25 years of quality of service research - where next?

Henning Schulzrinne
Columbia University
(with Omer Boyaci, Andrea Forte, Kyung-Hwa Kim)

Prologue

* Most keynotes are prospective – this one is (partially) retrospective and introspective
* Foil for reflection
  * applies just as well to P2P, mobility, multicast, sensor networks, social networks, …
    * but they are still (too) active to reflect
* How effective is our collective research?
* How do we choose and solve problems?
* When do we move on?

Preview

* What can we learn from 25+ years of QoS research?
* Some of my group's (semi-) QoS research
  * how good is industrial practice?
  * how can we diagnose QoS (and other problems) in the consumer Internet?
* Thoughts on QoS going forward

About (networking) research

My assumptions

* We're an engineering discipline
  * "Engineering is the discipline, art and profession of acquiring and applying technical, scientific, and mathematical knowledge to design and implement materials, structures, machines, devices, systems, and processes that safely realize a desired objective or invention."
* Other (good) possibilities:
  * we train future engineers
  * we train future researchers

Pasteur's quadrant

Quest for Fundamental Understanding?

Quest for Fundamental Understanding? / Considerations of Use?	No	Yes
Yes	Pure basic research (Bohr)	**Use-inspired basic research (Pasteur)**
No	*Guessing at problems (Infocom)*	**Pure applied research (Edison)**

Considerations of Use?

Most networking research wants to be here

Most networking research is here

Pasteur's Quadrant: Basic Science and Technological Innovation, Stokes 1997 (modified)

The $1B question

- How big a problem does your proposal solve?
  - Does it create new ones?
    - financial, management, …
- Can it be integrated into the existing Internet
  - or a plausible successor?
  - or 802.11, 802.16, …
- … without everybody changing their ways
  - the secret: **nobody** is in charge of the Internet
- Can it be understood by Cisco CNAs?
  - see IP multicast, PIM-SM

Useful research outcomes

* Standards
  * unfortunately, rarely cite papers
* Get Cisco, Google, Microsoft, … to adopt it
  * 3-4 QoS papers?
* Show what doesn't work
  * counteract industry shills
  * e.g., recently web site privacy
* Understand the Internet better
  * but not just your campus network
* Prior art in patent disputes
  * patents don't have a 90% rejection rate…

CS research to reality

CS as science

CS as engineering

~~CS as a soccer league~~

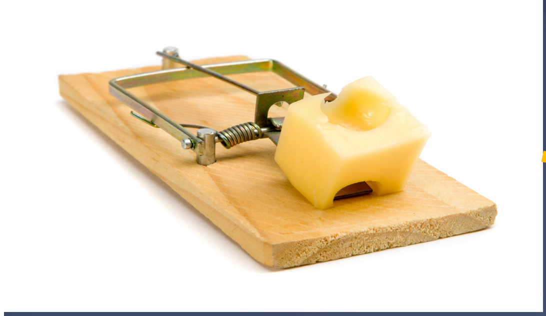

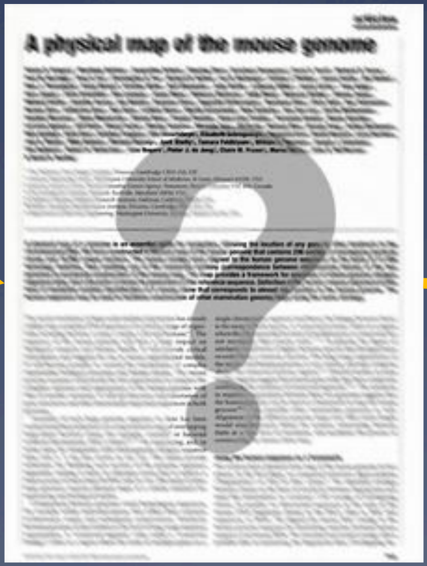

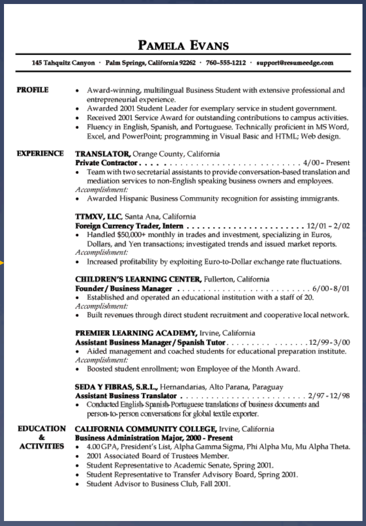

Pamela Evans

145 Tahquitz Canyon · Palm Springs, California 92262 · 760–555-1212 · support@resumeedge.com

PROFILE

- Award-winning, multilingual Business Student with extensive professional and entrepreneurial experience.
- Awarded 2001 Student Leader for exemplary service in student government.
- Received 2001 Service Award for outstanding contributions to campus activities.
- Fluency in English, Spanish, and Portuguese. Technically proficient in MS Word, Excel, and PowerPoint; programming in Visual Basic and HTML; Web design.

EXPERIENCE

TRANSLATOR, Orange County, California
Private Contractor . . . . 4/00 – Present
- Team with two secretarial assistants to provide conversation-based translation and mediation services to non-English speaking business owners and employees.

Accomplishment:
- Awarded Hispanic Business Community recognition for assisting immigrants.

TTMXV, LLC, Santa Ana, California
Foreign Currency Trader, Intern . . . . 12/01 – 2/02
- Handled $50,000+ monthly in trades and investment, specializing in Euros, Dollars, and Yen transactions; investigated trends and issued market reports.

Accomplishment:
- Increased profitability by exploiting Euro-to-Dollar exchange rate fluctuations.

CHILDREN'S LEARNING CENTER, Fullerton, California
Founder / Business Manager . . . . 6/00 - 8/01
- Established and operated an educational institution with a staff of 20.

Accomplishment:
- Built revenues through direct student recruitment and cooperative local network.

PREMIER LEARNING ACADEMY, Irvine, California
Assistant Business Manager / Spanish Tutor . . . . 12/99 - 3/00
- Aided management and coached students for educational preparation institute.

Accomplishment:
- Boosted student enrollment; won Employee of the Month Award.

SEDA Y FIBRAS, S.R.L., Hernandarias, Alto Parana, Paraguay
Assistant Business Translator . . . . 2/97 - 12/98
- Conducted English-Spanish-Portuguese translations of business documents and person-to-person conversations for global textile exporter.

EDUCATION & ACTIVITIES

CALIFORNIA COMMUNITY COLLEGE, Irvine, California
Business Administration Major, 2000 - Present
- 4.00 GPA, President's List, Alpha Gamma Sigma, Phi Alpha Mu, Mu Alpha Theta.
- 2001 Associated Board of Trustees Member.
- Student Representative to Academic Senate, Spring 2001.
- Student Representative to Transfer Advisory Board, Spring 2001.
- Student Advisor to Business Club, Fall 2001.

Network tech transfer, mode 1

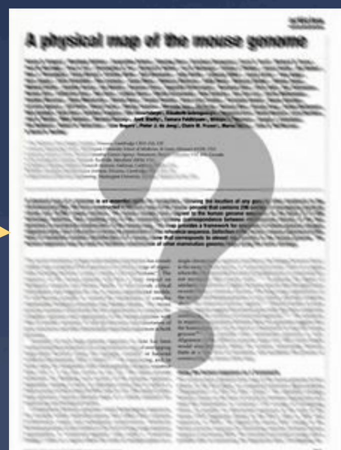
A physical map of the mouse genome

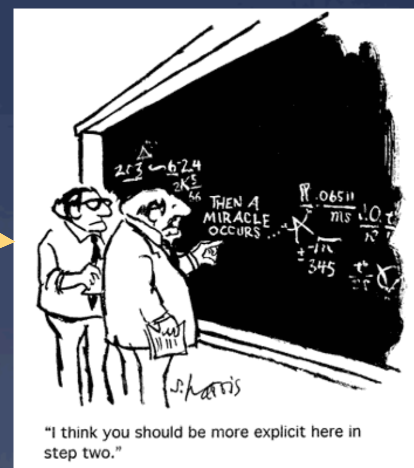

somebody else just waiting for your results

Network tech transfer, mode 2

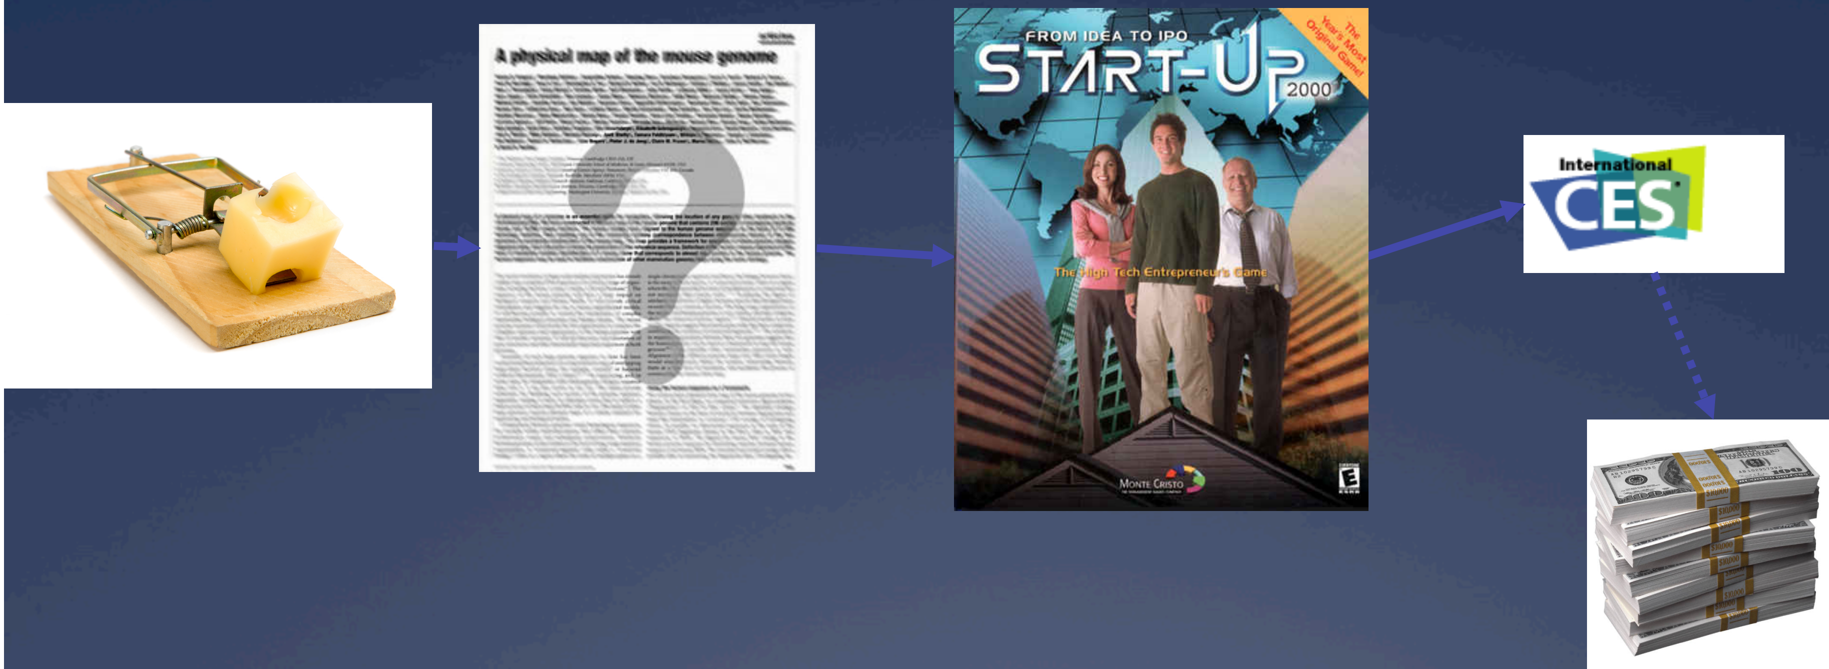

Or just measure citations

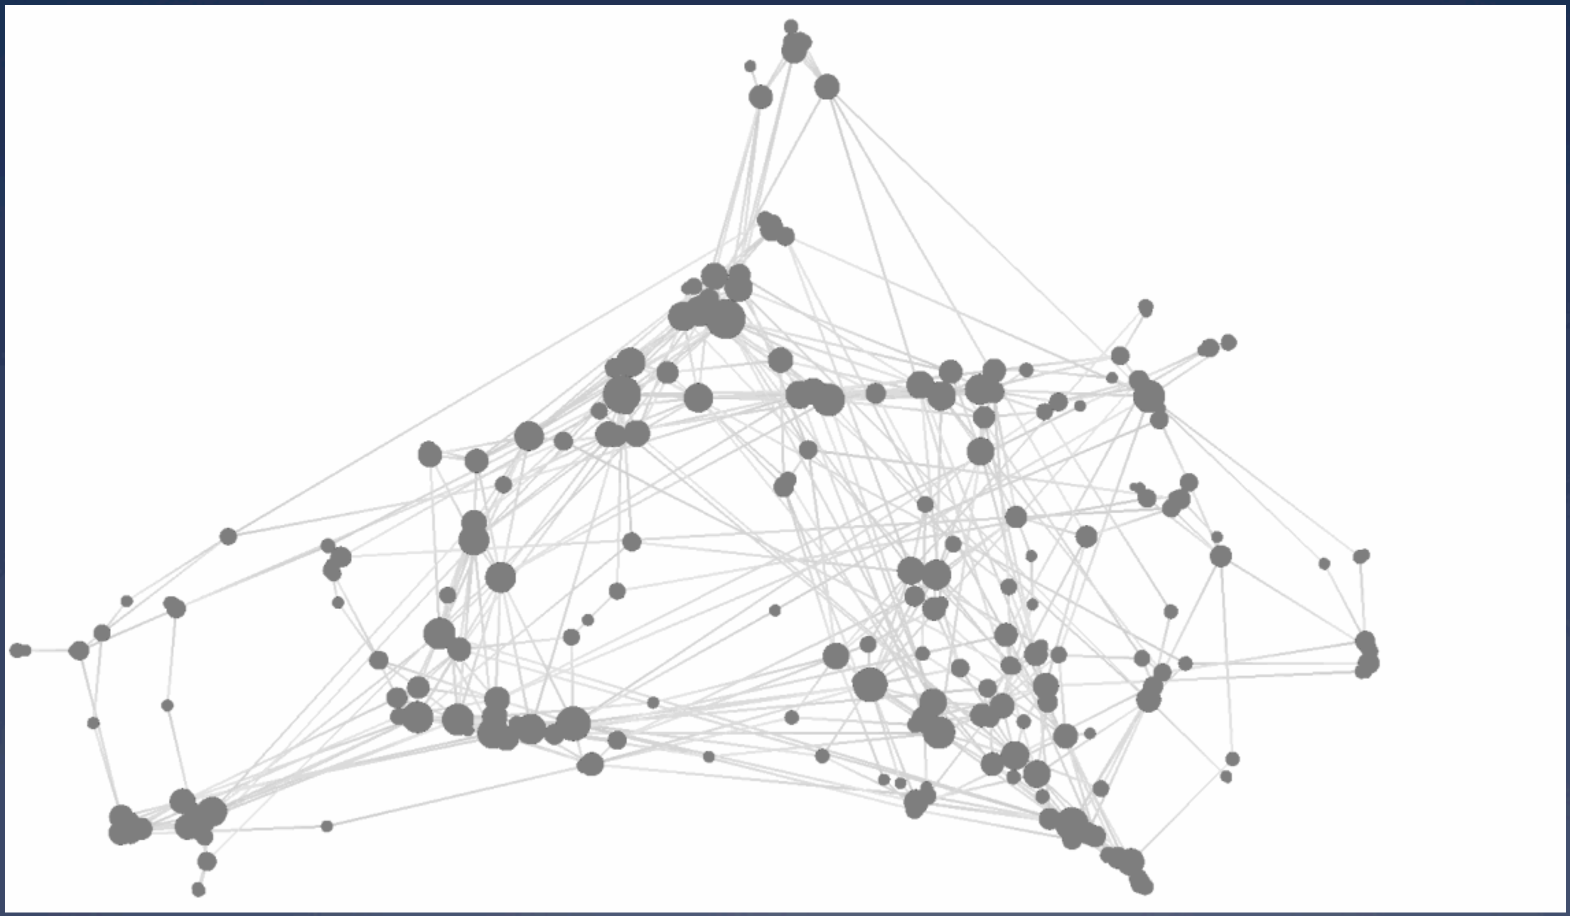

be sure to create enough conferences and workshops…

QoS research

Why is QoS attractive?

