
  - No per flow state

Properties

- Guaranteed quality of service (QoS)
  - No loss (due to congestion)
  - Known delay
  - Low jitter
    - ~100 μs
    - Independent of number of hops
  - Per flow (if needed)
- High resource utilization How?
  - No overprovisioning
  - ~ 90% (traffic with QoS guarantees)

Some considerations

- Resource reservation and admission control:
  - Always with guaranteed QoS provision
  - RSVP (Resource reSerVation Protocol)
- Connection oriented
  - MPLS (Multi-Protocol Label Switching)
- No new protocols are needed
- Packet switching features preserved
  - Statistical Multiplexing
  - Best effort
  - Differentiated Services

Some considerations

- Common Time Reference (CTR)
  - GPS (Global Positioning System), USA
  - Galileo, Europe
  - Generated and ristributed within
    - Work in progress
- Particularly suitable for streaming media
  - Do we care?
  - With what are networks going to be filled up?

Periodic Bursty Transmission

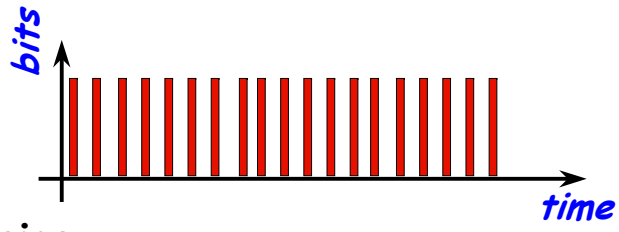

- Voice
- Videoconferencing

No shaping ➡ minimum delay

Let's see…

Complex Periodicity

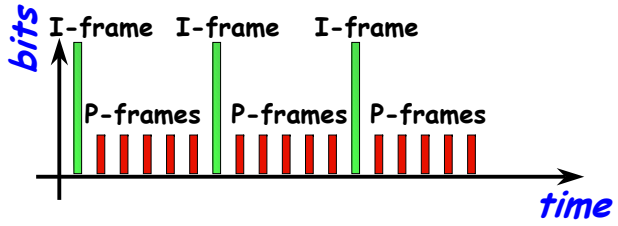

MPEG video

Ubiquitous Deployment?

Not necessarily

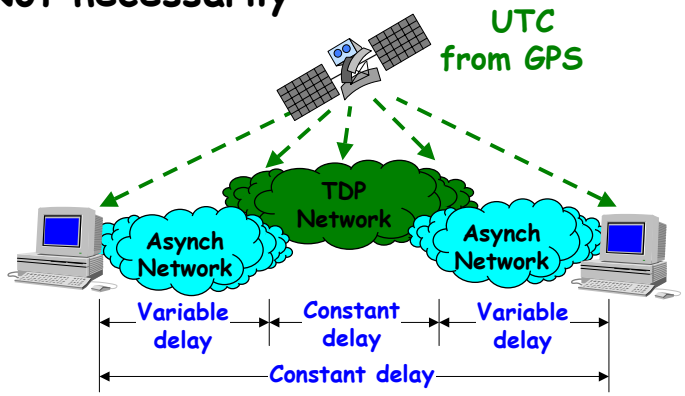

… even though the benefit would be higher

More on Scalability

- No per-flow state
- Simplified switching
  - Trivial non-blocking scheduling with 2x speed-up
  - Limited impact of a blocking switching fabric
    - E.g., low complexity Banyan

Want More Scalability?

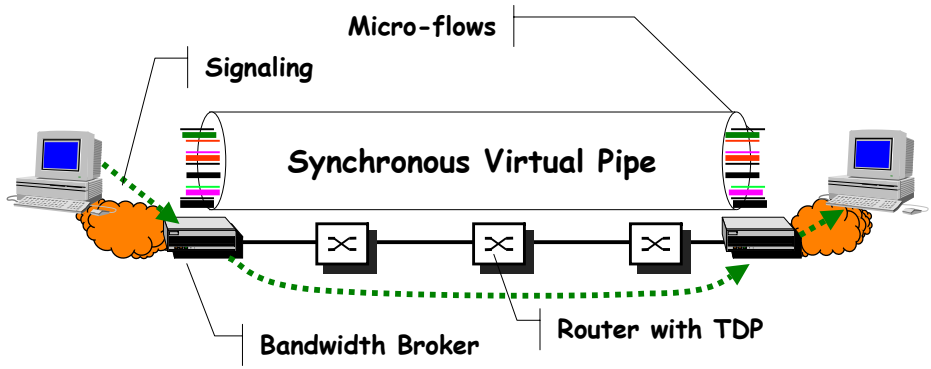

No need for hop-by-hop per flow signaling

Implementation

- Free BSD based prototype
  - ALT-Q
  - Dummynet
- ~1,400 lines of code
  - Routing software: ~15,000 lines
    - Does not include QoS support (RSVP, classification, queuing)
  - TDP (complexity): <10%