A Primer on Network Management

George Ou
Senior Analyst
ITIF
gou@ITIF.org
www.ITIF.org

The Information Technology & Innovation Foundation

ITIF

Summary

- Background information on the Internet
- Explain what network management is
- Quality of Service (QoS) and the Internet
- Wireless – The new frontier of the Internet
- Implications of regulatory proposals on the management and operation of networks

Why the engineering matters

- The debate over net neutrality has evolved in at least three main stages.
  - Issues of blocking or degrading (e.g., Madison River)
  - Should the Internet permit multiple tiers of priority and pricing?
  - What kind of network management, if any, should be permitted on the Internet?
- Effective Internet and telecom policy relies on a solid technical understanding of how the Internet works.

What is the Internet?

- The Internet is a network of networks
  - A federation of independent networks
  - Transmits data in little pieces called packets
- The Internet is also a packet delivery service
  - Network operators ship **packets;**
    FEDEX, UPS, and DHL ship **packages**
  - Network operators are similar to shipping companies; but the difference is that network operators hand off packets

Circuit switching vs. packet switching

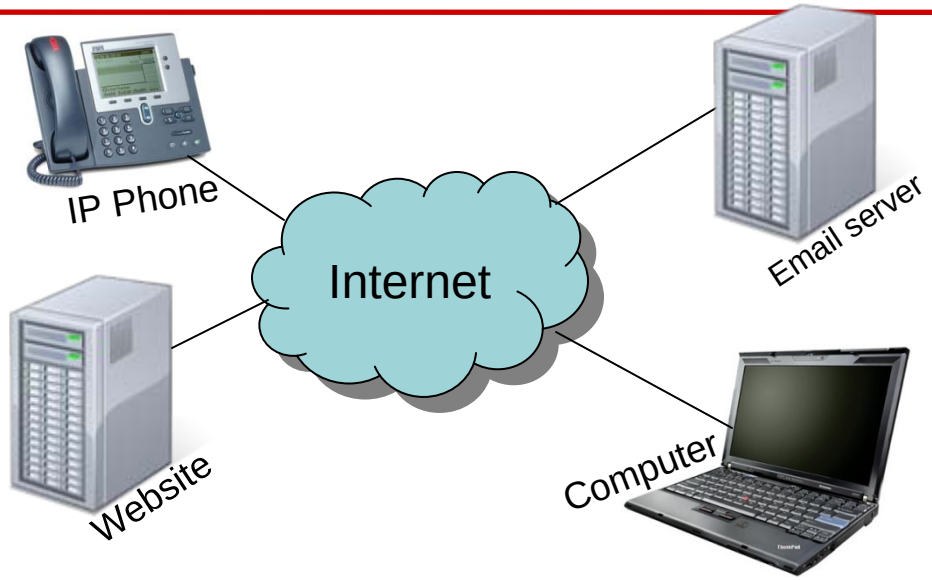

Circuit switching network	Packet switching network
Fixed bandwidth allocation	Variable/Dynamic bandwidth allocation
Inefficient bandwidth allocation	Efficient bandwidth allocation
Inherently predictable bandwidth and latency	Inherently unpredictable bandwidth and latency
Limited functionality	Extremely flexible with wide range of functionality

The Information Technology & Innovation Foundation

ITIF

PSTN = Public Switched Telephone Network

George Ou
December 11, 2008

Internet communication standards

- The communication standard of the Internet is Transmission Control Protocol/Internet Protocol (TCP/IP)
- TCP handles data transmissions
  - Connection establishment and error correction
  - Transmission rates and congestion control
- IP handles the addressing and routing of packets

Three distribution models of the Internet

- Client-Server
- Peer-to-Peer (P2P)
- Content Delivery Network (CDN)

Client-Server

- Pros
  - Quick and easy to set up
  - Good for few users or low bandwidth applications

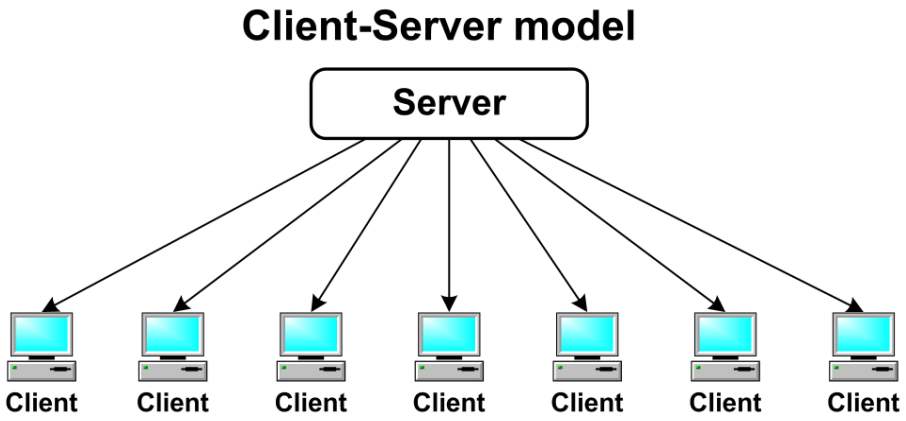

- Cons
  - Limited bandwidth
  - Limited content delivery scalability
  - Higher latency to distant locations

Peer-to-peer (P2P)

- Pros
  - Unlimited file distribution scalability
  - Shifts most costs to end-users
  - P4P (improved P2P) helps decrease strain on core of Internet
- Cons
  - End-user pays distribution costs
  - 2x more bandwidth load on broadband
  - Low resolution video-on-demand due to out-of-order deliver

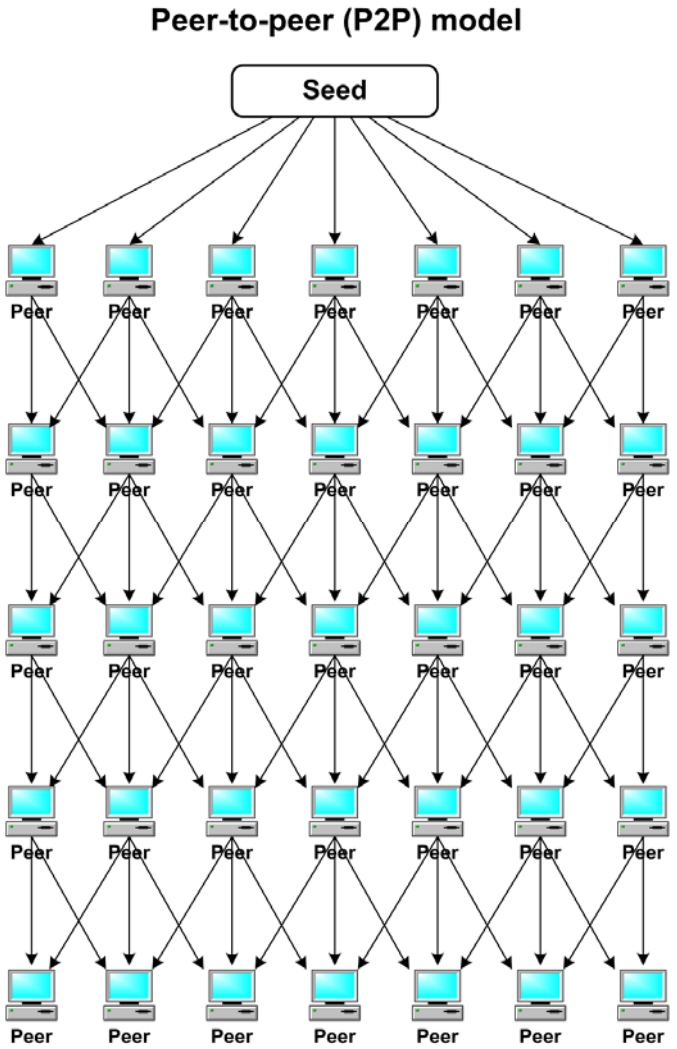

The Information Technology & Innovation Foundation

ITIF

Content Delivery Network (CDN)

- Pros
  - Unlimited file distribution scalability
  - Least bandwidth load on the core or edge of Internet
  - High quality video-on-demand
  - Doesn't offload costs to end-users

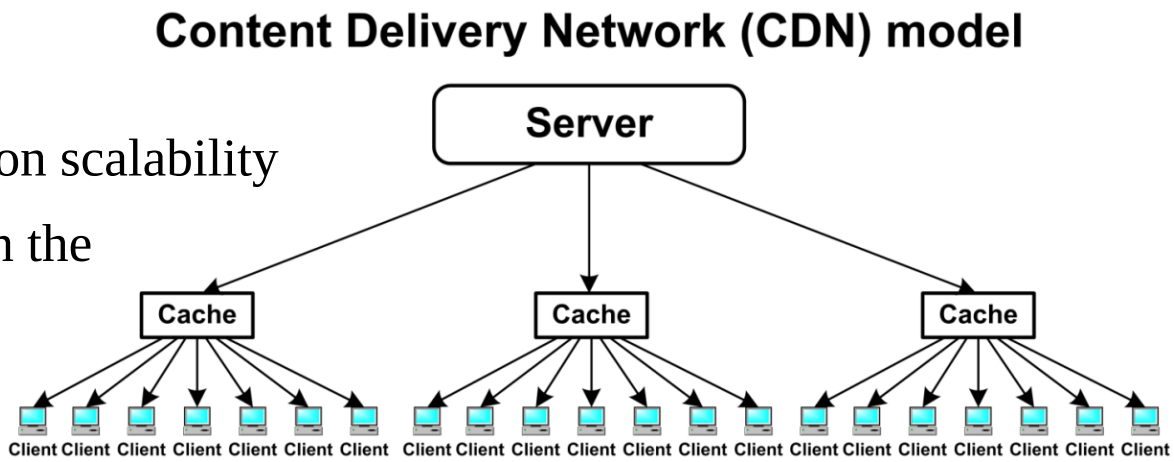

- Cons
  - Content provider must pay to use the service

The Information Technology & Innovation Foundation

ITIF

The history of network management

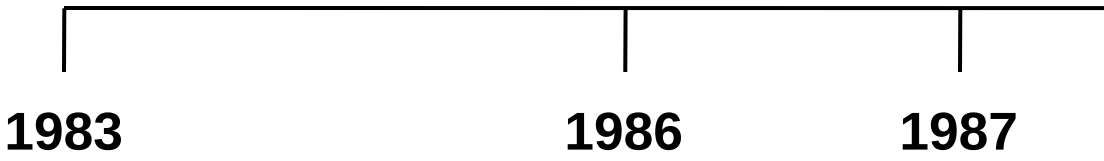

Birth of the Modern Internet

The Information Technology & Innovation Foundation

ITIF

The goal of network management

1. Fair and equitable bandwidth allocation
   - Customers who pay for the same service should have the same bandwidth for the same duration of time