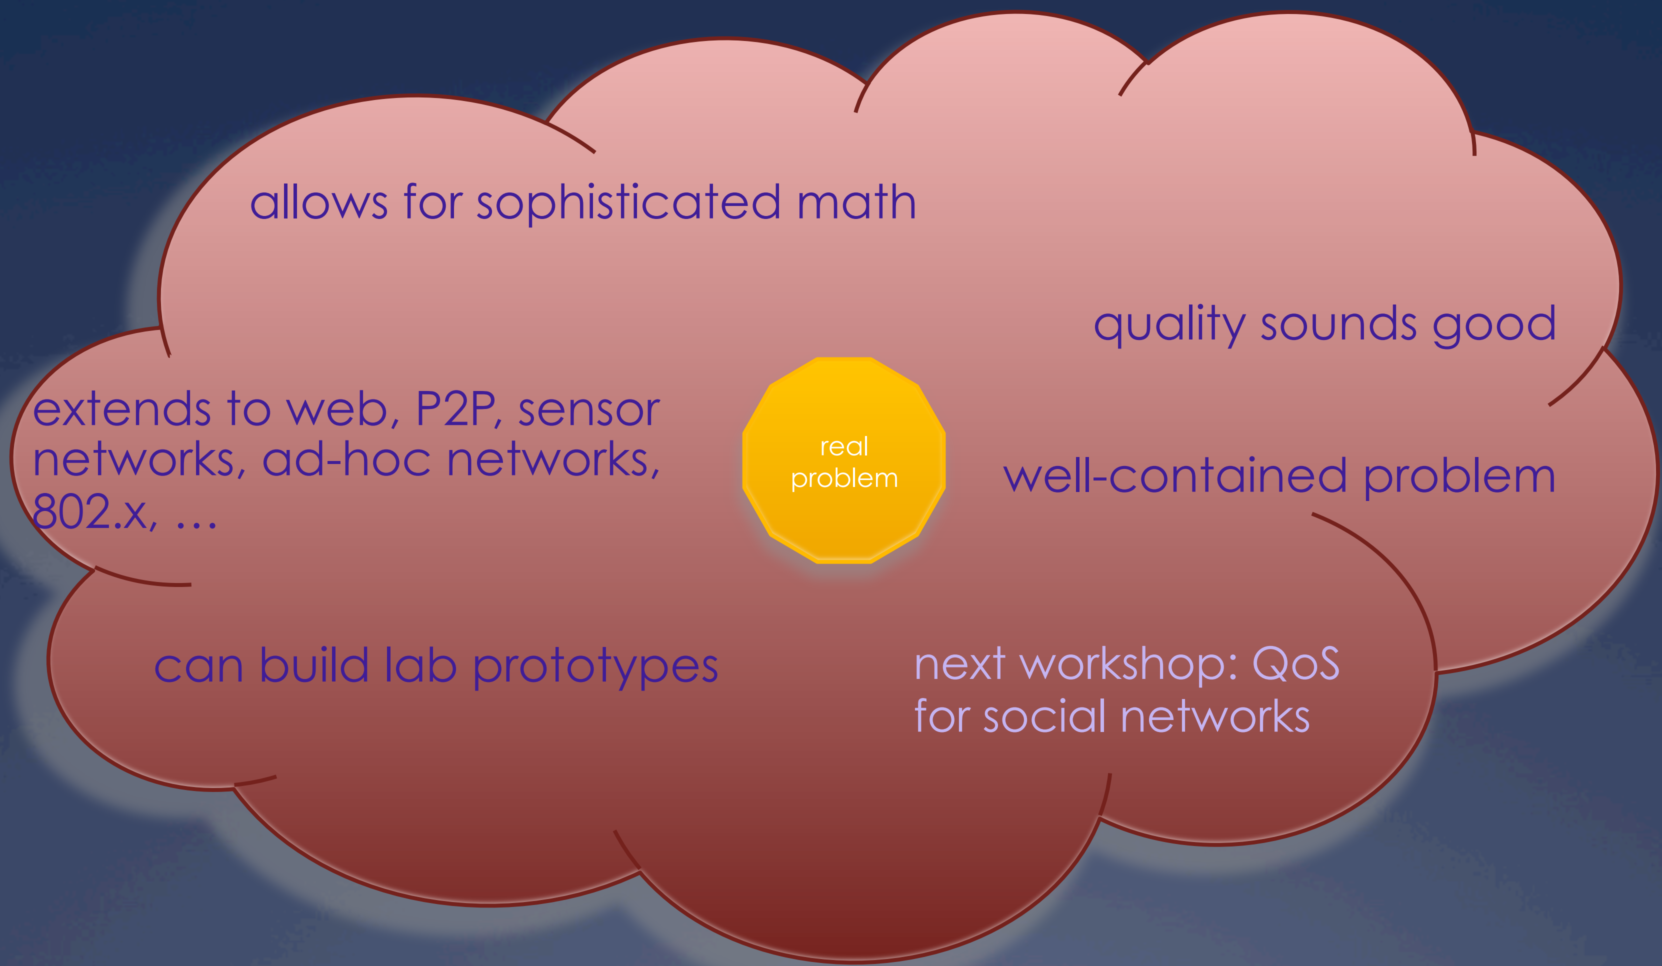

Old, old joke

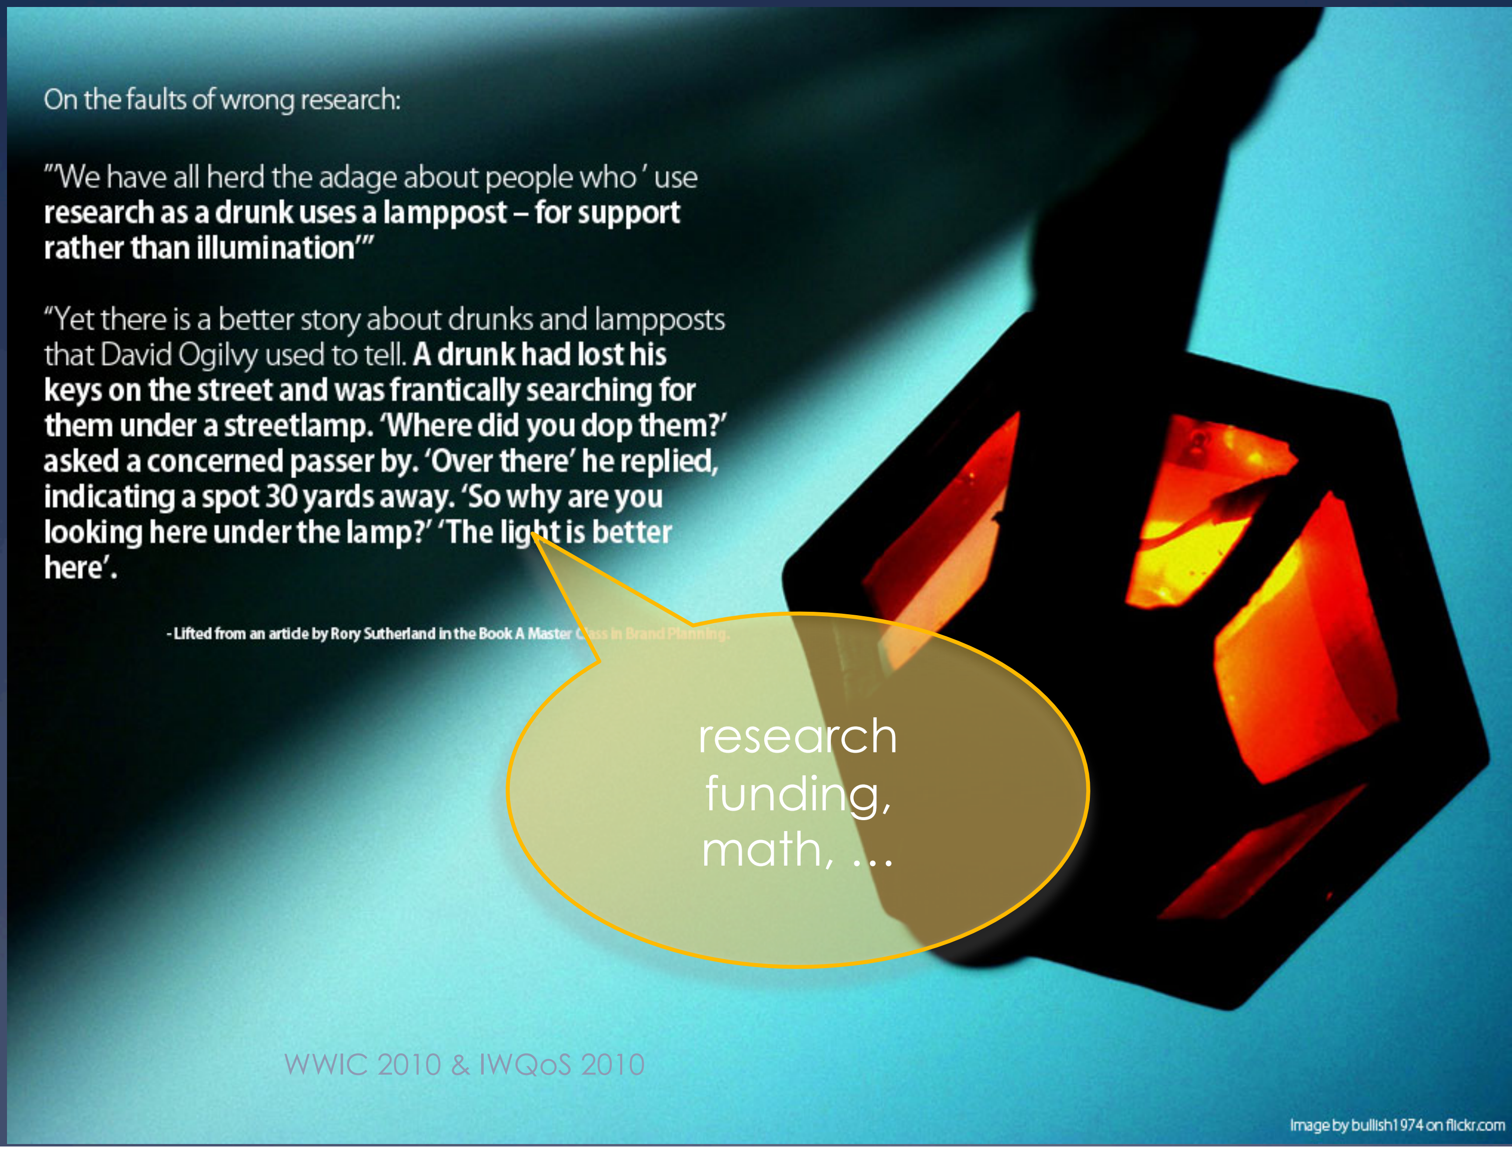
On the faults of wrong research:
"'We have all herd the adage about people who ' use **research as a drunk uses a lamppost – for support rather than illumination'"**
"Yet there is a better story about drunks and lampposts that David Ogilvy used to tell. **A drunk had lost his keys on the street and was frantically searching for them under a streetlamp. 'Where did you dop them?' asked a concerned passer by. 'Over there' he replied, indicating a spot 30 yards away. 'So why are you looking here under the lamp?' 'The light is better here'.**
- Lifted from an article by Rory Sutherland in the Book A Master Class in Brand Planning.
research funding, math, …
Image by bullish1974 on flickr.com

Almost ~~25~~100 years of QoS

Toll Telephone Traffic

Experiments are described to determine the relationship between telephone circuit loads and the corresponding delay to traffic. The operating methods employed and the number of circuits available determine in general the number of messages per day which can be handled over a single toll circuit. The average delay to traffic obviously depends upon the number of messages per circuit per day, or the circuit loads. With a given load factor, increase in the circuit loads will increase the average delay to traffic. At the same time the revenue per circuit mile will correspondingly increase. The practical limit, however, is approached when the delays to traffic reach a point where the service is unsatisfactory. The results of the experiments described illustrate the fact that increasing circuit loads increase the delay to traffic, and vice versa. The revenue per circuit mile is directly proportional to the product of the circuit load and the toll rate per minute-mile; consequently the relationship between the quality of service and the toll rate is generally obvious, assuming a certain rate of return on the plant investment.

Frank Fowle, *Transactions of the American Institute of Electrical Engineers*, June 1914

More early QoS work

Second generation computer control procedures for dial-a-ride

Based on operational experience with initial computer control procedures, more sophisticated procedures have been developed designed to provide a greater variety of services simultaneously and to allow the operator more discretion in the quality of service provided. This paper describes these second generation control procedures and analyses their effectiveness in the light of previous operational experience and in a simulation context.

Nigel H. M. Wilson, *Decision and Control including the 14th Symposium on Adaptive Processes*, **1975**

First (?) QoS (+ security) paper

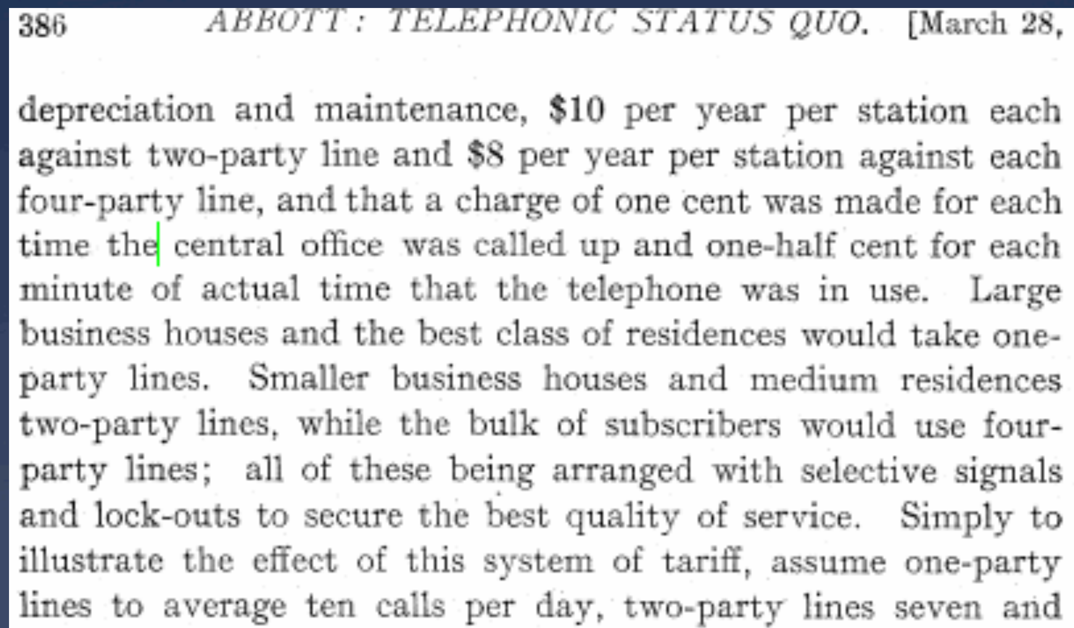

386 ABBOTT: TELEPHONIC STATUS QUO. [March 28,

depreciation and maintenance, $10 per year per station each against two-party line and $8 per year per station against each four-party line, and that a charge of one cent was made for each time the central office was called up and one-half cent for each minute of actual time that the telephone was in use. Large business houses and the best class of residences would take one-party lines. Smaller business houses and medium residences two-party lines, while the bulk of subscribers would use four-party lines; all of these being arranged with selective signals and lock-outs to secure the best quality of service. Simply to illustrate the effect of this system of tariff, assume one-party lines to average ten calls per day, two-party lines seven and

Abbott, Arthur Vaughan, "The Telephonic Status Quo," *American Institute of Electrical Engineers, Transactions of the, vol.XIX, pp. 373-388, Jan. 1902*

DiffServ v0: IP 791 (1981)

```
Bits 0-2:  Precedence.
Bit    3:  0 = Normal Delay,      1 = Low Delay.
Bits   4:  0 = Normal Throughput, 1 = High Throughput.
Bits   5:  0 = Normal Relibility, 1 = High Relibility.
Bit  6-7:  Reserved for Future Use.

   0     1     2     3     4     5     6     7
+-----+-----+-----+-----+-----+-----+-----+-----+
|                 |     |     |     |     |     |
|   PRECEDENCE    |  D  |  T  |  R  |  0  |  0  |
|                 |     |     |     |     |     |
+-----+-----+-----+-----+-----+-----+-----+-----+

  Precedence

    111 - Network Control
    110 - Internetwork Control
    101 - CRITIC/ECP
    100 - Flash Override
    011 - Flash
    010 - Immediate
    001 - Priority
    000 - Routine
```

QoS and energy - 1984

Energy Saving the "Record" System

A study is presently being conducted at the French Telecommunications Research Centre (CNET) in order to optimize the power consumption of air conditioning equipment in time-division exchanges. It is conducted within the frame of an "Energy Saving" campaign started by the French Administration. The so-called RECORD system (research for continuous optimal conditions of the air-conditioning system) was developed. This system enables the following functions to be performed: - acceptance and maintenance operations in air conditioning systems, - checking of power consumption, - evaluation of possible energy savings, provided the regulation instructions are modified within limits giving the same quality of service and reliability of the exchange.

Telecommunications Energy Conference, 1984. INTELEC '84.

Early packet QoS paper: 1986

This paper first examines quality of service as it applies to the Transport Service of the Open Systems Interconnection (OSI) Reference Model. Quality of service and some of the quality of service parameters applicable to the Transport Service are discussed. Also presented is a new concept concerning the "building-up" of quality of service. These ideas are then used to discuss the concept of robustness.

Table 2. Some Quality of Service Parameters for the Individual Layers of the Transport Service - Performance

	PERFORMANCE								
	SPEED				ACCURACY	RELIABILITY			
		DELAY							
TRANSPORT	THROUGHPUT	CONNECTION ESTABLISHMENT DELAY	TRANSIT DELAY	CONNECTION RELEASE DELAY	RESIDUAL ERROR RATE	CONNECTION ESTABLISHMENT FAILURE PROBABILITY	TRANSFER FAILURE PROBABILITY	CONNECTION RELEASE FAILURE PROBABILITY	CONNECTION RESILIENCE
NETWORK	THROUGHPUT	CONNECTION ESTABLISHMENT DELAY	TRANSIT DELAY	CONNECTION RELEASE DELAY	RESIDUAL ERROR RATE	CONNECTION ESTABLISHMENT FAILURE PROBABILITY	TRANSFER FAILURE PROBABILITY	CONNECTION RELEASE FAILURE PROBABILITY	CONNECTION RESILIENCE
DATA LINK	THROUGHPUT	*CONNECTION ESTABLISHMENT DELAY	TRANSIT DELAY	*CONNECTION RELEASE DELAY	RESIDUAL ERROR RATE	*CONNECTION ESTABLISHMENT FAILURE PROBABILITY	*TRANSFER FAILURE PROBABILITY	*CONNECTION RELEASE FAILURE PROBABILITY	CONNECTION RESILIENCE
PHYSICAL	TRANSMISSION RATE	*CONNECTION ACTIVATION DELAY	TRANSIT DELAY	*CONNECTION DEACTIVATION DELAY	ERROR RATE	*CONNECTION ACTIVATION FAILURE PROBABILITY	*TRANSFER FAILURE PROBABILITY	*CONNECTION DEACTIVATION FAILURE PROBABILITY	*CONNECTION RESILIENCE

*PARAMETERS NOT SPECIFIED IN OSI DOCUMENTATION

Pardue, M. D. ; Harvey, J. A. ; Haupt, K. D. ; Orlando, T. A., *MILCOM 1985*

QoS research activity

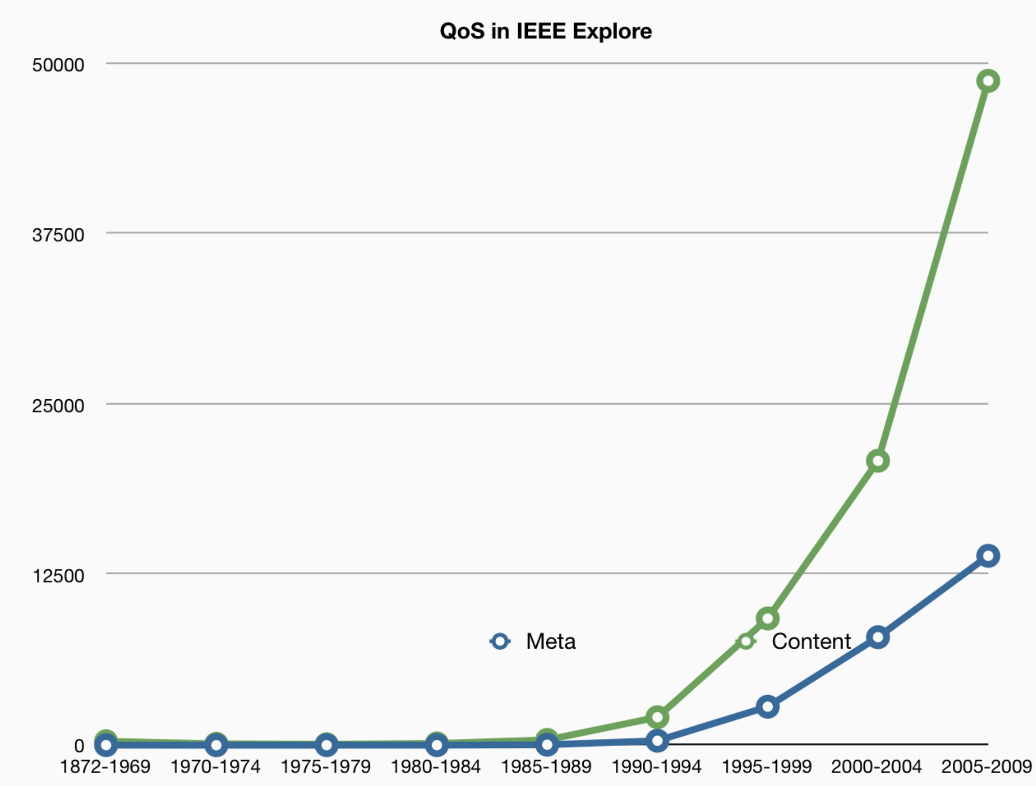

QoS research

- IEEE: 25,583 papers with "QoS" in metadata through 5/2010
  - 84,257 with QoS in meta data or text
- 2 papers/PhD year
- $50,000/PhD year
- → $640M in QoS research

What might we learn?

Cause of death for the next big thing

	QoS	multi-cast	mobile IP	active networks	IPsec	IPv6
not manageable across competing domains	✝	✝	✝	✝		
not configurable by normal users (or apps writers)	✝			✝	✝	
no business model for ISPs	✝	✝	✝	✝	✝	✝
no initial gain	✝	✝	✝	✝		✝
80% solution in existing system	✝	✝	✝	✝	✝	✝ (NAT)
increase system vulnerability	✝	✝	✝	✝		

Why did e2e QoS fail?

* Trivial issue: No uniform DiffServ code points
  * manual configuration of applications and home gateways
* No clearing house or end-to-end identity
* No economic model
  * flat, peak-rate based charging common
  * interesting model: *metro pricing*
* Lots of factors outside carrier control
  * home & enterprise network

Users don't care about QoS

* not even QoE
* they do care about service reliability:
  * consumer grade ~ electricity (99%?)
    * 99.5% → 43.8 hours outage/year
  * commercial grade
    * e.g., web server
    * Google Apps: 99.9% uptime SLA
    * Verizon business DSL SLA: 99%
  * critical grade
    * e.g., tele-surgery
    * typical by redundancy

QoS problems are real

The New York Times

Companies

WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HEALTH | SPORTS | OPINION

Search Technology Go

Inside Technology
Internet | Start-Ups | Business Computing | Companies

Bits Blog

Customers Angered as iPhones Overload AT&T

By JENNA WORTHAM
Published: September 2, 2009

Slim and sleek as it is, the iPhone is really the Hummer of cellphones.

Enlarge This Image

Michael Appleton for The New York Times
AT&T monitors its network from its operations center in Bedminster, N.J.
More Photos »

It's a data guzzler. Owners use them like minicomputers, which they are, and use them a lot. Not only do iPhone owners download applications, stream music and videos and browse the Web at higher rates than the average smartphone user, but the average iPhone owner can also use 10 times the network capacity used by the average smartphone user.

"They don't even realize how much data they're using," said Gene Munster, a senior securities analyst with Piper Jaffray.

Multimedia

SIGN IN TO RECOMMEND
TWITTER
COMMENTS (322)
E-MAIL
SEND TO PHONE
PRINT
REPRINTS
SHARE

… but traditional QoS research unlikely to help

QoS: more than L2 + L3

DNS lookup
IPv6/IPv4

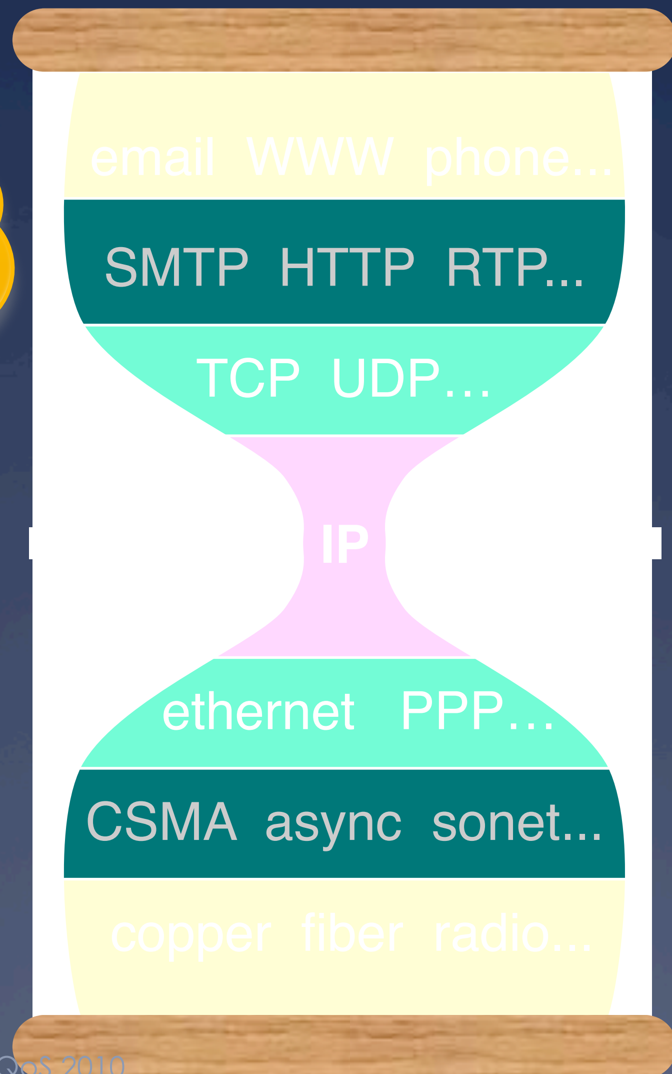

20% of the problem,
80% of the effort

DNS delays

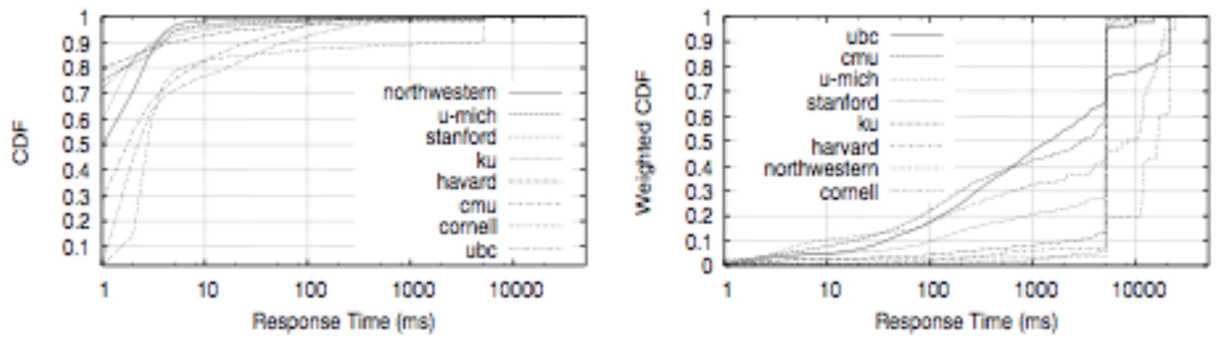

(a) Fraction of Lookups Taking < X ms

(b) Fraction of the Sum of Lookups Taking < X ms

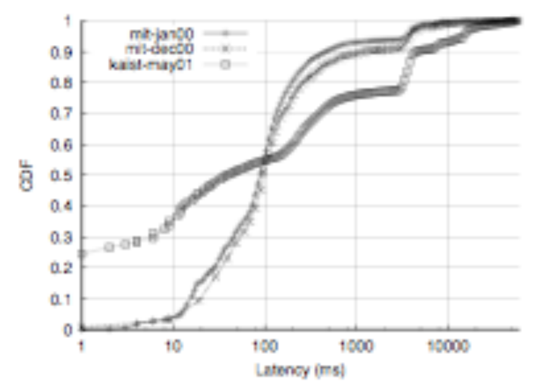

Fig. 3. Cumulative distribution of DNS lookup latency.

Park, Pai, Peterson, Wang (OSDI04)

Jung et al (ToN 2002)

Google vs. OpenDNS

likely exceeds page transfer delay

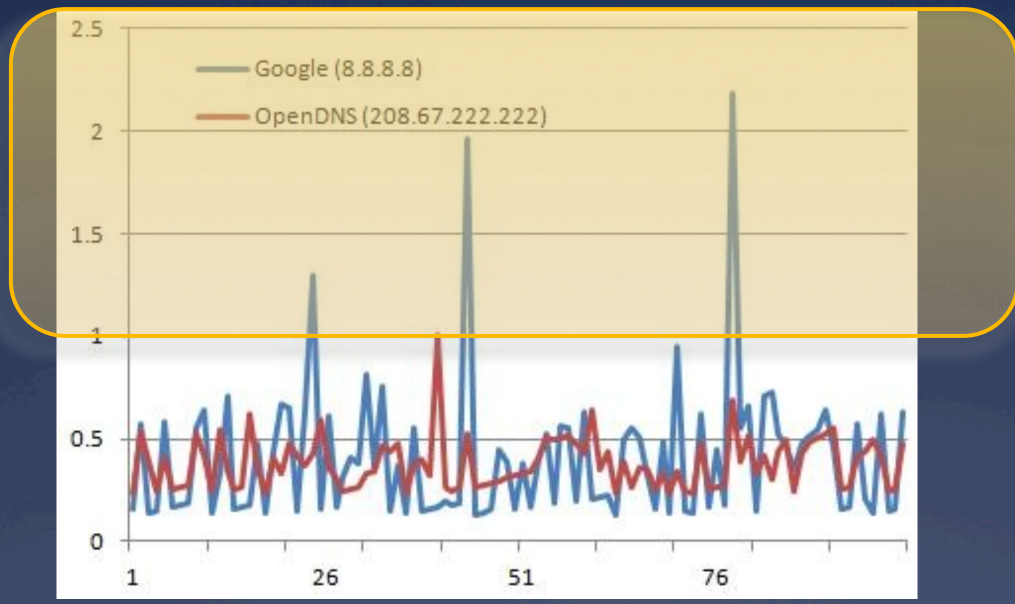

Dec. 2009 -- http://blog.gadodia.net/performance-comparison-of-opendns-and-google-dns/

What happens to the QoS losers?

- low priority

"TCP is elastic"

but demand is inelastic

- look at so many web pages
- watch that YouTube video **now**
- download email

- defer demand
- reduce demand

two choices

Deferring demand

* Capacity need is driven by peak demand
  * thus, useful to defer peak
* Cf. electric utilities
  * peak electricity costs >> baseload costs
  * but peak bandwidth costs = average costs
* Peak deferral
  * μs to ms:
    * node & router queues
  * minutes:
    * scheduling VoIP > TCP at home
    * Dad's phone call beats son's Hulu show
  * hours:
    * download OS patches
    * back-up
    * → *scavenger service*

Diurnal variation of traffic demand

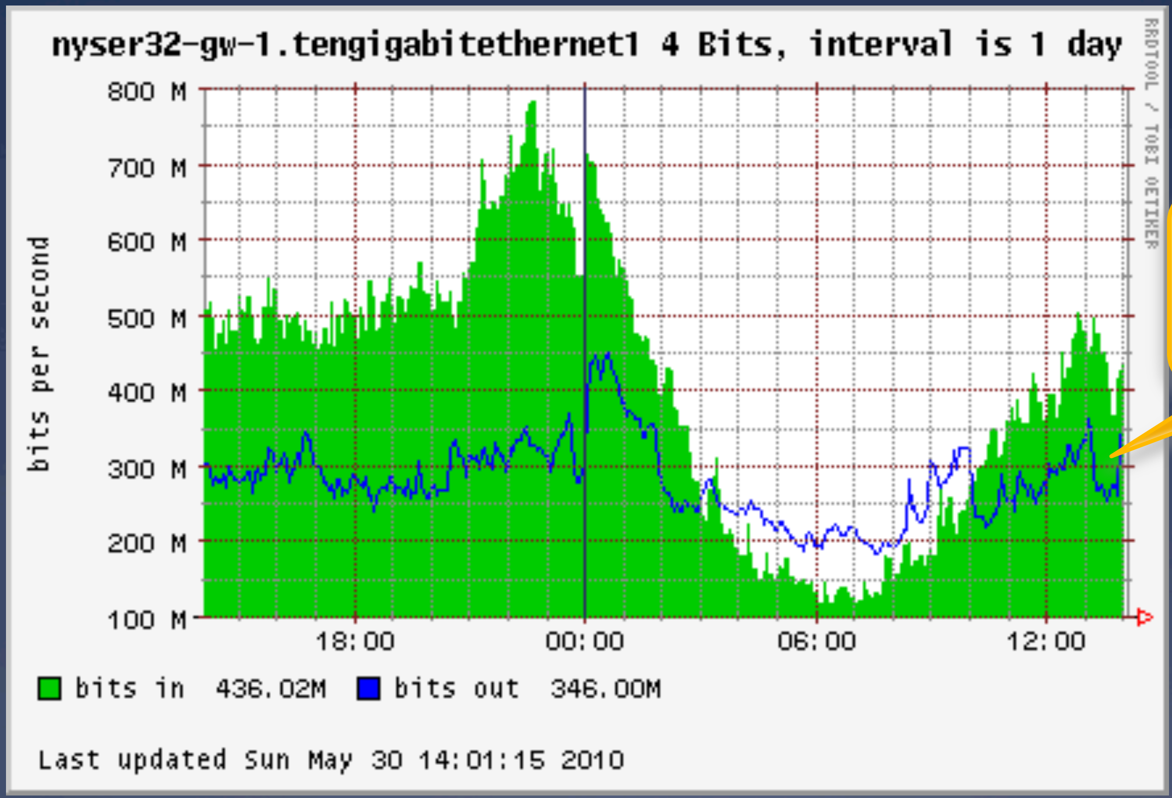

peak ≅ 1.7 avg.