2. Improve multi-tasking capability of network
   - Better simultaneous application usage
   - Minimize jitter at any bottleneck on the network

Shared bandwidth is good for consumers

- Why are networks even shared to begin with?
- Networks will always be shared somewhere
- Shared networks are faster and cheaper
  - Dedicated 1.5 Mbps T1 circuit is at least $180 per Mbps
  - Shared 6 Mbps broadband is $7 per Mbps

P2P bypasses Jacobson's algorithm

Fair

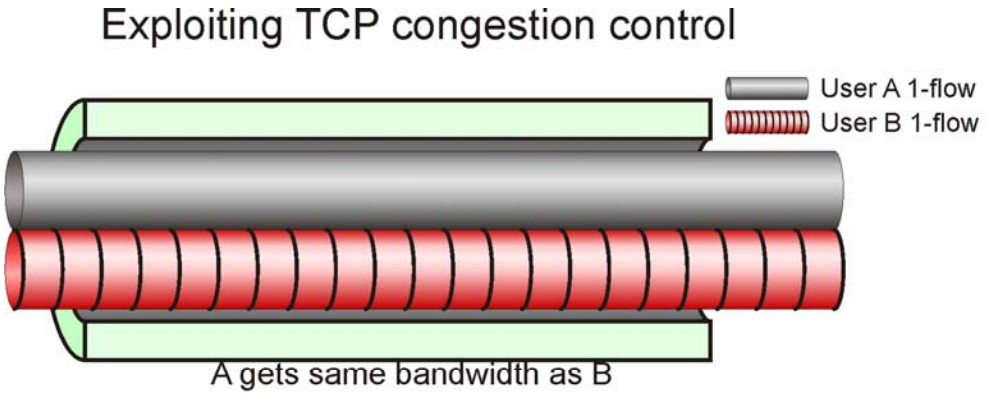

Unfair

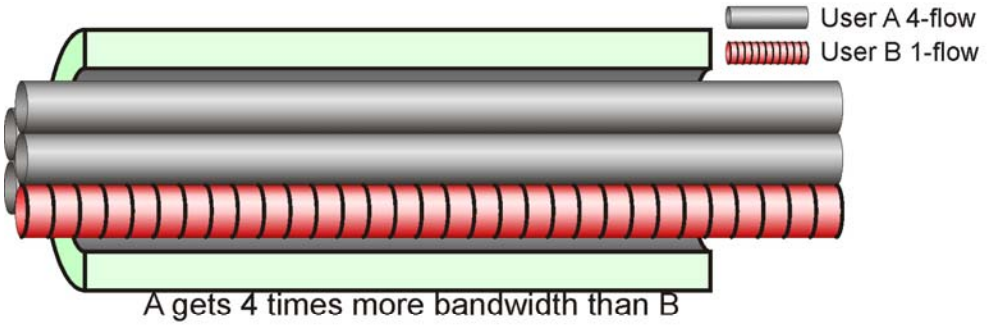

Really unfair

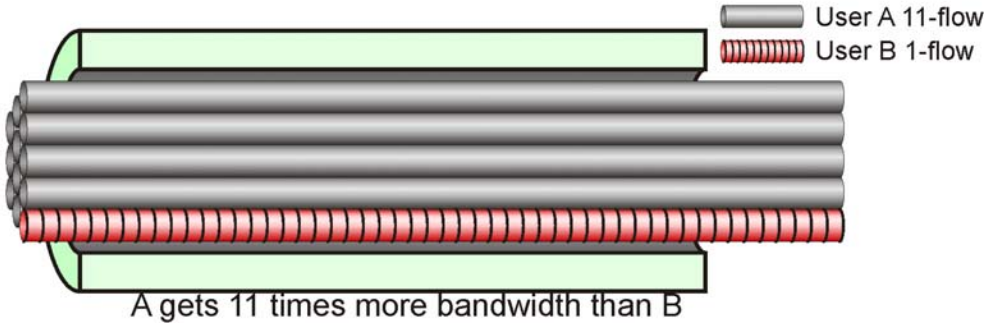

The Information Technology & Innovation Foundation

ITIF

Protocol agnostic solutions restore fairness

Unmanaged & unfair

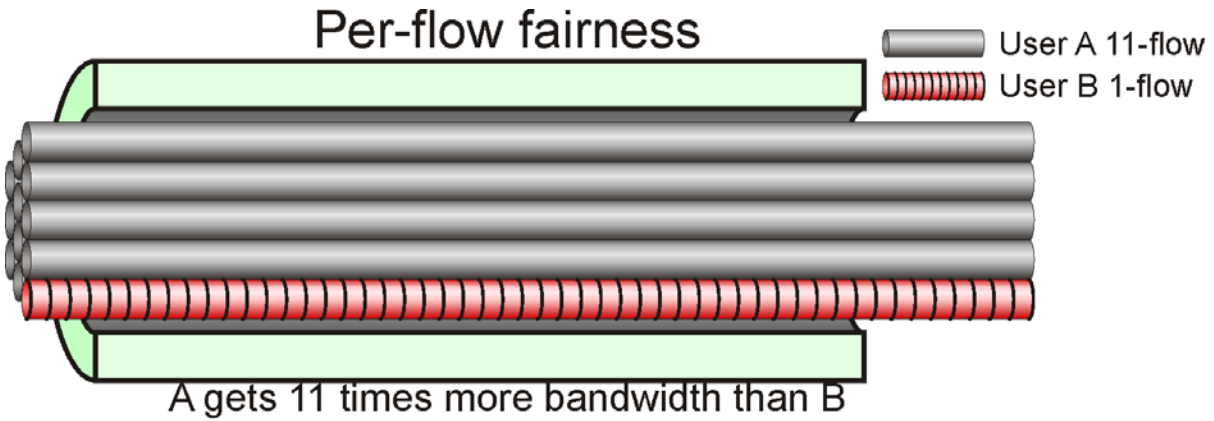

Managed & fair

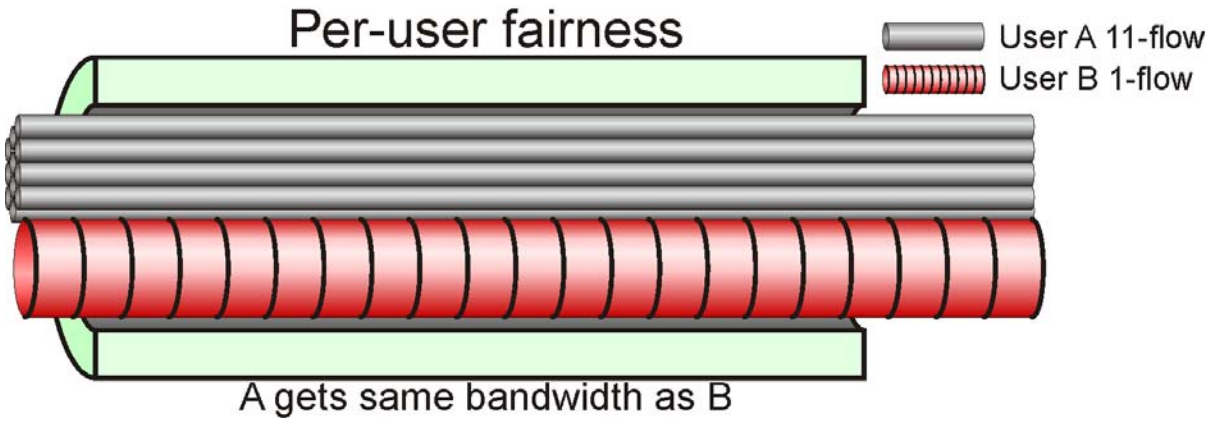

The Information Technology & Innovation Foundation

ITIF

The goal of network management

1. Fair and equitable bandwidth allocation
   - Customers who pay for the same service should have the same bandwidth for the same duration of time
2. Improve multi-tasking capability of network
   - Better simultaneous application usage
   - Minimize jitter at any bottleneck on the network

QoS and the Internet

- Quality of Service (QoS) fixes the inherently unstable bandwidth and packet delay of packet switching
- There are many Internet standards for QoS
  - Type of Service (ToS) standard in 1981
  - Integrated Services (IntServ) in 1994
  - Differentiated Services (DiffServ) in 1998
- QoS is a **very** broad term
  - Used in telephony, other data networks, and Internet Protocol (IP) networks
- Alternate names
  - "Enhanced QoS" or "Prioritization" or "Premium service"

Why do we need QoS?

- Internet Protocol (IP) networks are inherently bad at application multitasking
- Multiple computers, Internet enabled TVs or set top boxes with P2P capability will soon become common
  - "Honey, can you shut the TV download so I can make a phone call" will become more common at home
- Voice over IP (VoIP) and online gaming are extremely "allergic" to P2P without QoS