Columbia University commercial Internet access (10 GigE), May 30, 2010

Electric Load Duration Curve

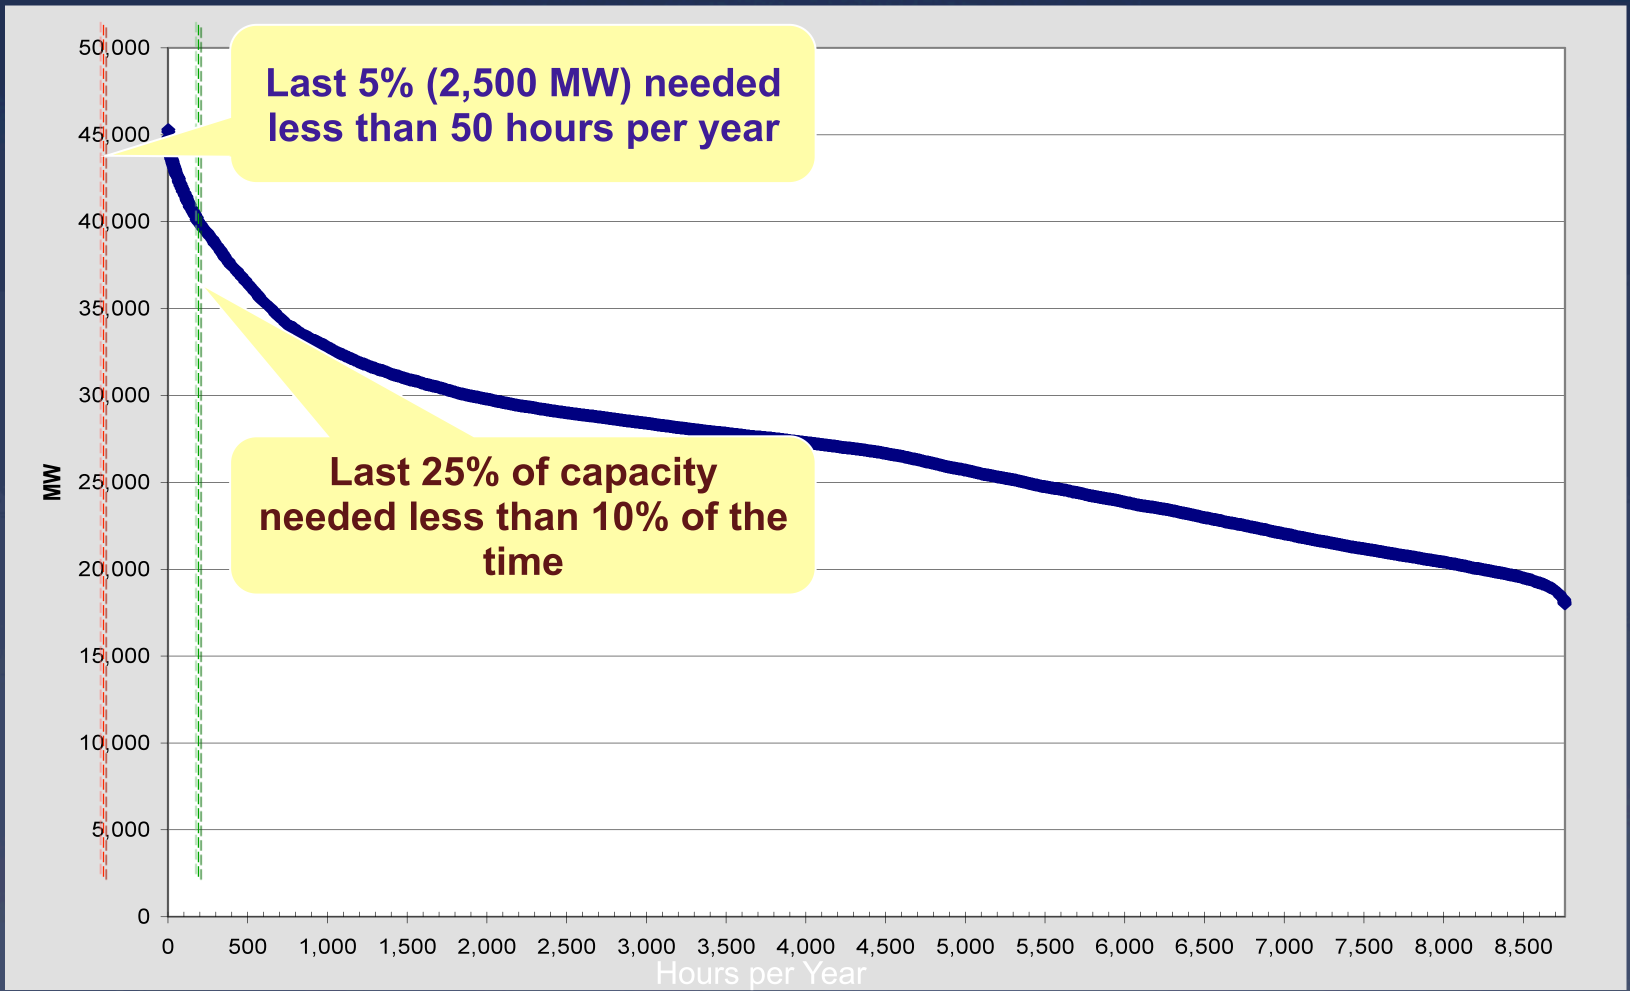

Source: California Independent System Operator Corporation

Electricity diurnal demand

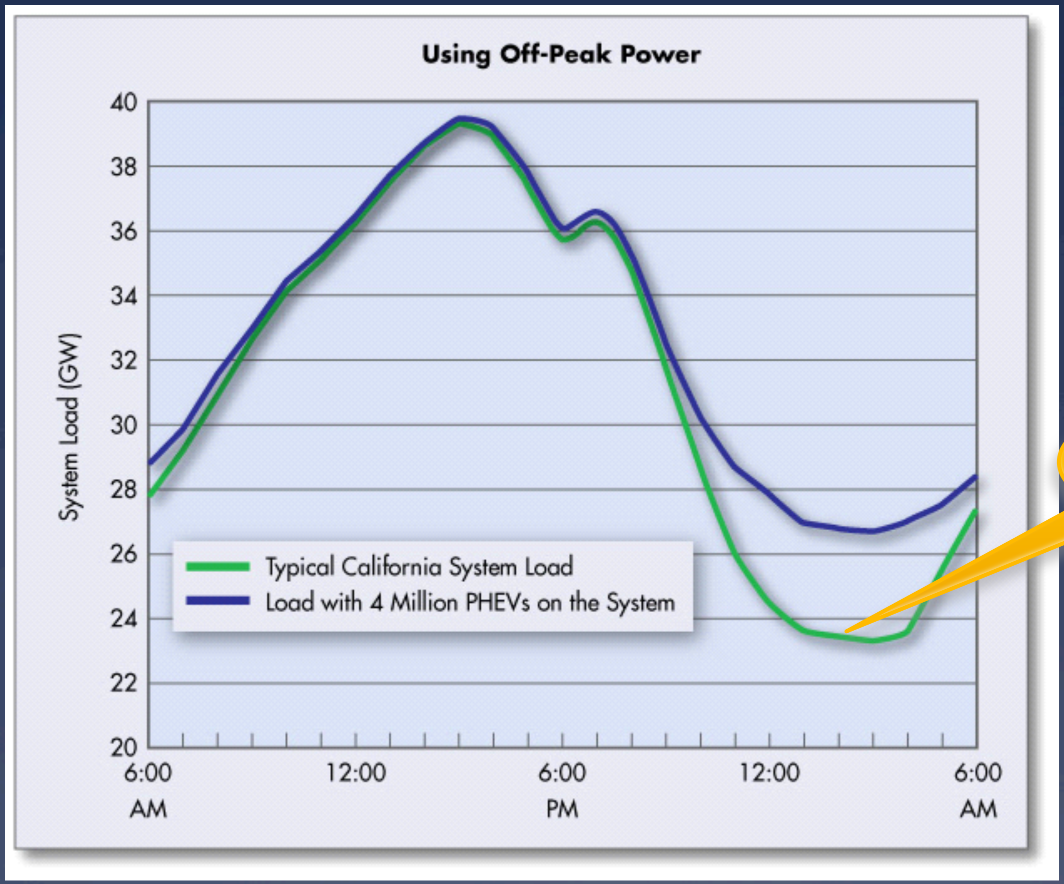

about 60%

"Environmental Assessment of Plug-In Hybrid Electric Vehicles (PHEVs)", June 2009

QoS vs. flying business class

economy vs. business class	QoS
always more leg room and better (any) food	only during congestion
flights are mostly full (load factor 80%+)	networks are mostly empty (20-30%)
better food & nicer flight attendants	looks the same
airline doesn’t get blamed for traffic jam on the way to the airport	packet loss at home looks just the same
more frequent flyer miles	there’s an idea…

ITU-T Y.1541 QoS classes

	0	1	2	3	4	5
IPTD (transfer delay)	100 ms	400 ms	100 ms	400 ms	1 s	U
IPDV (jitter)	50 ms		Unspecified			
IPLR (loss ratio)	0.1%					U
IPER (error rate)	0.01%					U
Usage	Voice		Signaling	Interactive Data	Streaming video	Best-effort data

doesn't provide necessary rate to watch movies

Application changes

Applications

Rank	Application	2007	2009	Change
1	**Web**	41.68%	52.00%	**24.76%**
2	**Video**	1.58%	2.64%	**67.09%**
3	**VPN**	1.04%	1.41%	**35.58%**
4	**Email**	1.41%	1.38%	**-2.13%**
5	**News**	1.75%	0.97%	**-44.57%**
6	**P2P** (*)	2.96%	0.85%	**-71.28%**
7	**Games**	0.38%	0.49%	**28.95%**
8	**SSH**	0.19%	0.28%	**47.37%**
9	**DNS**	0.20%	0.17%	**-15.00%**
10	**FTP**	0.21%	0.14%	**-33.33%**
	Other	2.56%	2.67%	**4.30%**
	Unclassified	46.03%	37.00%	**-19.62%**

(*) 2009 P2P Value based on 18% Payload Inspection
Weighted average percentage of all Internet traffic using well-known ports

probably includes RT traffic

- **Growing volume of Internet traffic uses port 80 / 443**
  - Includes significant video component and source of most growth
- **Unclassified includes P2P and video**
  - Payload matching suggests P2P at 18%
  - P2P is fastest declining

Craig Labovitz, "Internet Traffic and Content Consolidation", IETF March 2010.

P2P declining

P2P

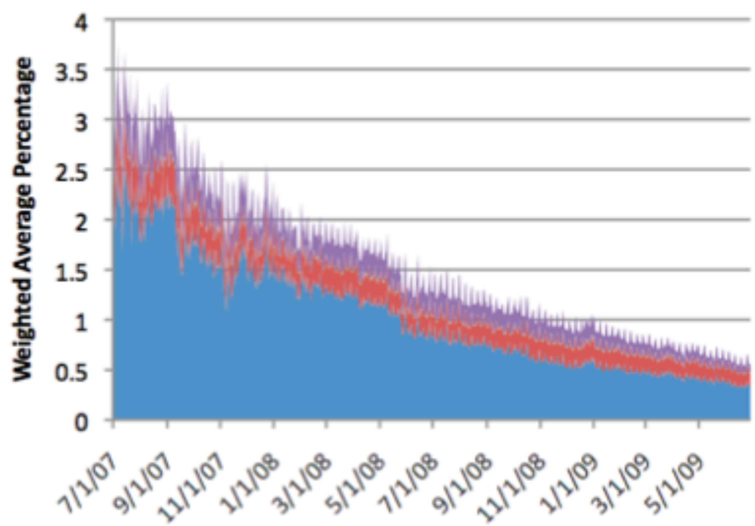

Graph of weighted average traffic using well-known P2P ports

- **In 2006, P2P one of largest threats facing carriers**
  - Significant protocol, engineering and regulatory effort / debate
- **In 2010, P2P fastest declining application group**
  - Trend in both well-known ports and payload based analysis
  - Still significant volumes
  - Slight differences in rate of decline by region (i.e. Asia is slower)

Craig Labovitz, "Internet Traffic and Content Consolidation", IETF March 2010.

Cisco's traffic prediction

Table 3. Global Consumer Internet Traffic, 2008–2013

Consumer Internet Traffic, 2008–2013						
	2008	**2009**	**2010**	**2011**	**2012**	**2013**
By Sub-Segment (PB per month)						
Web/Email	1,239	1,595	2,040	2,610	3,377	3,965
File Sharing	3,345	4,083	5,022	6,248	7,722	9,629
Internet Gaming	47	87	135	166	217	239
Internet Voice	103	129	152	174	183	190
Internet Video Communications	36	57	94	160	239	354
Internet Video to PC	1,112	2,431	4,268	6,906	9,630	12,442
Internet Video to TV	29	149	381	1,004	1,711	2,594
Ambient Video	110	224	634	1,332	2,089	2,715

Ambient video = nannycams, petcams, home security cams, and other persistent video streams

Cisco traffic prediction

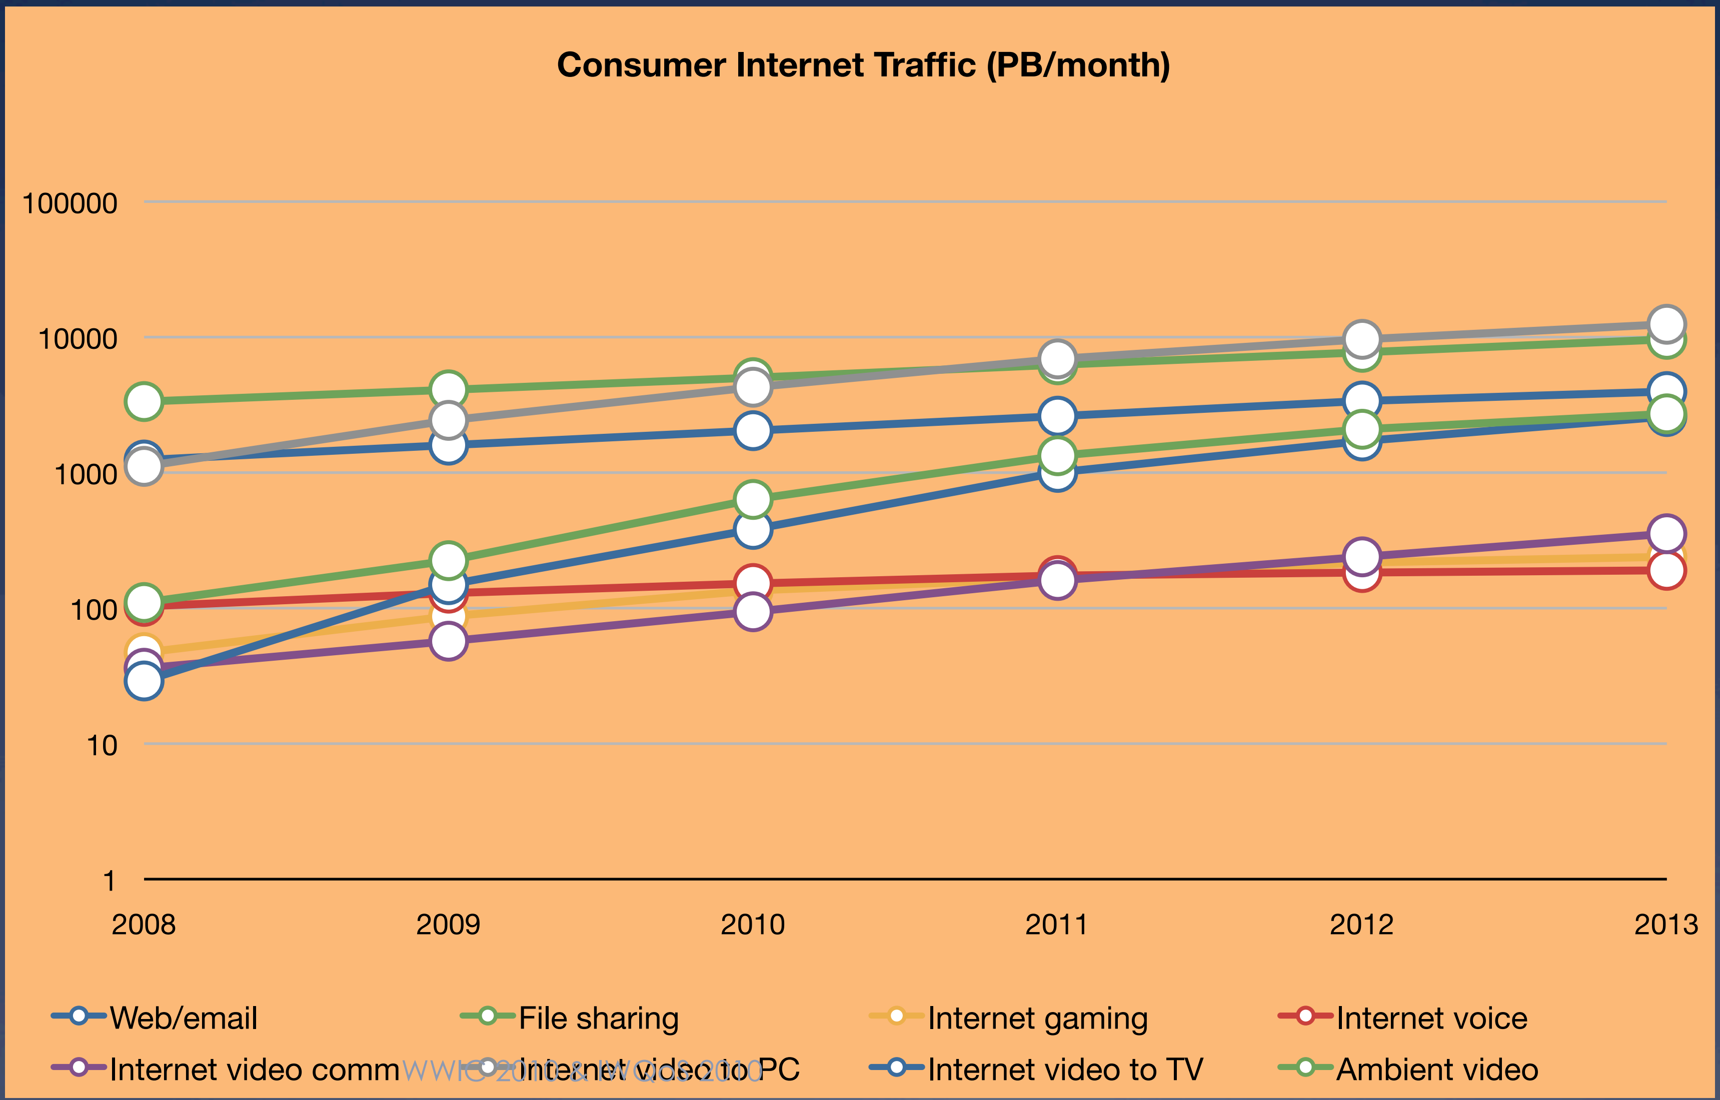

The race against abundance

- resource scarcity → QoS
  - Soviet model of economic planning: manage scarcity
- But turning away paying customers is not good business
- Few people will use unpredictable networks
  - “sorry, the Internet is sold out today”

What did we end up with?

* 1997: RFC 2205 (*Resource ReSerVation Protocol (RSVP)*)
* 1998: RFC 2474 (*An Architecture for Differentiated Services*)
* DiffServ
  * typically, priority for VoIP
    * access, transport to PSTN gateway
  * RSVP for traffic engineering
* 802.11e
  * essentially DiffServ
* Volume limits (Comcast = 250 GB/month) or per-MB charges (mobile)
* Works well as long as highest priority is small fraction of total

The mantra of TCP fairness

* TCP-friendly: non-TCP traffic needs to be TCP-fair
  * back off under loss
  * RFC XXXX
* Problematic:
  * RTT-sensitive
    * good – may encourage local access
  * it's per session – but one web browser may open 4 connections
  * it's instantaneous only
    * what if I haven't sent for a week and you've been downloading 3 GB of YouTube?
  * assumes that all bits are worth the same to the user
* Bob Briscoe's work

Some QoS research issues

* How can a user tell where things are breaking?
* Subscriber-level QoS measurements
  * not just in academic networks
* What pricing models work for users?
  * congestion pricing: too unpredictable
    * how many MB are in that web page?
    * nice phone call – would you like to continue for $3/minute?
  * maybe content provider pays?
  * per-minute pricing for VoIP service + QoS
    * see Skype Access
  * tiered service, capturing 90% of customer group
    * see web server pricing
    * include some account of priority traffic

Performance of video chat clients under congestion

* Residential area networks (DSL and cable)
  * Limited uplink speeds (around 1Mbit/s)
  * Big queues in the cable/DSL modem(600ms to 6sec)
  * Shared more than one user/application
* Investigate applications' behavior under congestion
  * Whether they are increasing the overall congestion
  * Or trying to maintain a fair share of bandwidth among flows

How good is industrial practice?