Three basic services of the Internet

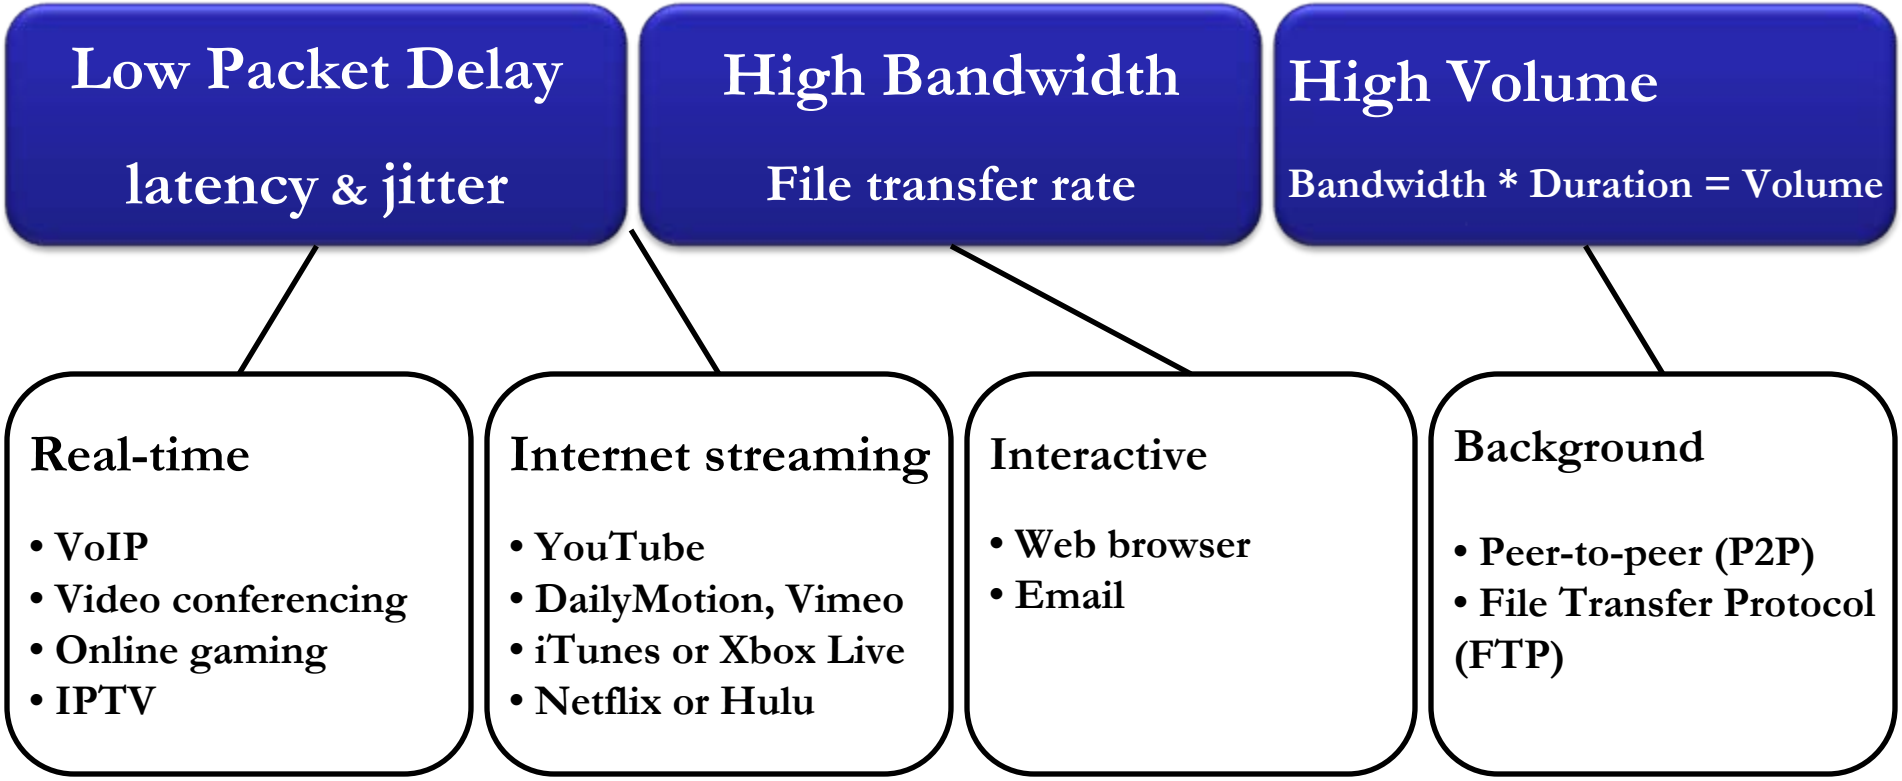

The logical order of packet priority

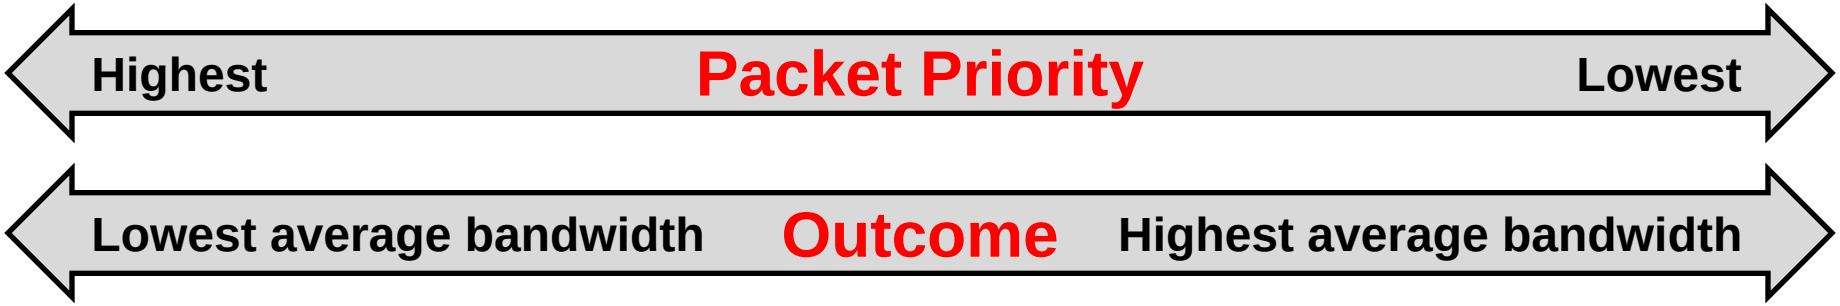

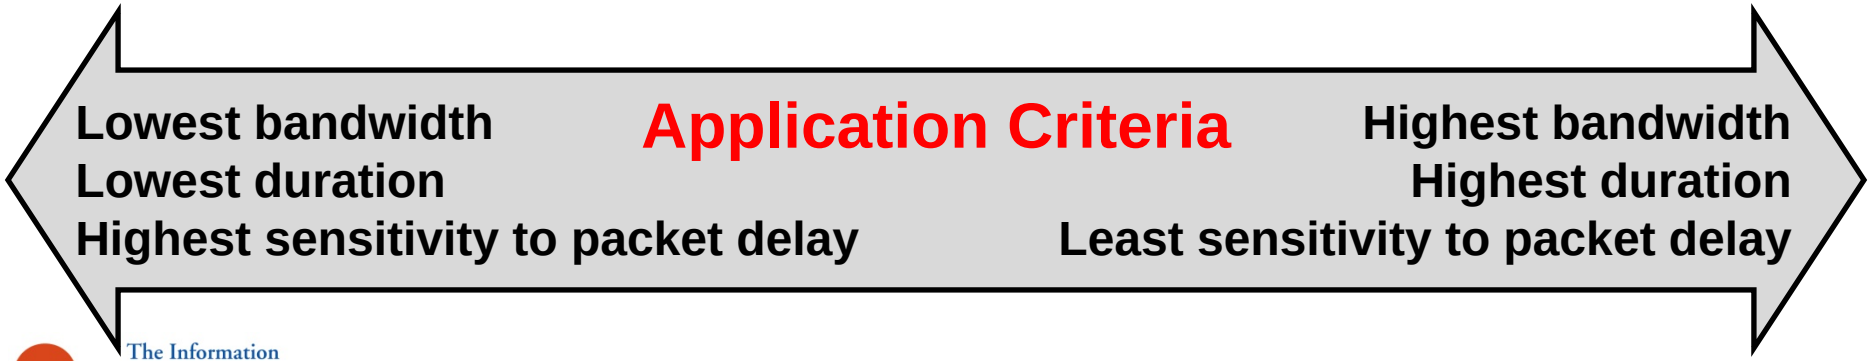

The Information Technology & Innovation Foundation

ITIF

Dumb networks multitask poorly

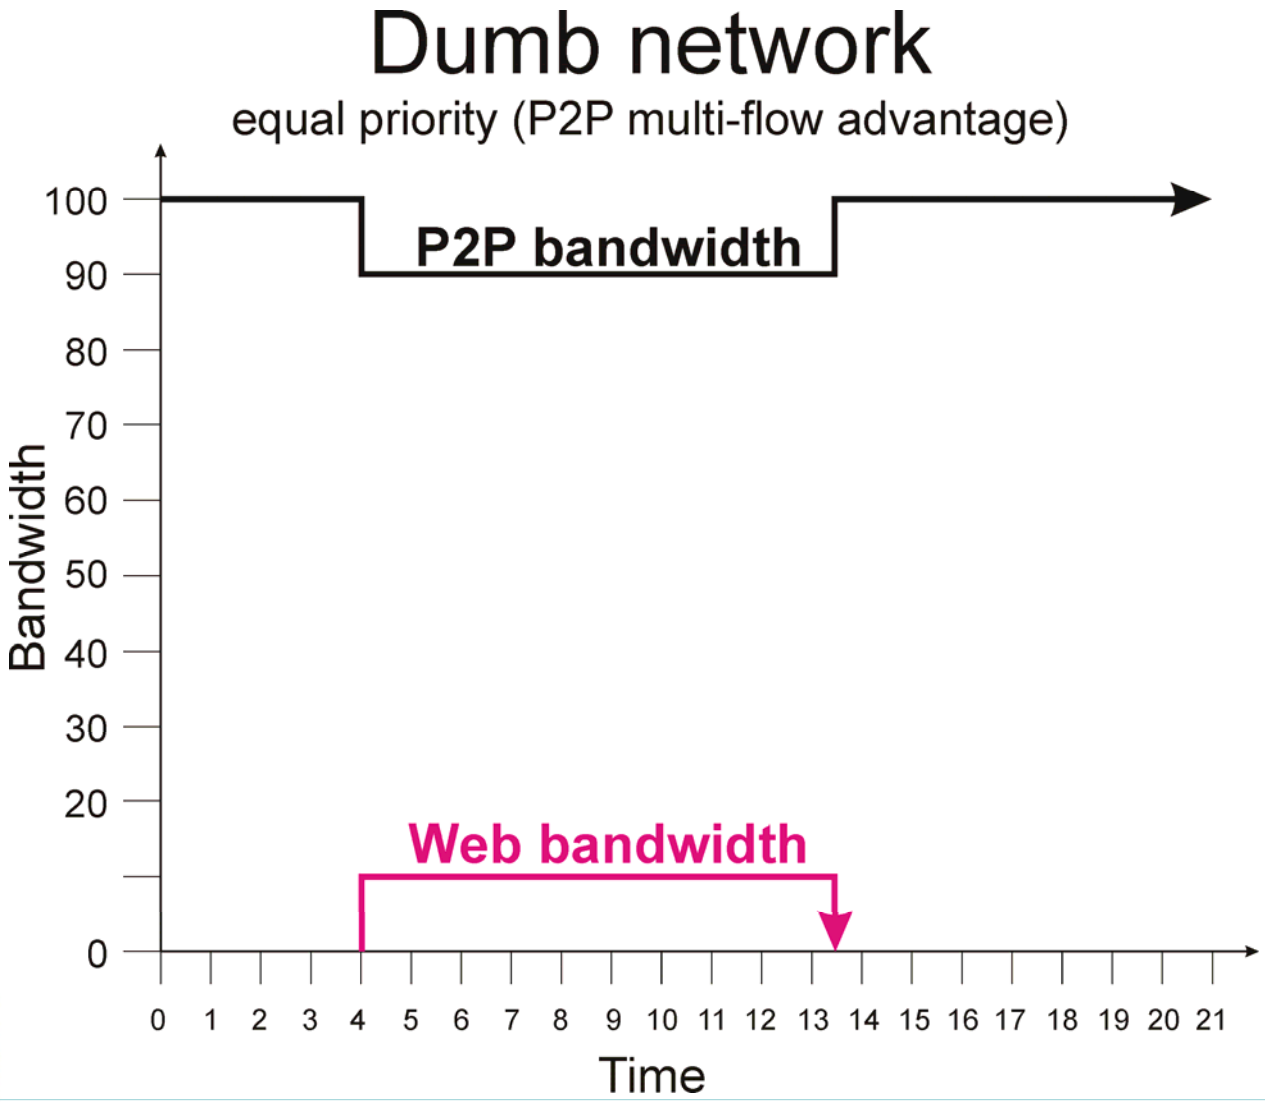

The Information Technology & Innovation Foundation

ITIF

Smart networks multitask better