Experimental setup

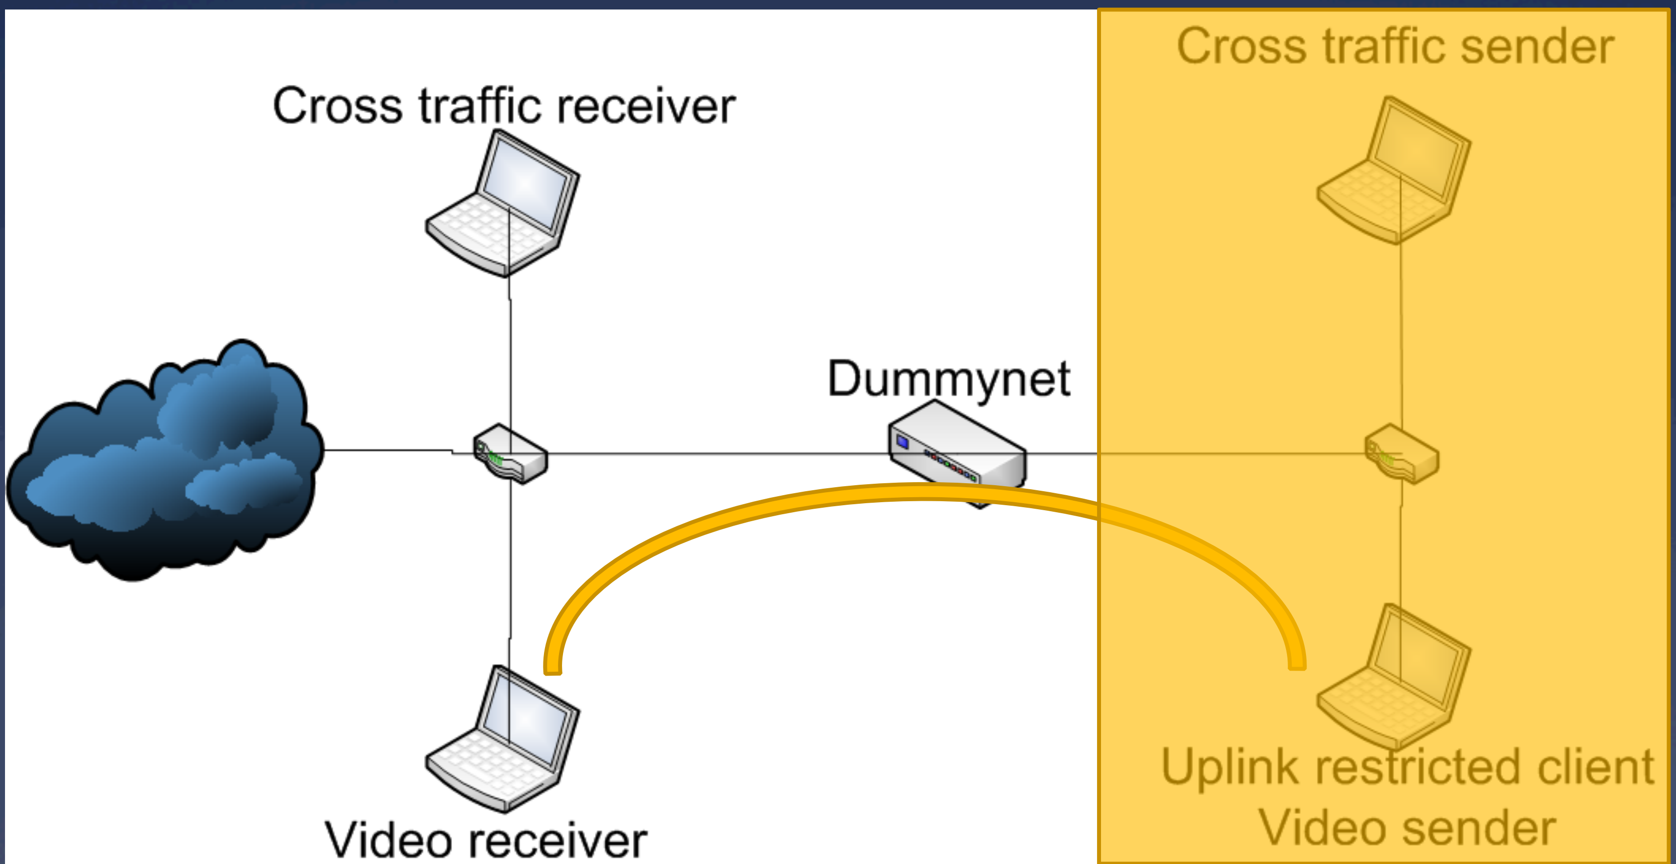

Skype

Step 10 s, 100 kb/s

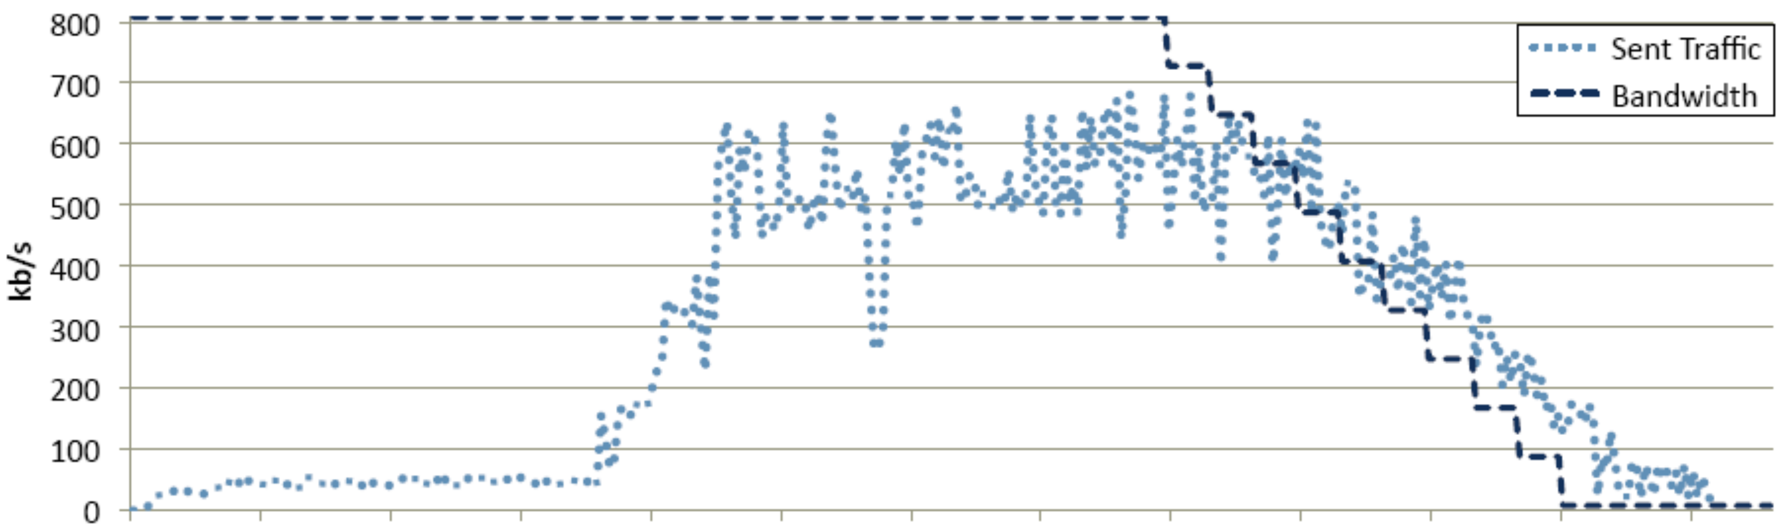

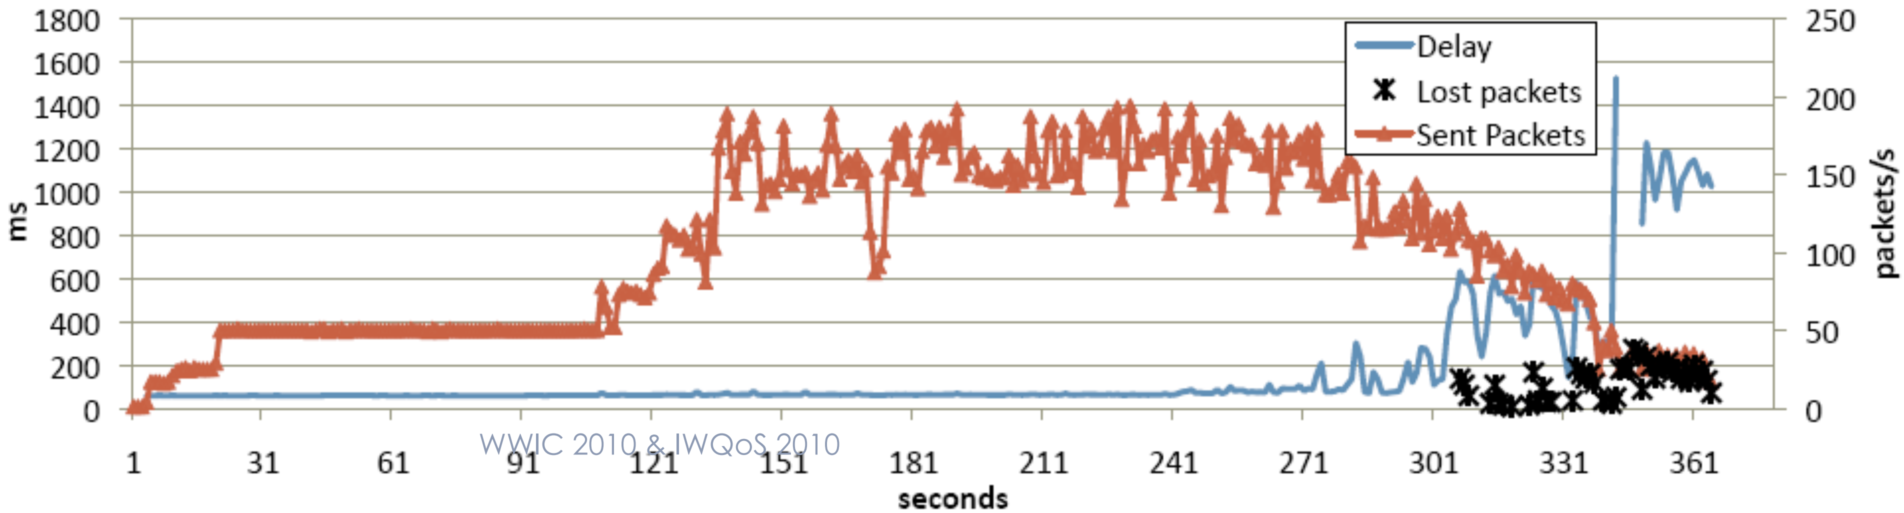

Live

Step 10 s 100 kb/s

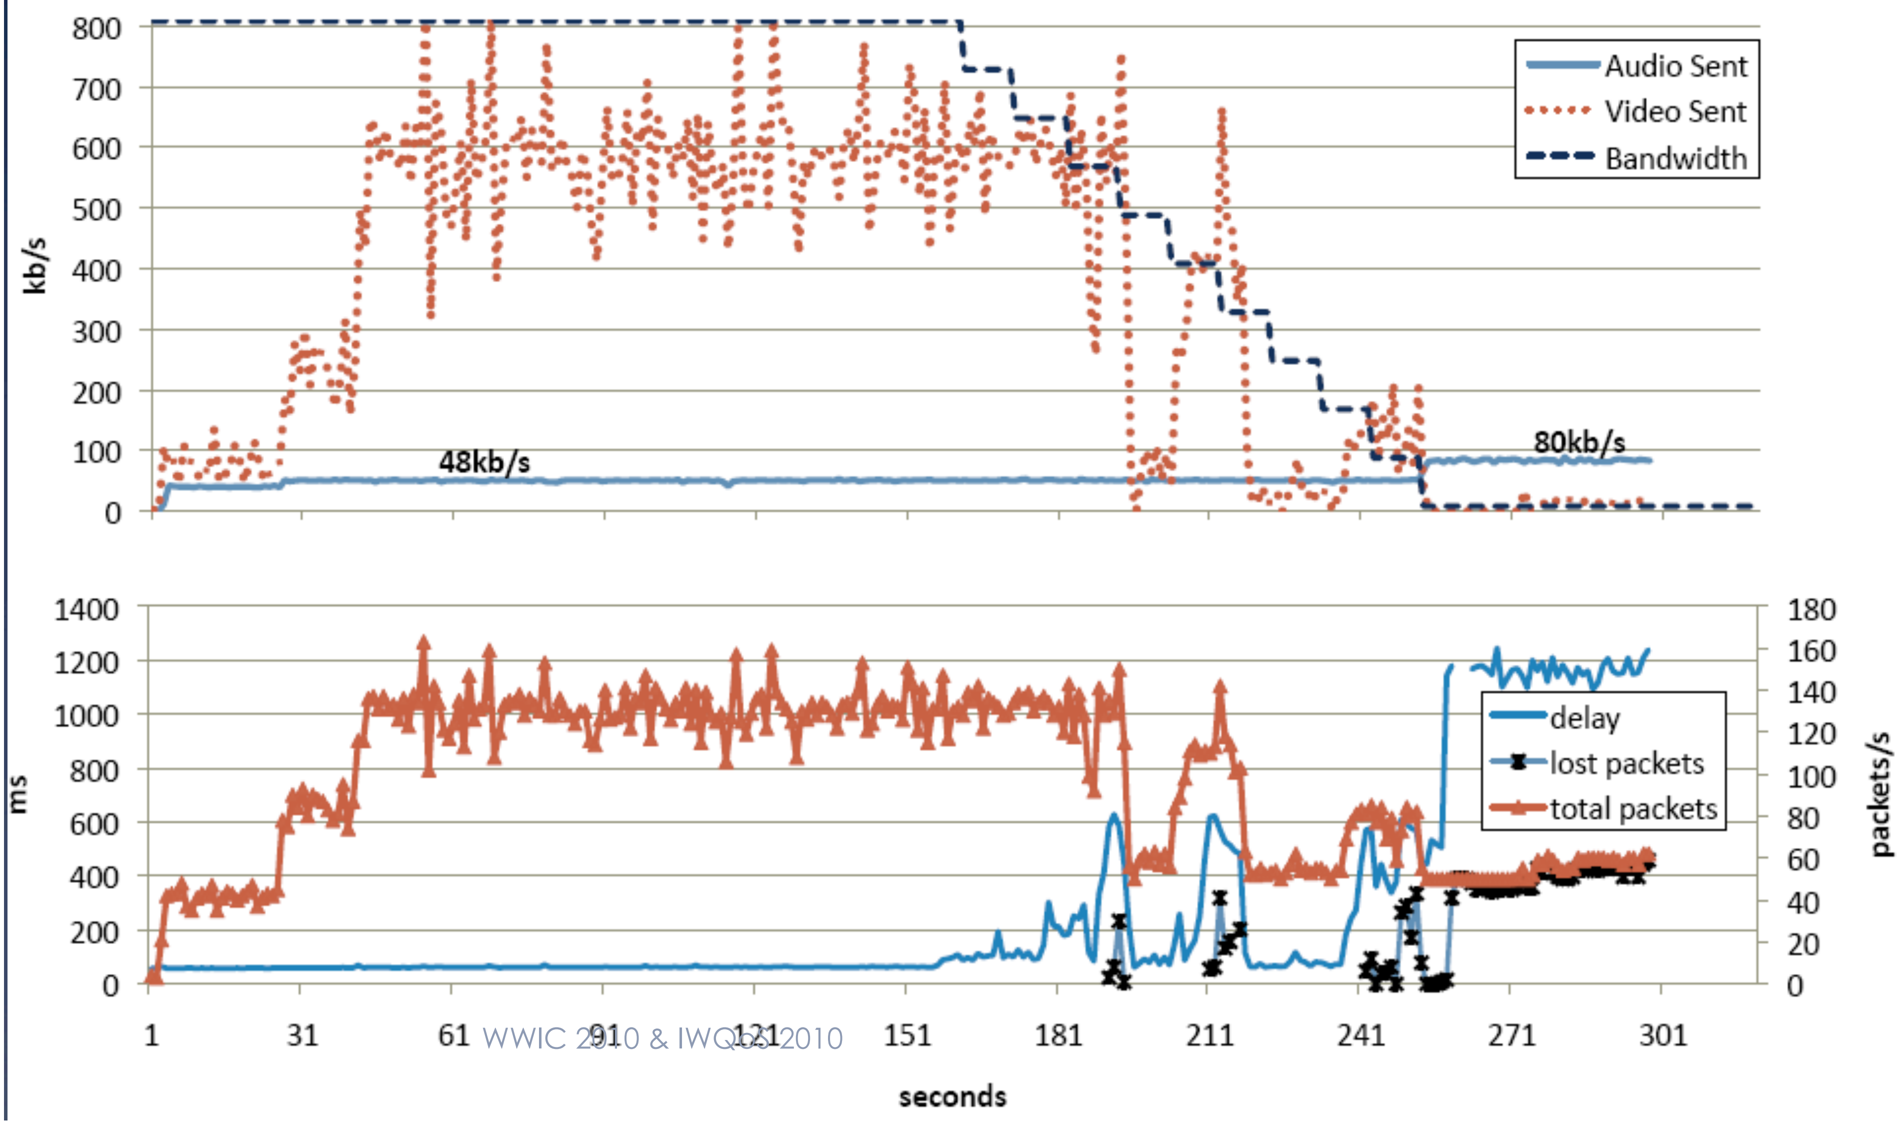

X-Lite

Step 10 s 100 kb/s