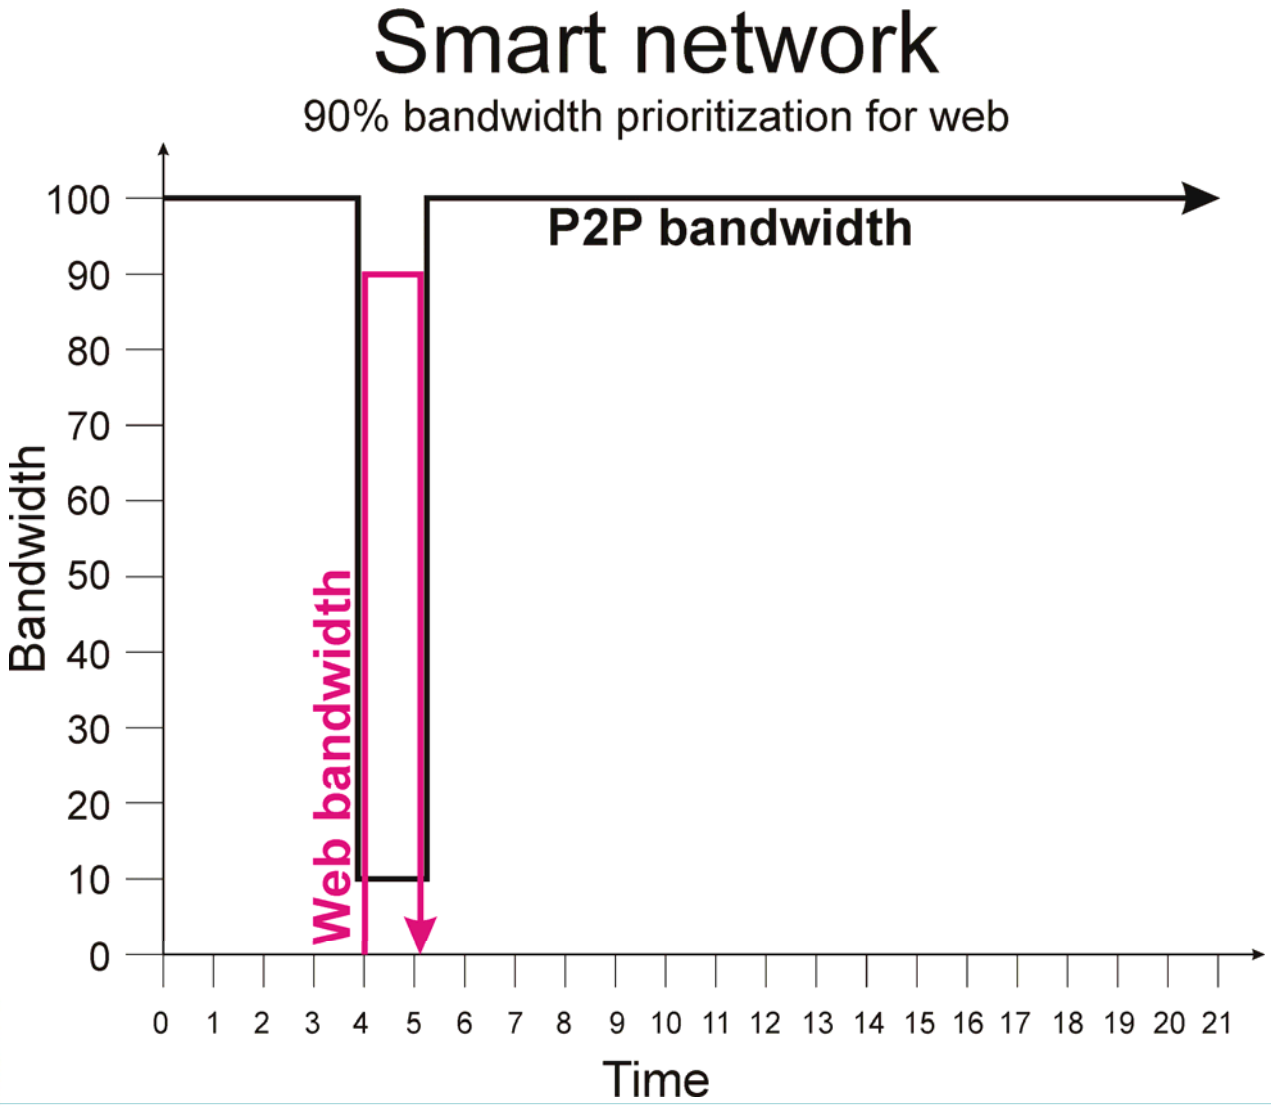

The Information Technology & Innovation Foundation

ITIF

The goal of network management

1. Fair and equitable bandwidth allocation
   - Customers who pay for the same service should have the same bandwidth for the same duration of time
2. Improve multi-tasking capability of network
   - Better simultaneous application usage
   - Minimize jitter at any bottleneck on the network