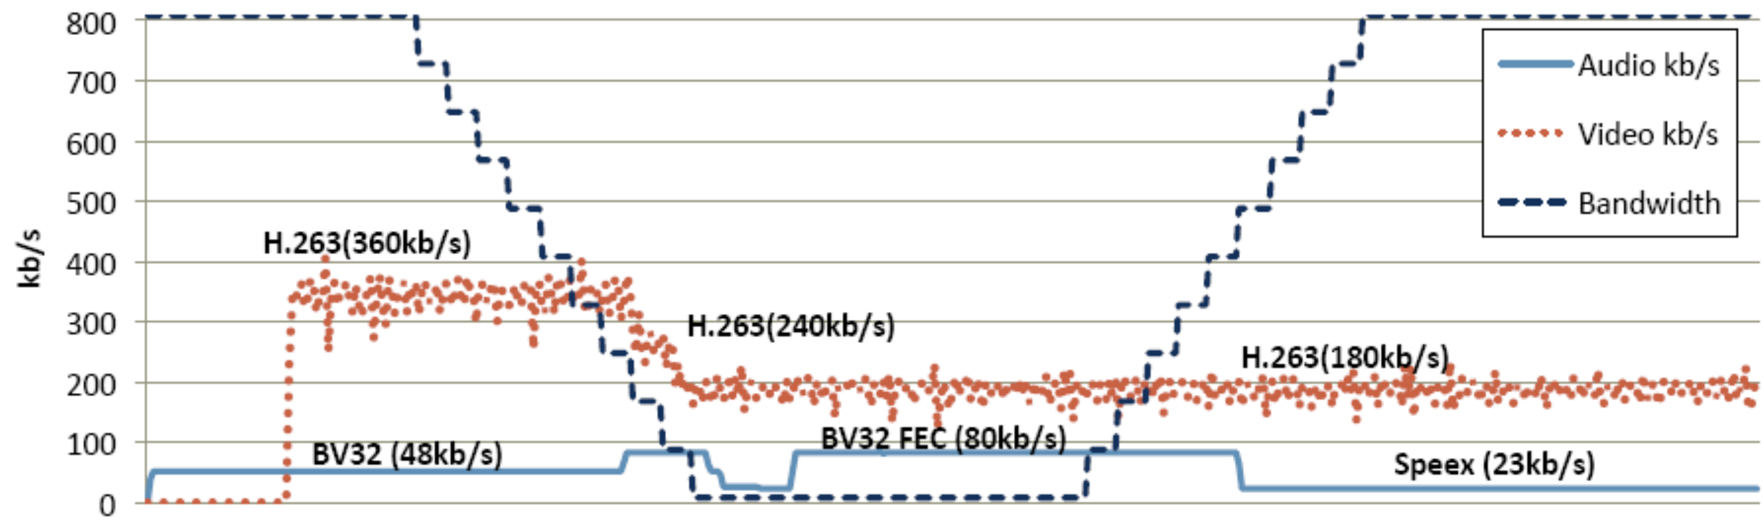

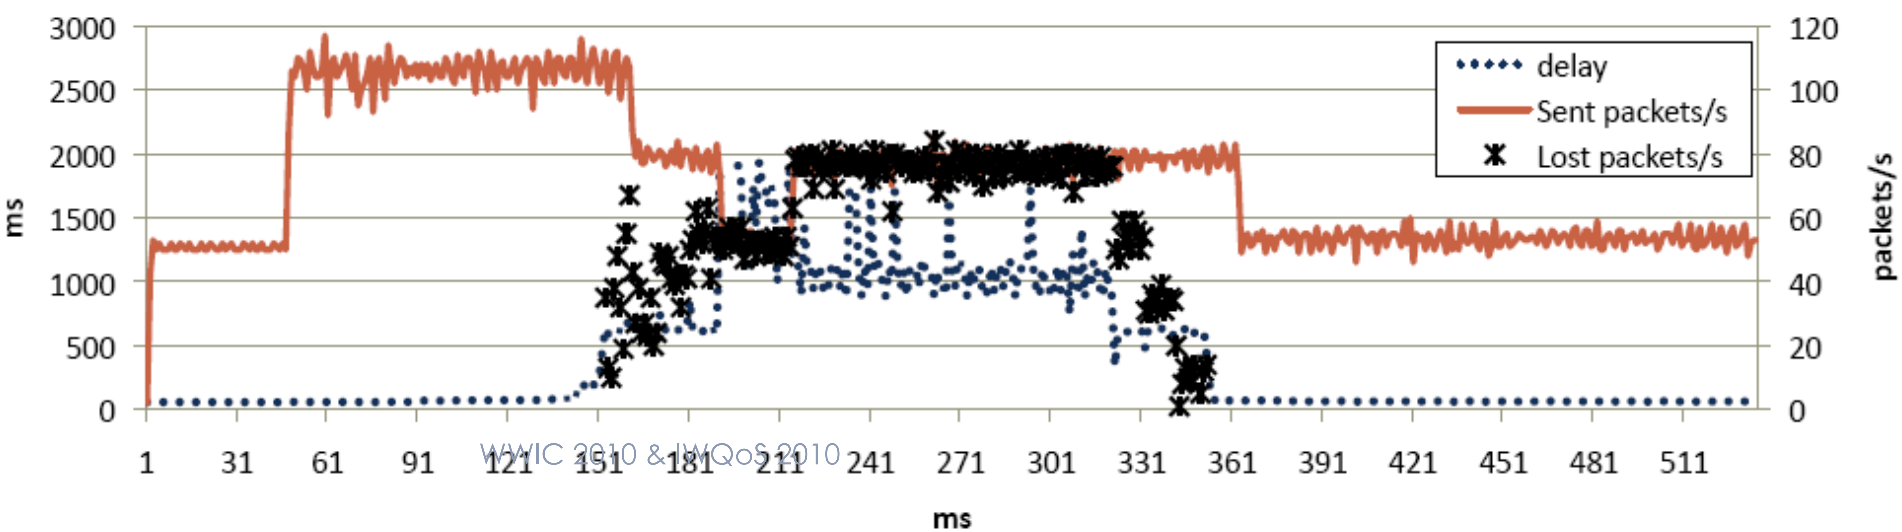

Skype File Transfer

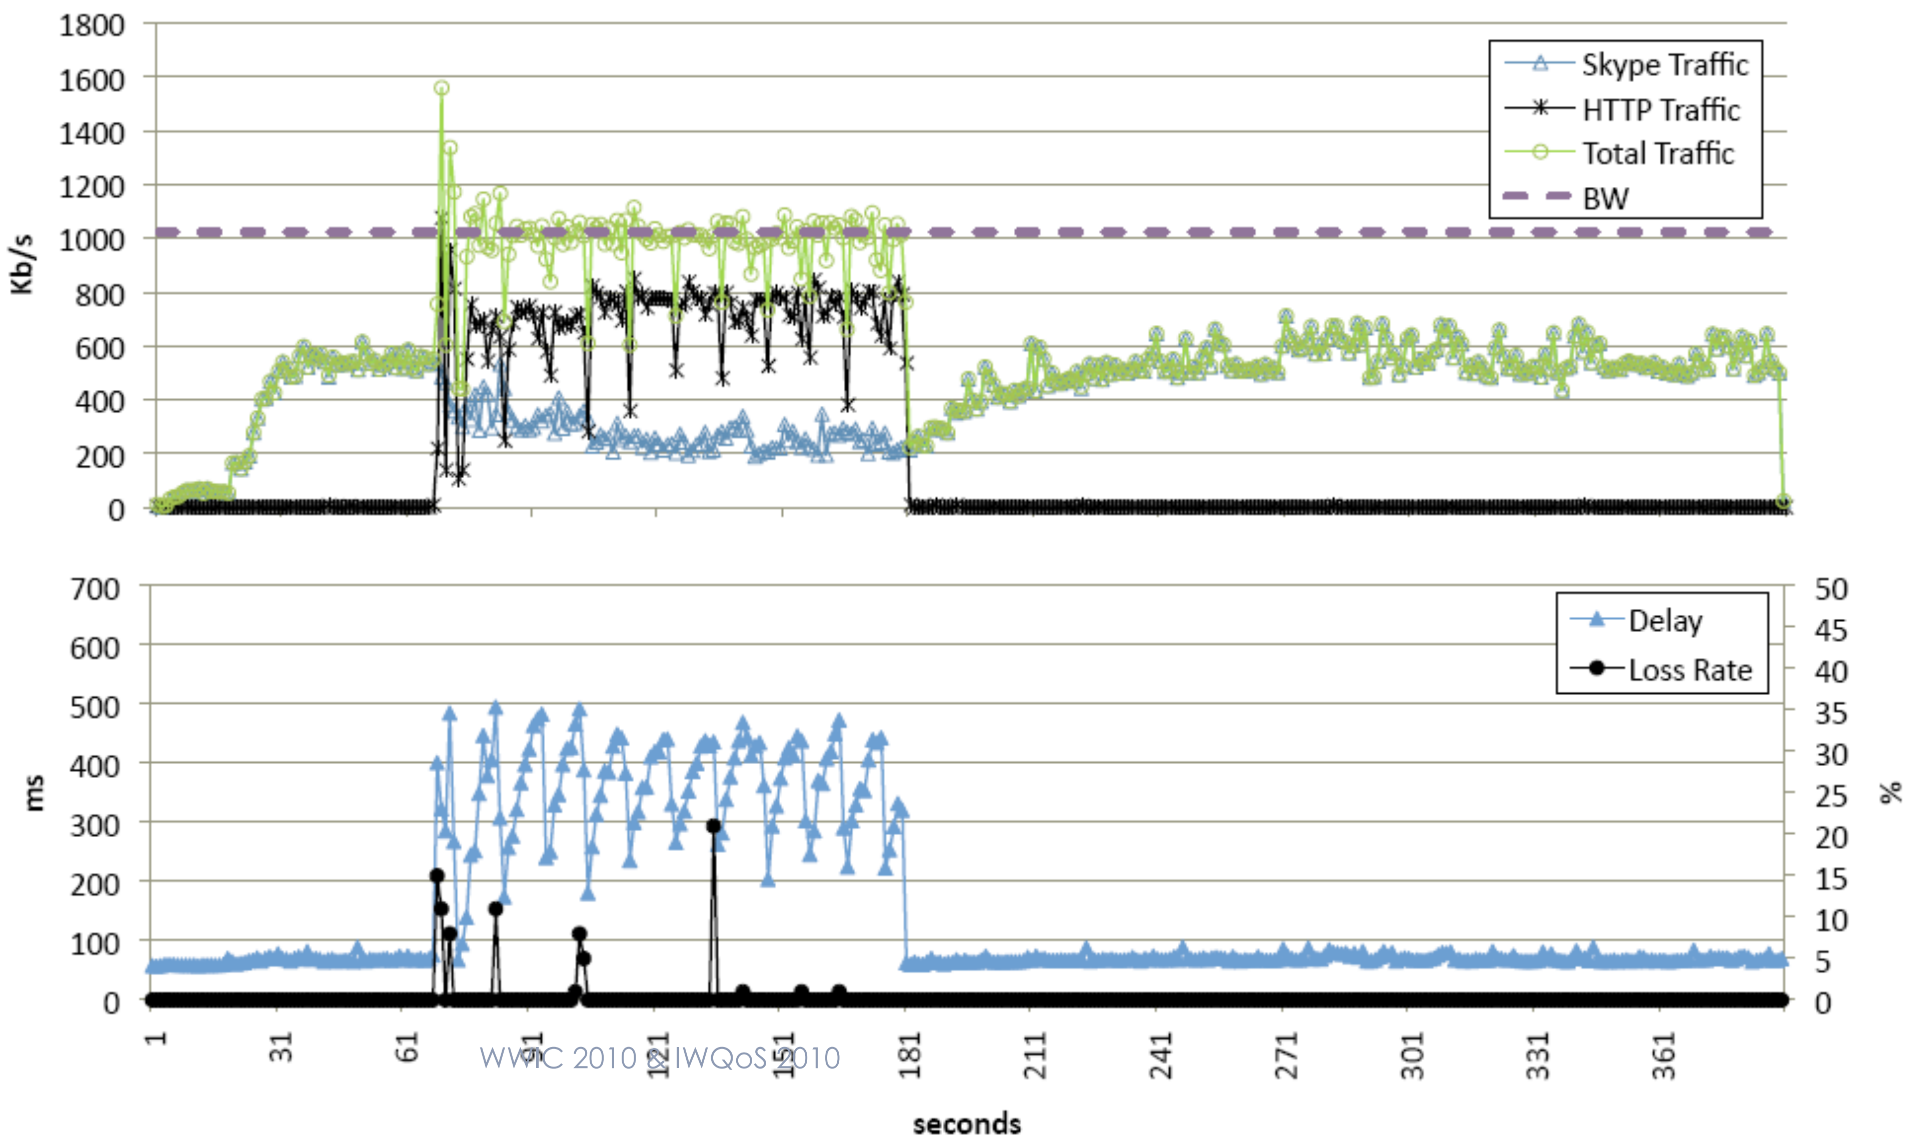

Live

File Transfer

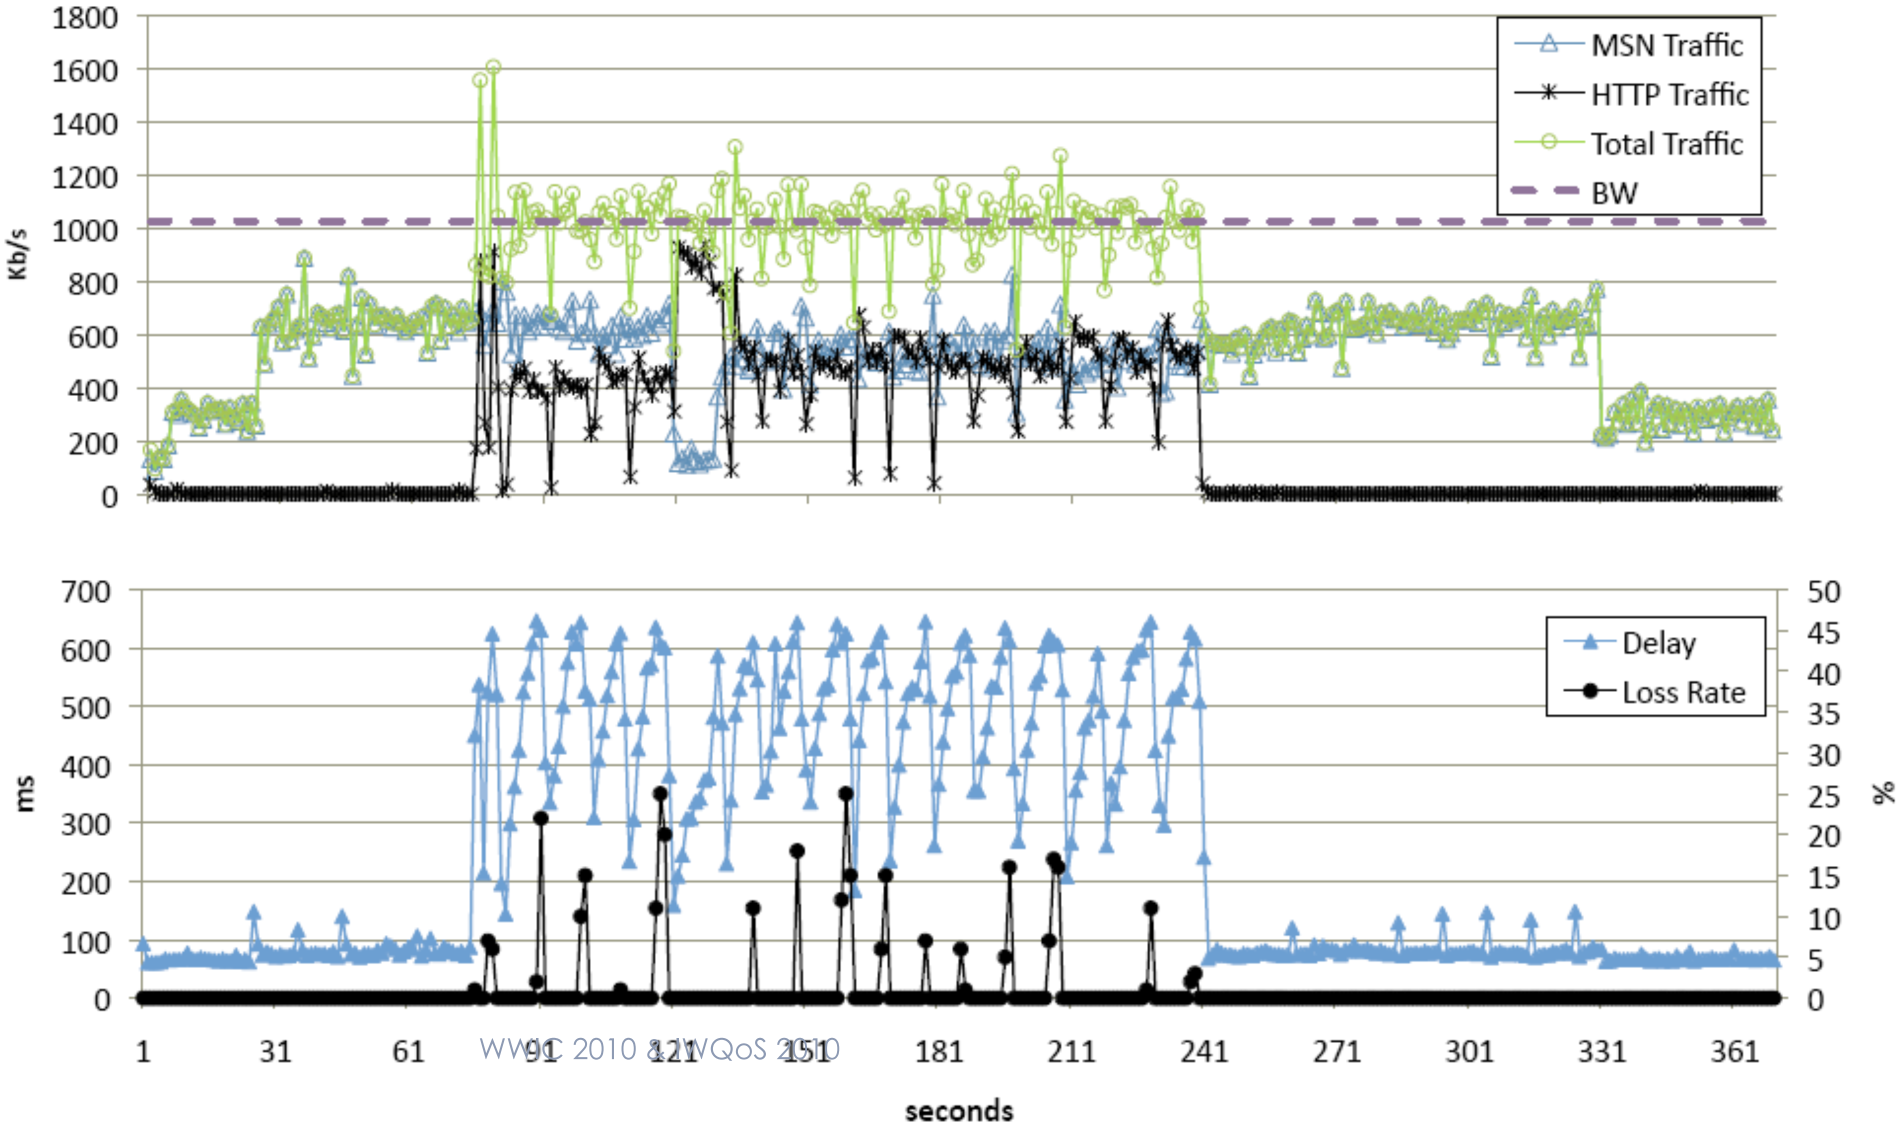

X-Lite File Transfer

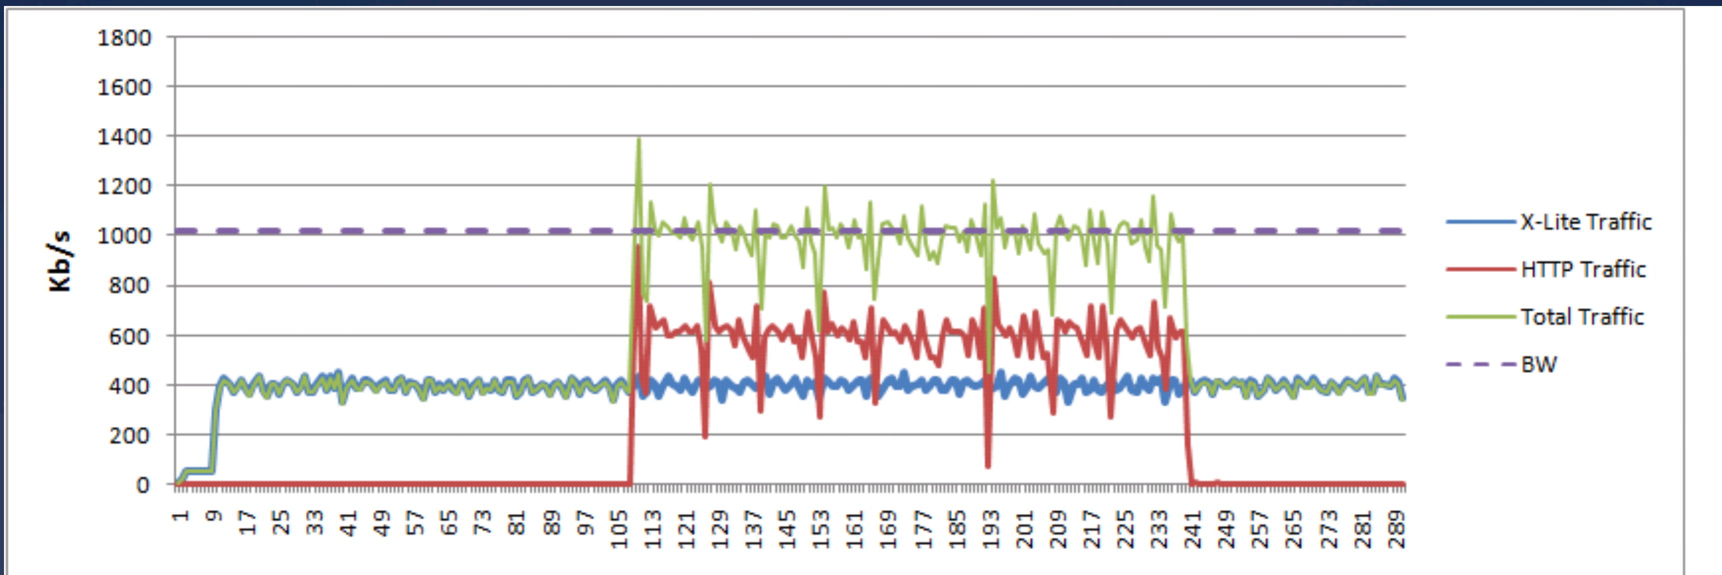