High jitter at low traffic levels

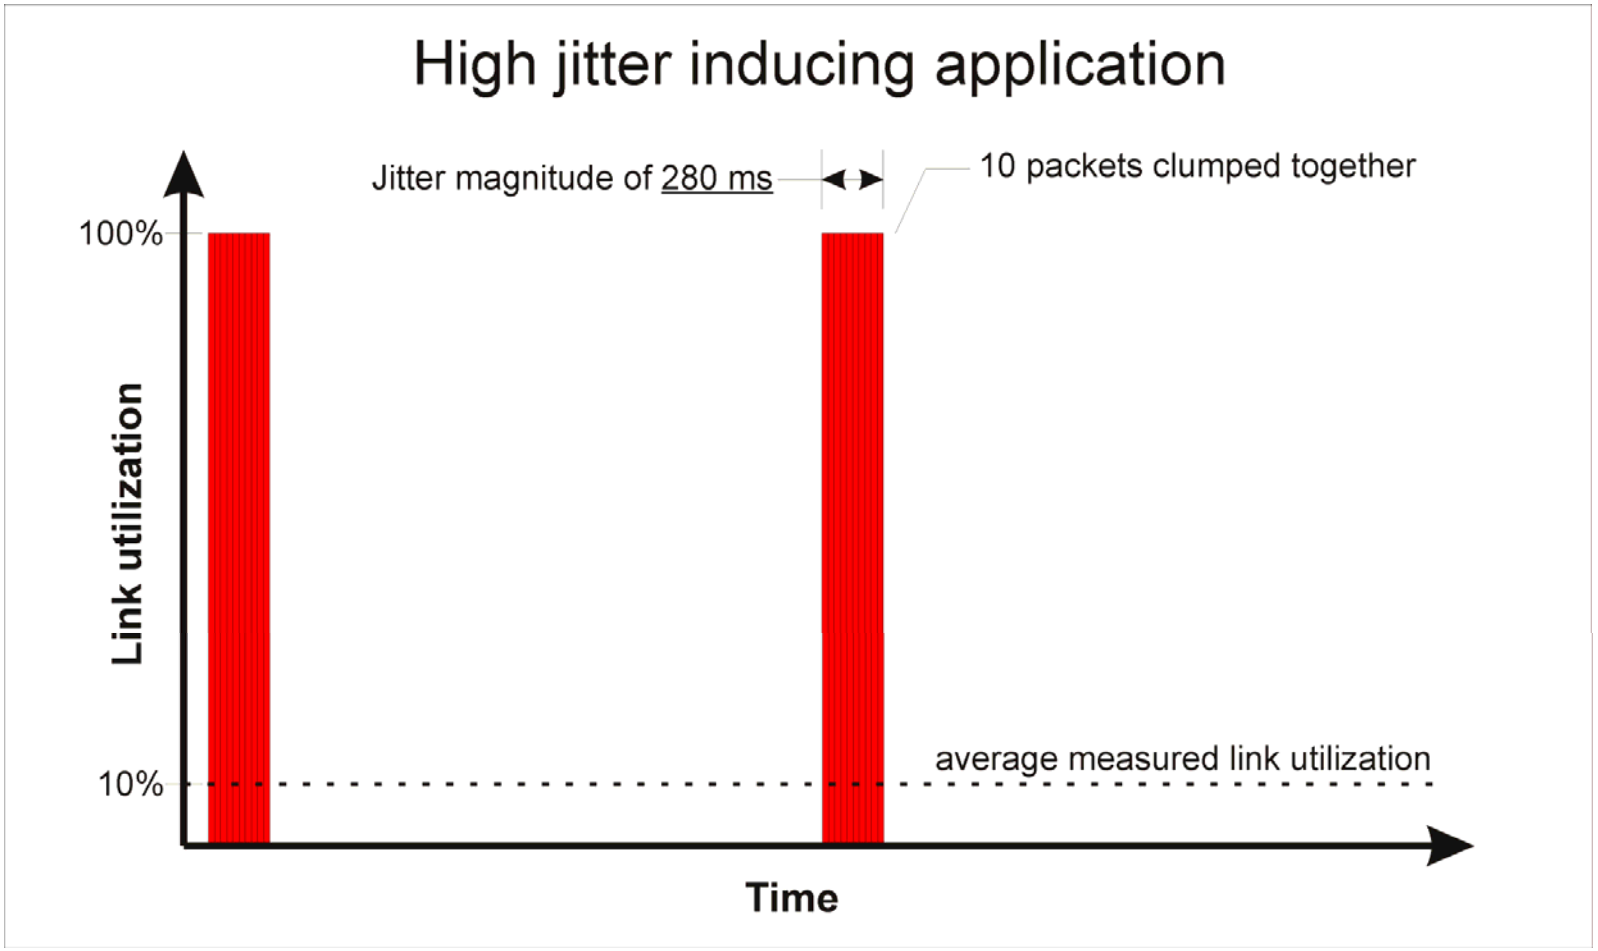

The Information Technology & Innovation Foundation

ITIF

Spaced out packets don't produce jitter

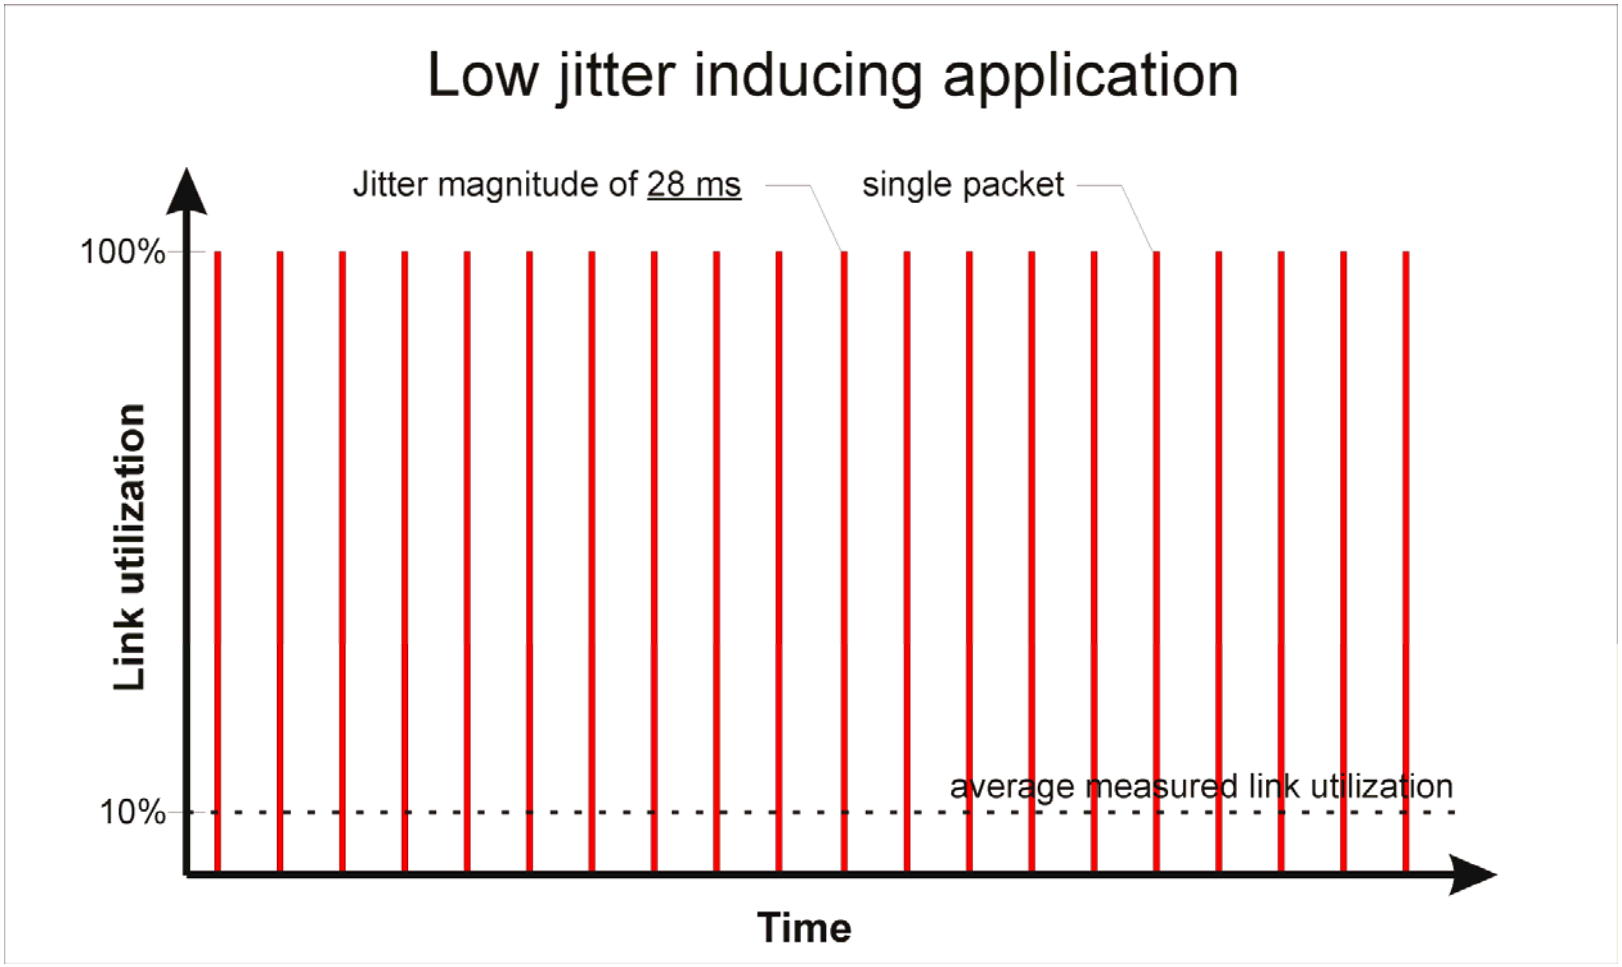

The Information Technology & Innovation Foundation

ITIF

P2P doesn't mix well with VoIP

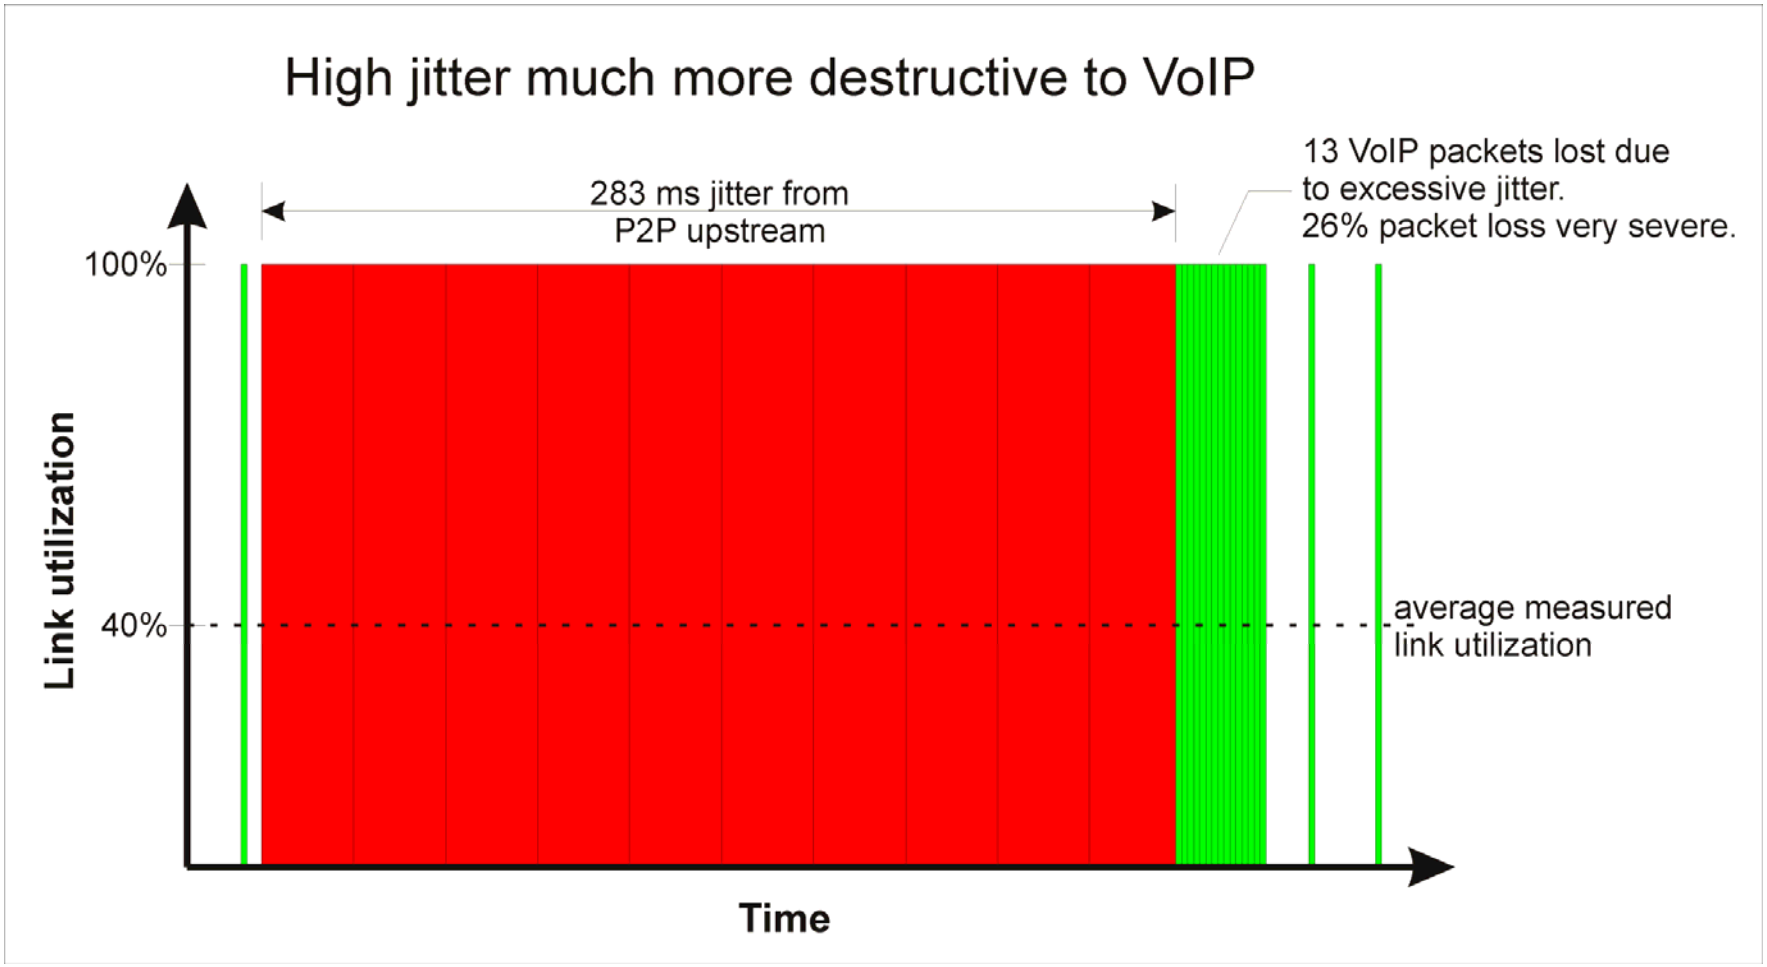

The Information Technology & Innovation Foundation

ITIF

P2P usage causes massive jitter

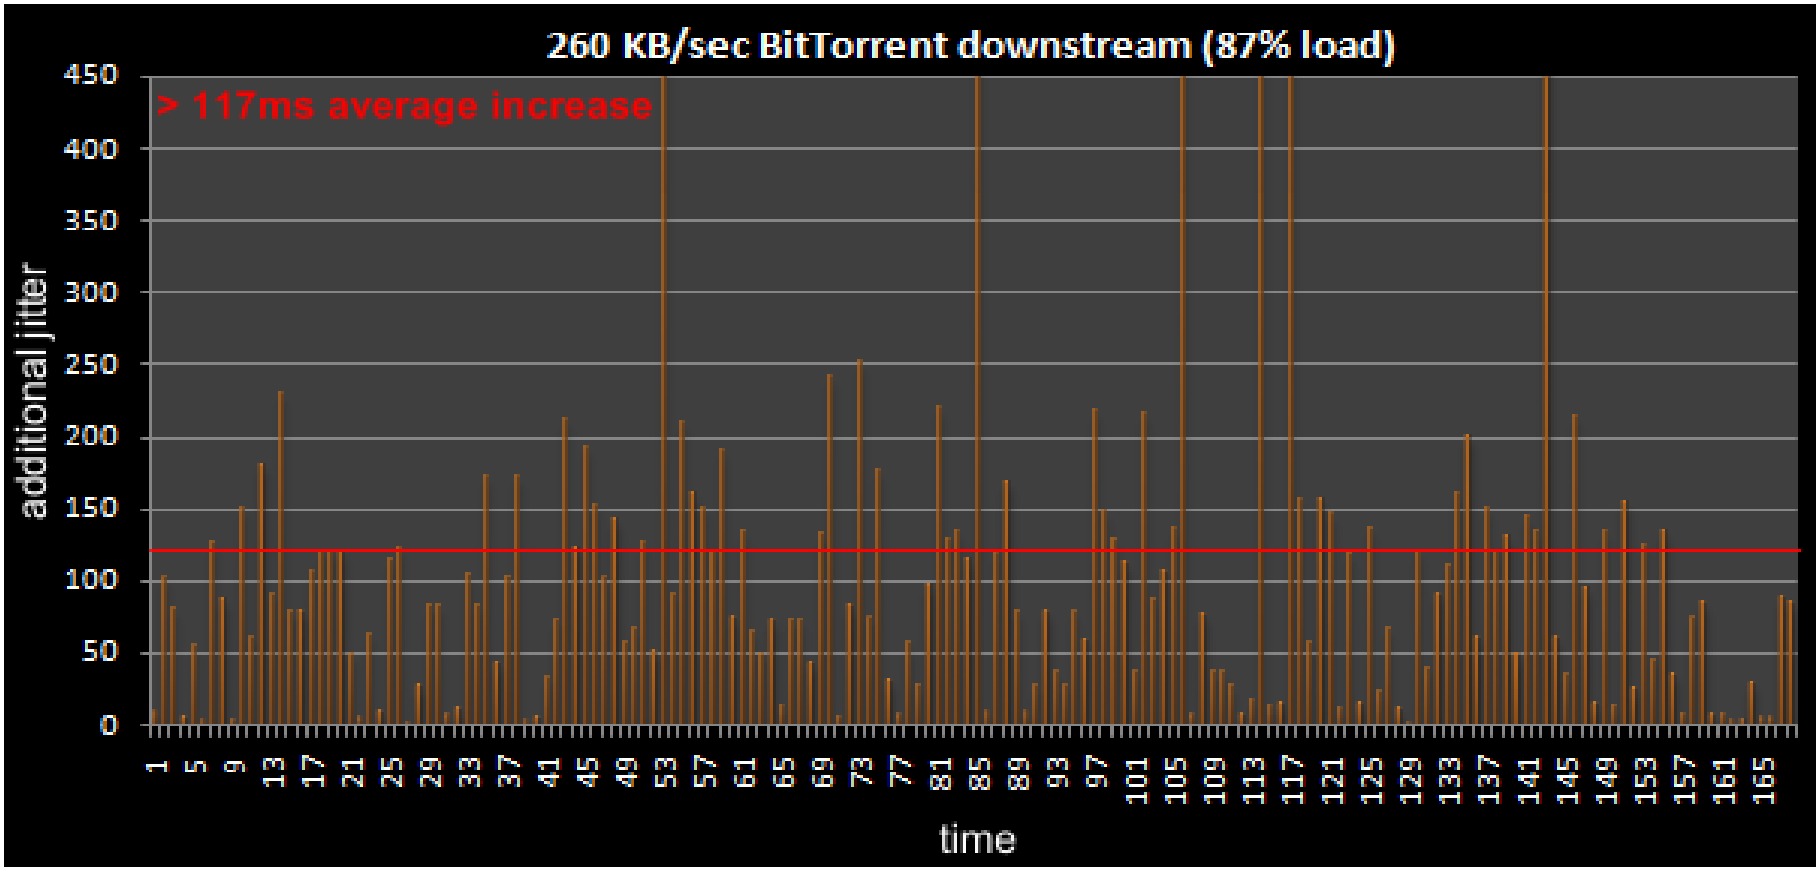

1000+ ms jitter above baseline from P2P usage

The Information Technology & Innovation Foundation

ITIF

Measured by George Ou on DSL broadband in May 2008

Even mild P2P usage causes jitter

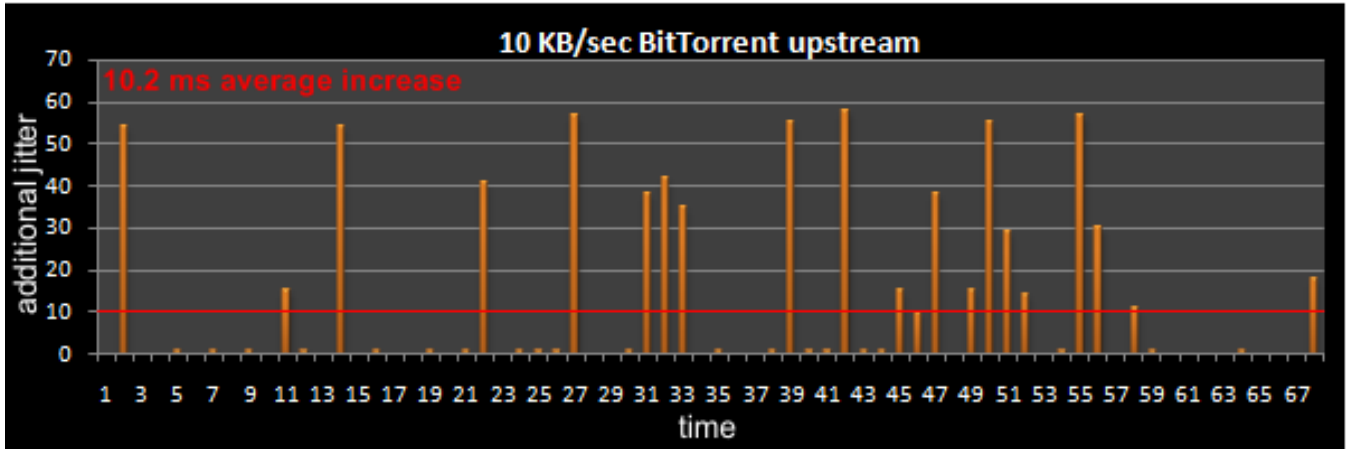

60 ms jitter
from just 10 KB/sec P2P usage

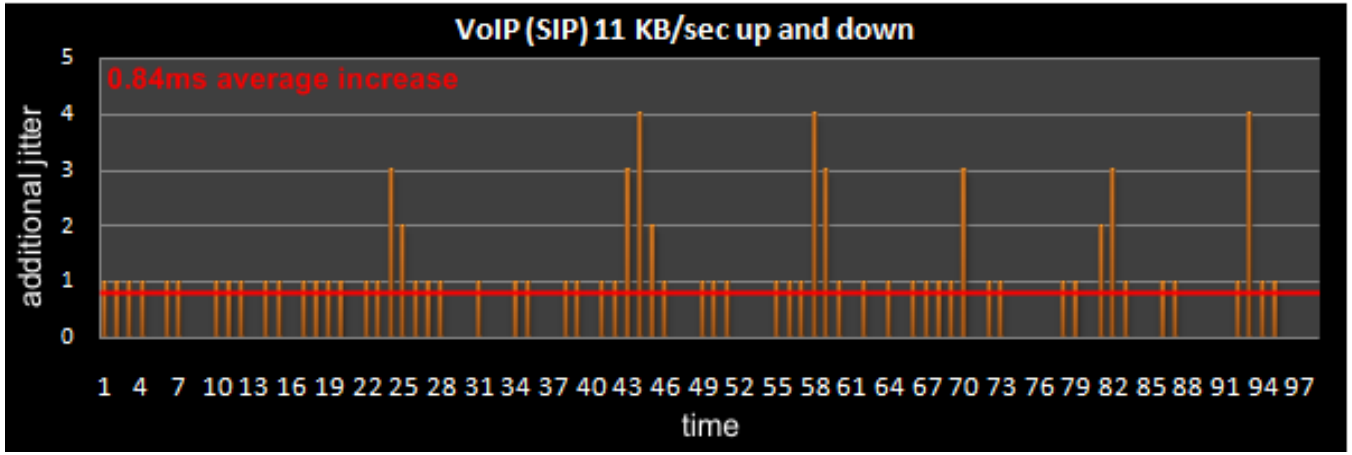

4 ms jitter
from 11 KB/sec VoIP usage

The Information Technology & Innovation Foundation
ITIF

Measured by George Ou on DSL broadband in May 2008

Mitigating jitter with QoS

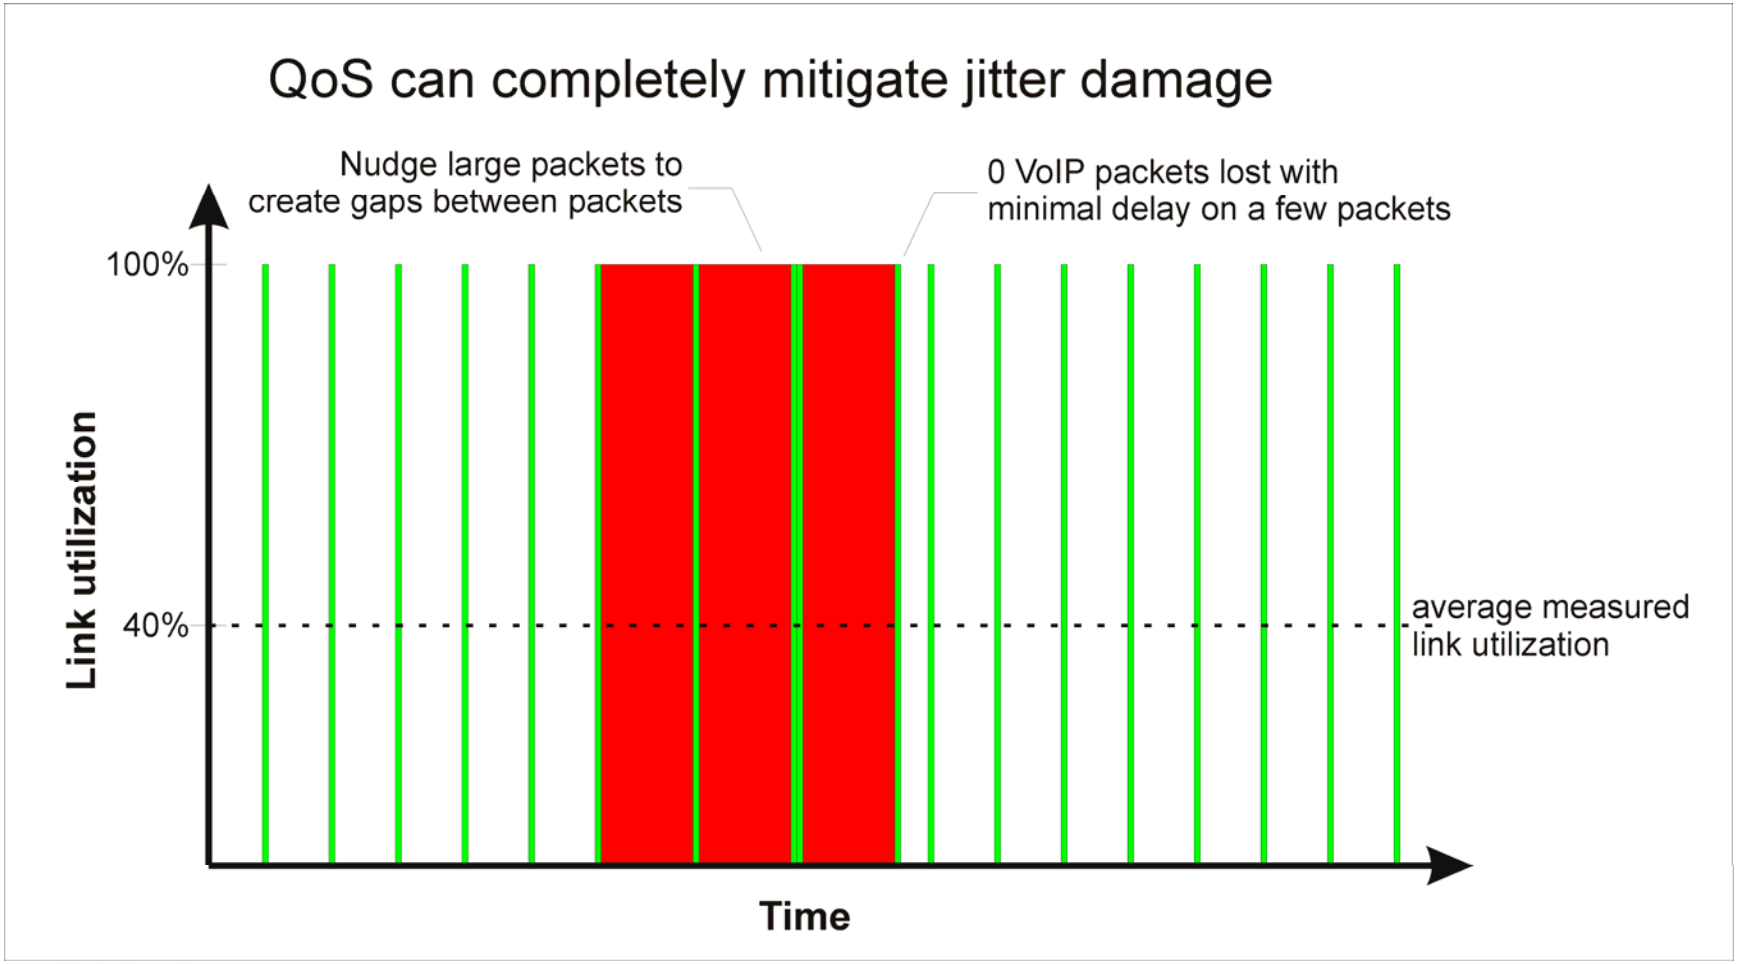

The Information Technology & Innovation Foundation

ITIF

Mitigating jitter with QoS

Dumb network

First In First Out (FIFO)

Click to view animation

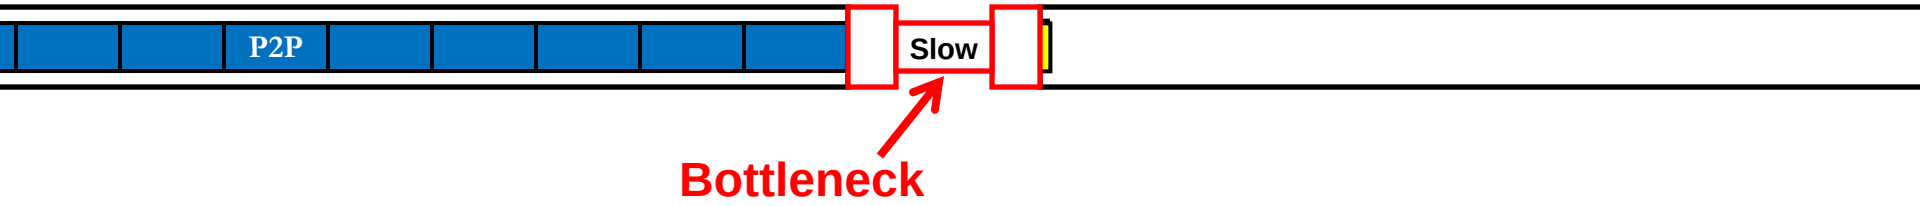

Bottleneck

Network with QoS

Multiple Queues or “onramps”

Click to view animation

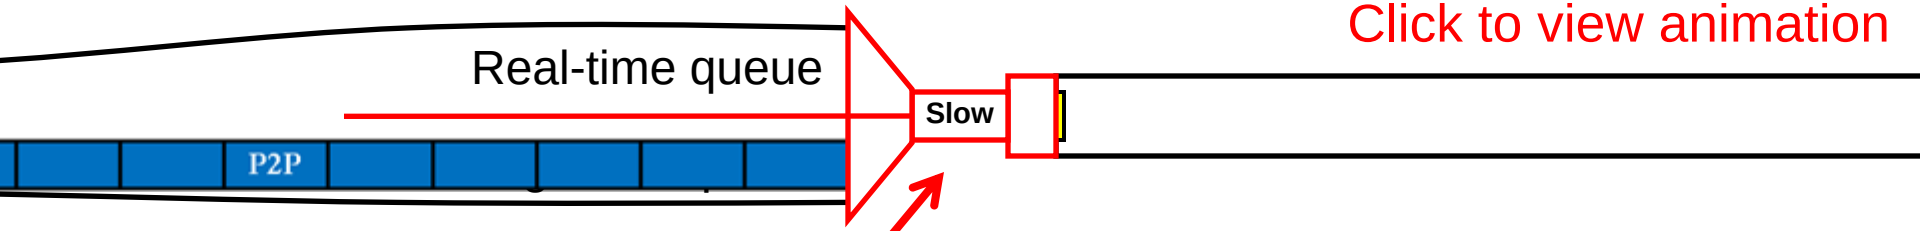

Bottleneck

Small VoIP packet (80 - 220 BYTES) VoIP

Large P2P packet (1472 BYTES) P2P

The Information Technology & Innovation Foundation

ITIF

QoS needed on both ends to fight jitter

Home modem must manage upstream jitter

ISP has to eliminate downstream jitter

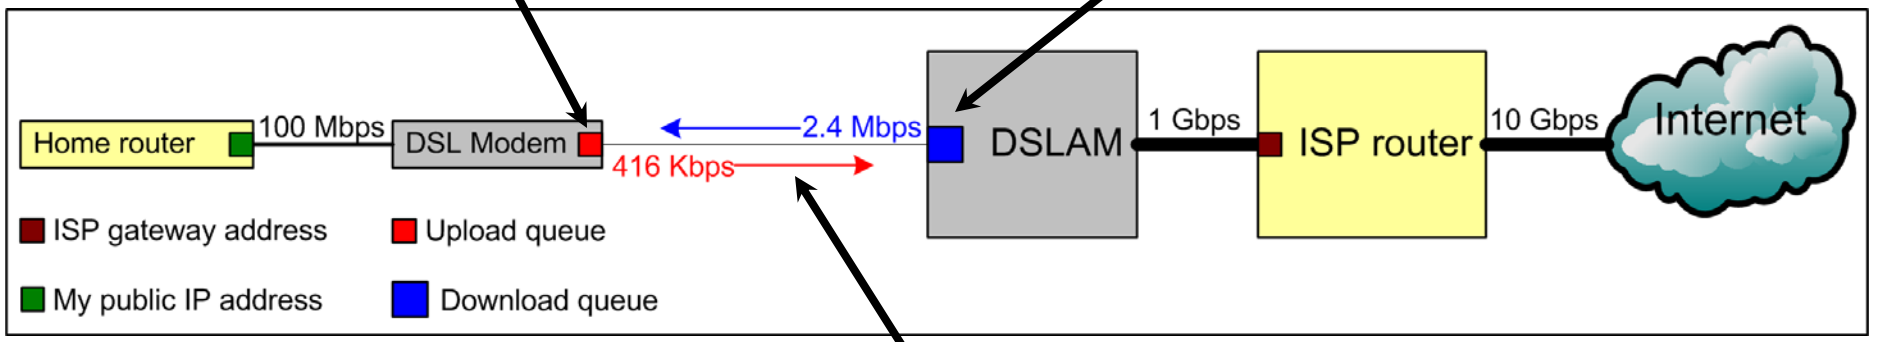

Will always remain a relative bottleneck

The Information Technology & Innovation Foundation

ITIF

Wireless – The new frontier of the Internet

The Information Technology & Innovation Foundation

ITIF

George Ou
December 11, 2008

Network management is critical for wireless

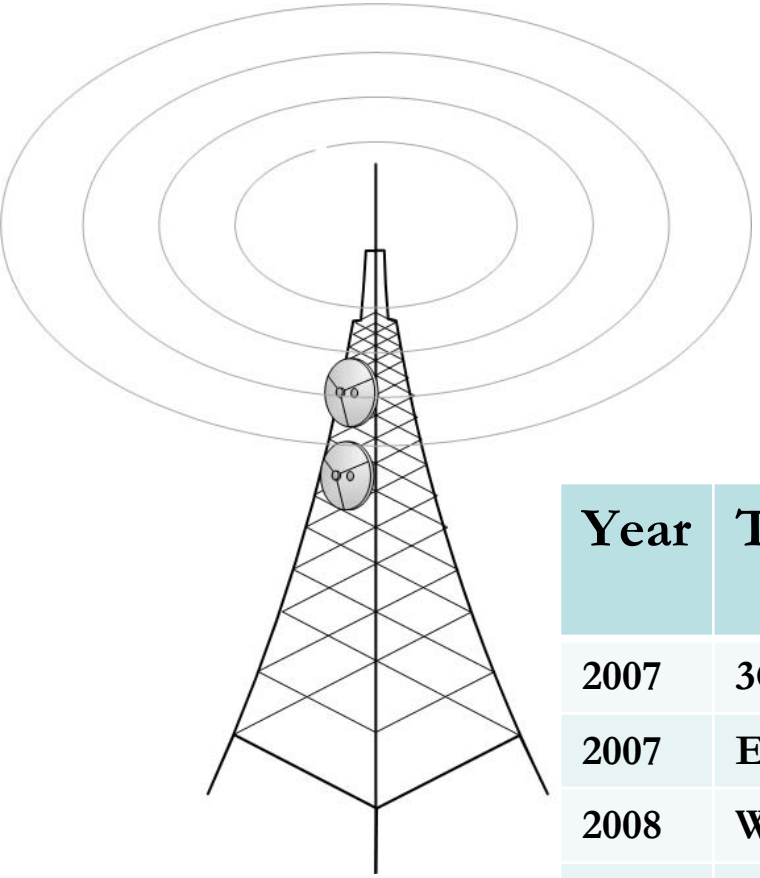

- Average of $650,000 per 3G cell tower
- One radio shared between 100 to 1000 people

Year	Technology	Bandwidth (mbps)		Latency (ms round trip)
		Up	Down	
2007	3GPP R5 – HSDPA	0.375	14.4	150
2007	EVDO Rev A (5 MHz)	7.2	12.4	100
2008	WiMAX (10 MHz)	8	40	60
2009	3GPP R7 – HSPA+	22	42	90
2010	LTE (20 MHz 2xMIMO)	50	150+	20

The Information Technology & Innovation Foundation

ITIF

Managed versus unmanaged wireless

- Wi-Fi 802.11b unscheduled access
  - 20 MHz of spectrum per radio
  - 70 simultaneous VoIP calls in theory;
    4 simultaneous VoIP in practice
  - 5th phone on network causes all 5 VoIP phones to suffer breakup
  - Unscheduled packets colliding randomly are the culprit
- LTE scheduled access
  - 200 active users per radio using 5 MHz of spectrum
  - 200 times more users per MHz
  - 200 times better spectral efficiency than dumb Wi-Fi

Common misconceptions about QoS

- ***"QoS violates the end-to-end architecture of the Internet"***
  - QoS is an Internet standard
  - End-to-end never mandated a dumb Internet
- ***"Capacity is a cheaper substitute for QoS"***
  - Capacity is never cheap enough
  - Jitter can occur on "fat pipes" with very little traffic

Common misconceptions about QoS

- ***"QoS doesn't work on the Internet"***
  - Based on misconception that QoS must work on every leg of Internet (across multiple network providers) to be useful
  - Reality is that QoS is useful especially for broadband
- ***"Internet2 concluded that QoS isn't necessary"***
  - Based on a paper by Shalunov and Teitelbaum
    - Admitted QoS was even necessary on high capacity Internet2
    - Wrongly concluded that capacity is cheaper than QoS
    - Admitted QoS works well when targeted at congested links

Common misconceptions about QoS

"Capacity is cheaper than QoS"

- Japan's 100 Mbps fiber broadband network is often congested
- 10% users using P2P accounted for 65% to 90% of all traffic
- ISPs implemented 30 GB daily upstream caps
- Implemented 3 warnings for piracy before account termination

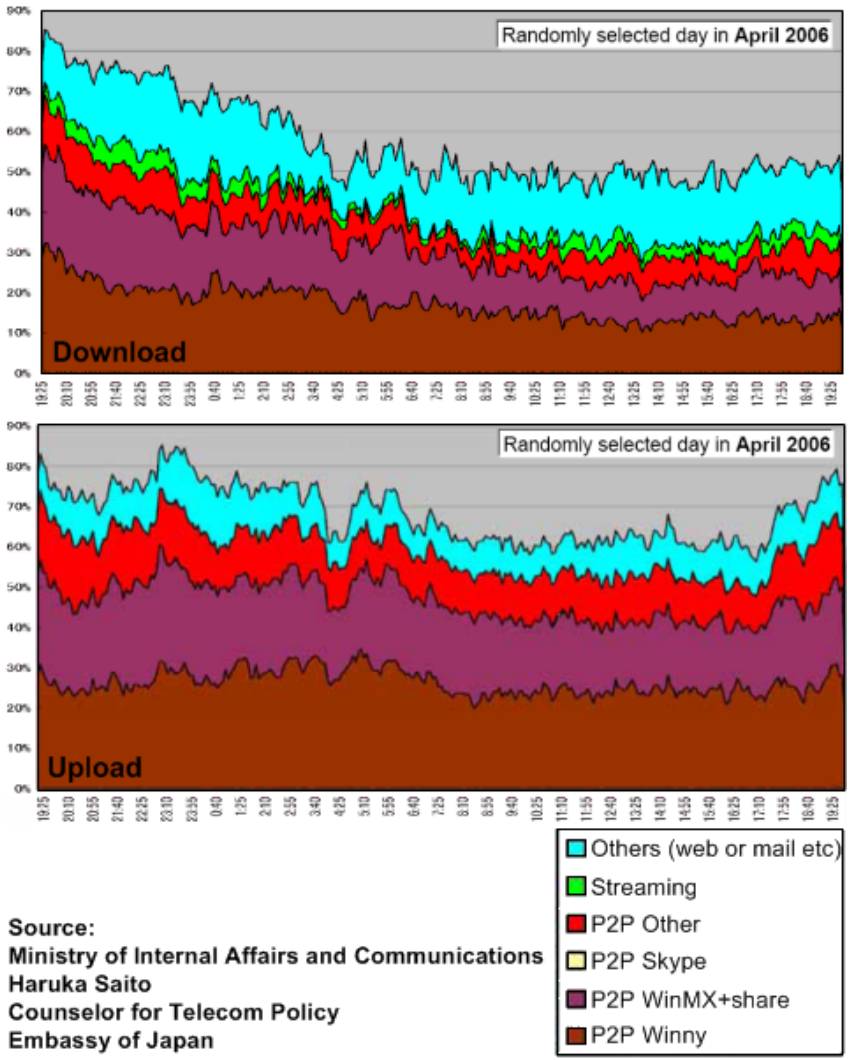

Source:
Ministry of Internal Affairs and Communications
Haruka Saito
Counselor for Telecom Policy
Embassy of Japan

The Information Technology & Innovation Foundation

ITIF

Regulatory Implications of Net Neutrality

- Bill that attempted to ban prioritization
  - S.2360 - Internet Non-Discrimination Act of 2006 – Wyden (D-Oregon)
- What this bill mandates
  - Prohibit broadband providers from prioritizing bandwidth and allocating bandwidth
- Implications
  - Reduces the quality and utility of broadband
  - May force more use of private circuits for IPTV resulting in less bandwidth for the Internet

Regulatory Implications of Net Neutrality

- Bills that ban multi-tiered Quality of Service (QoS)
  - H.R. 5273 - Net Neutrality Act of 2006 - Markey (D-MA)
  - S. 215 - Internet Freedom and preservation act of 2007
    Snowe (R-ME) and Dorgan (D-ND)
  - H.R. 5417 - Internet Freedom and Nondiscrimination Act of 2006
    Sensenbrenner (R-WI) and Conyers (D-MI). (Reintroduced in 2008)
- What these bills mandate
  - Would prohibit broadband providers from charging for "enhanced QoS"
  - Permits traffic type prioritization but not based on traffic source
- Implications
  - Effectively mandates equal service for unequal payment
  - May force more use of private circuits resulting in less bandwidth for Internet

Why exclusive QoS is necessary

Exclusive QoS

0-64%

36-100%

DSL/IPTV
head-end unit
Fiber to the Node (FTTN)
DSLAM

Ban exclusive QoS

64%

36%

25 Mbps total

FTTN

Priority IPTV data
(0-16 Mbps)

VDSL2
modem & router

Internet data
0-25 Mbps

George Ou
December 11, 2008

Policy implications

- ISPs and application providers need to be more transparent
  - Some ISPs are advertising “unlimited” service it isn’t unlimited
  - Some ISPs aren’t disclosing usage caps
  - Some ISPs aren’t explaining minimal bandwidth clear enough
  - Consumers don’t always understand that costs from some applications are offloaded to them
- Government oversight
  - FCC should ensure broadband providers don’t abuse power
  - Industry-wide standard on transparency and disclosure to create a level playing field

Conclusion

- The Internet is so valuable because it is open to anyone, any use, and any business model
  - But participation always required varying levels of payment for varying levels of service between willing parties
- We always need more capacity (and policies to spur more capacity), but more capacity isn't a substitute for network management
- Network management results in higher performance for everyone at lower prices

Website and contact information

- **George Ou**
- **www.ITIF.org**
- **gou@itif.org**

The Information Technology & Innovation Foundation

ITIF