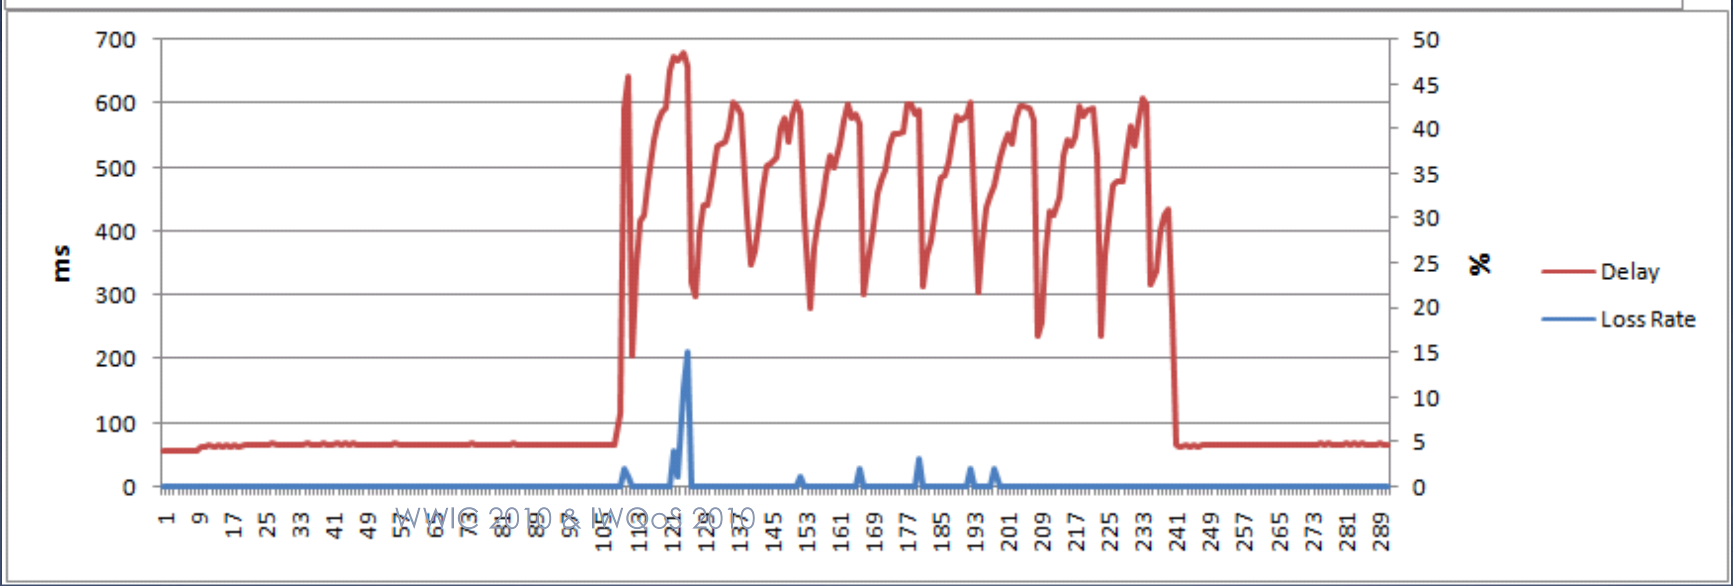

Eyebeam File Transfer

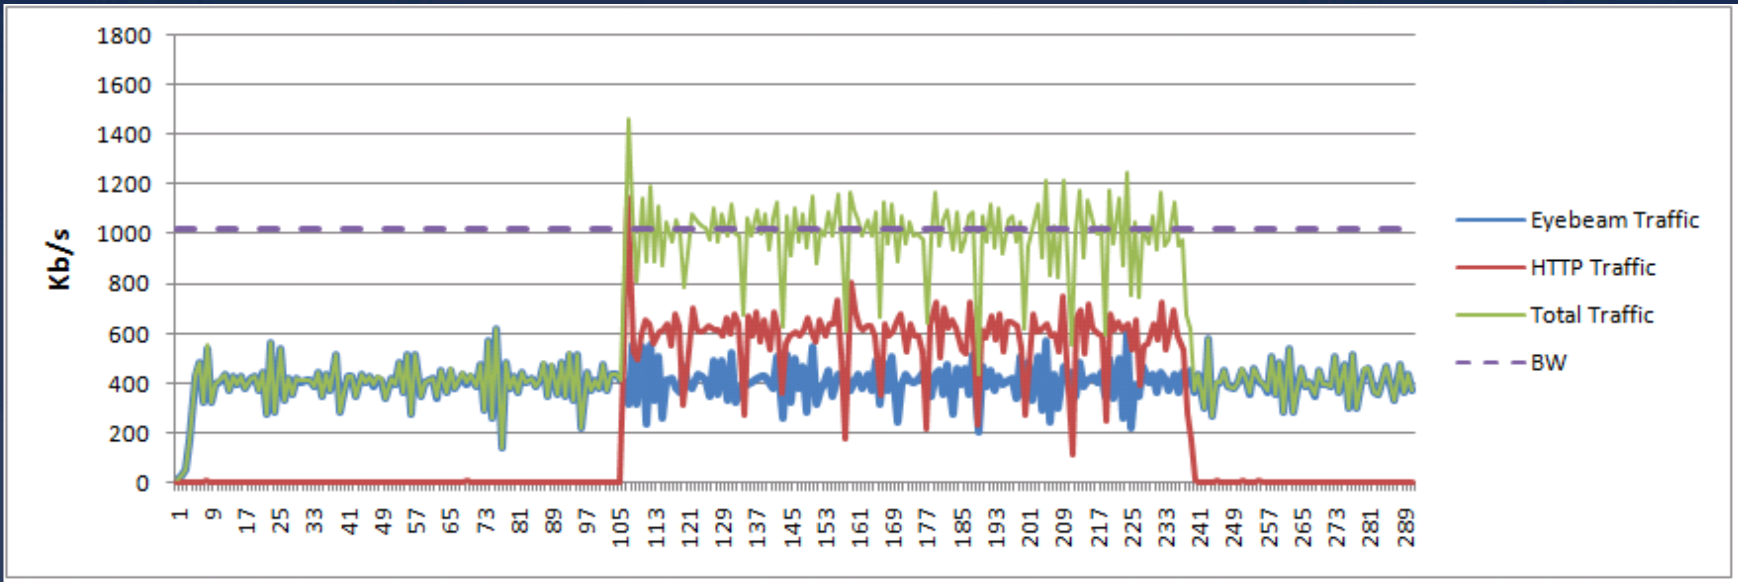

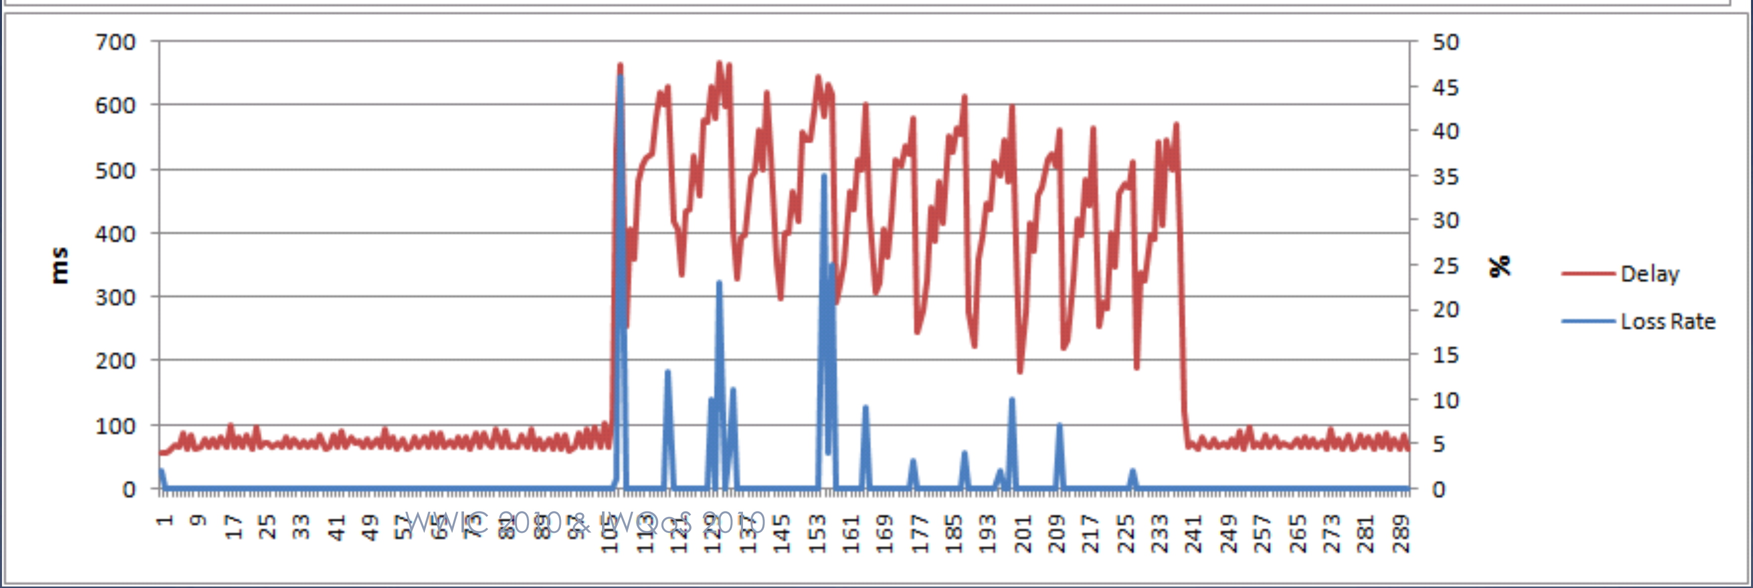

Skype

Bittorrent

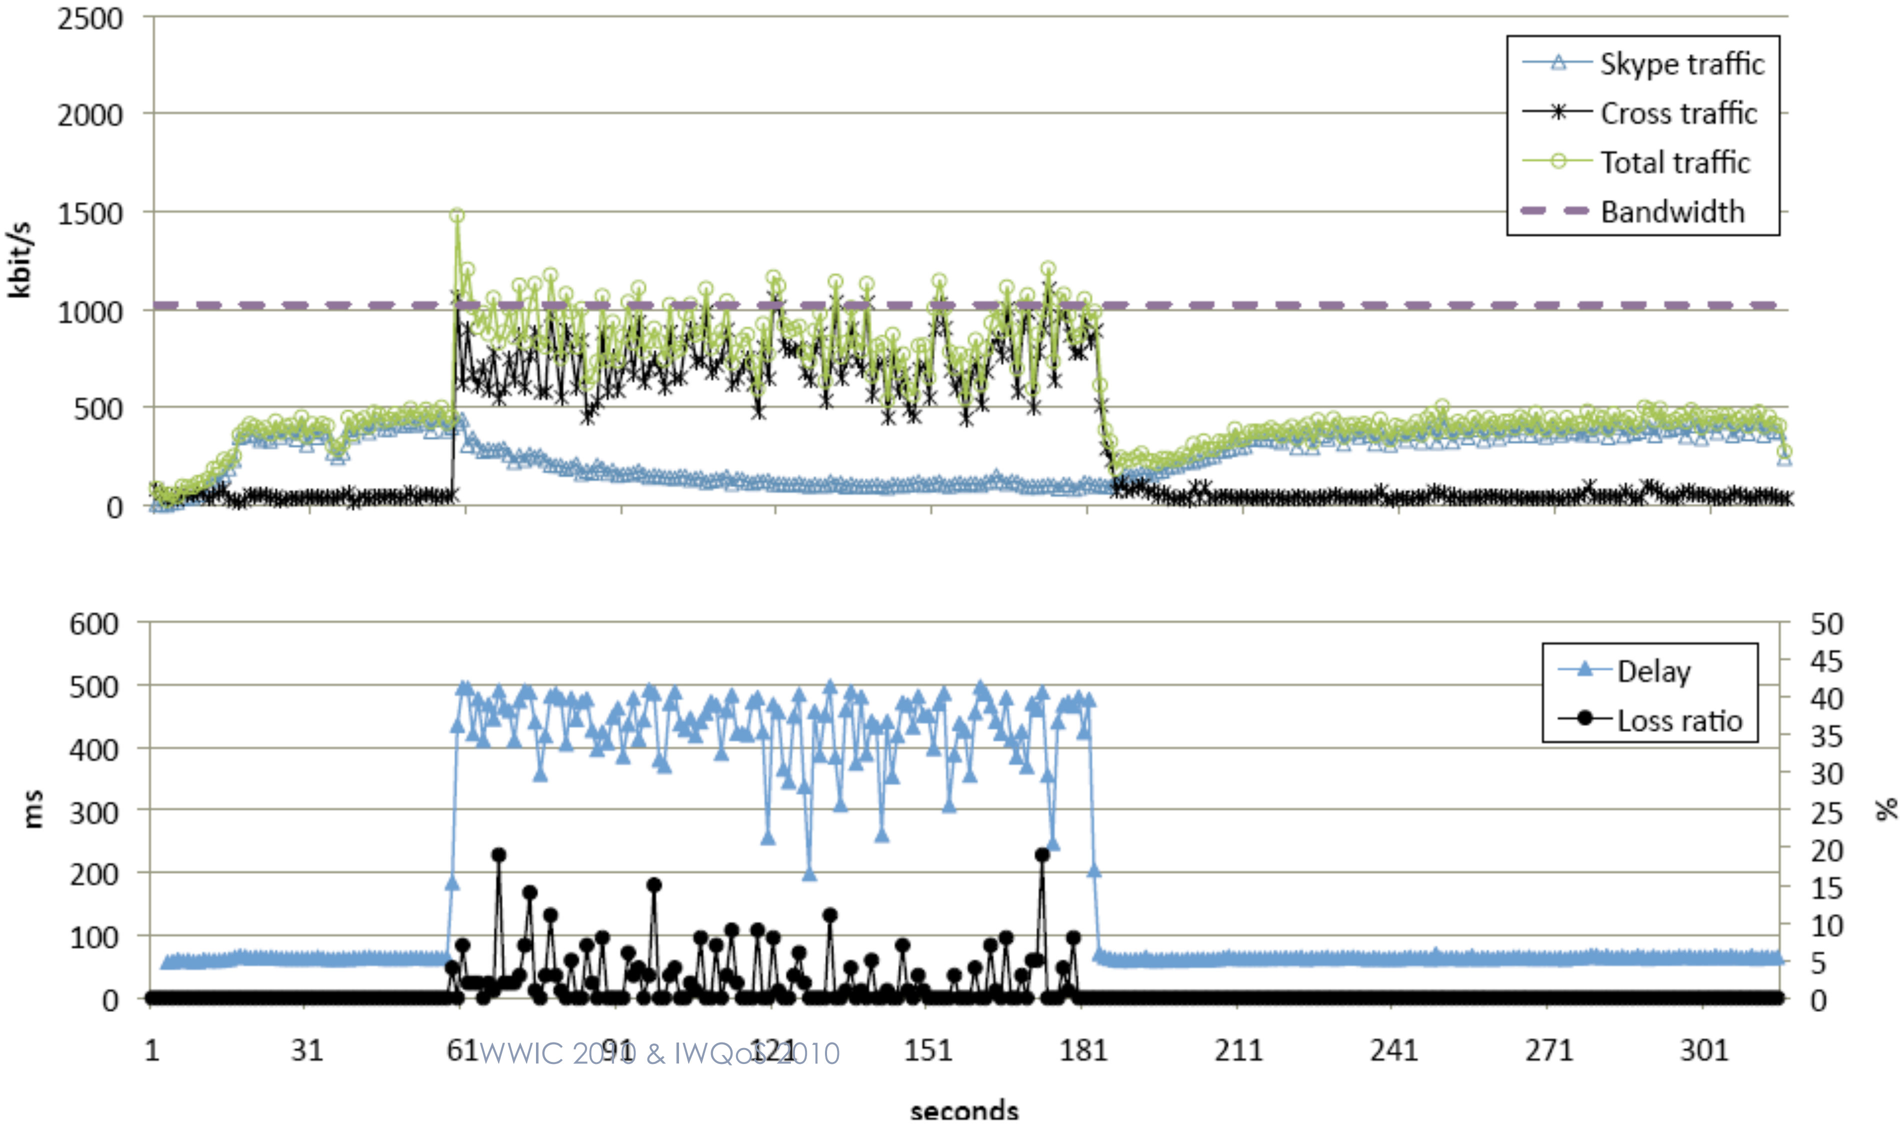

Live

Bittorrent

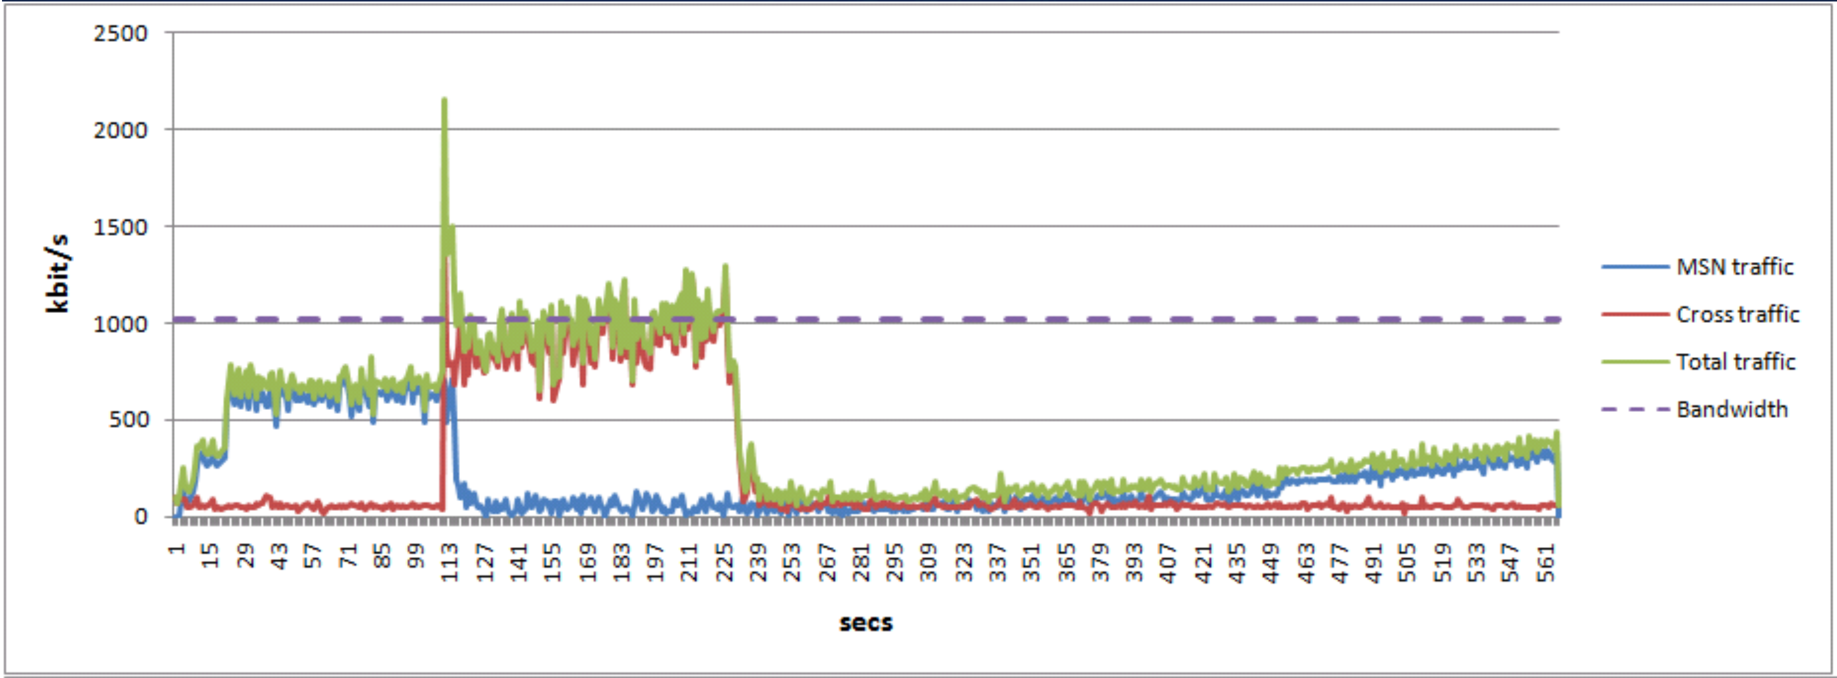

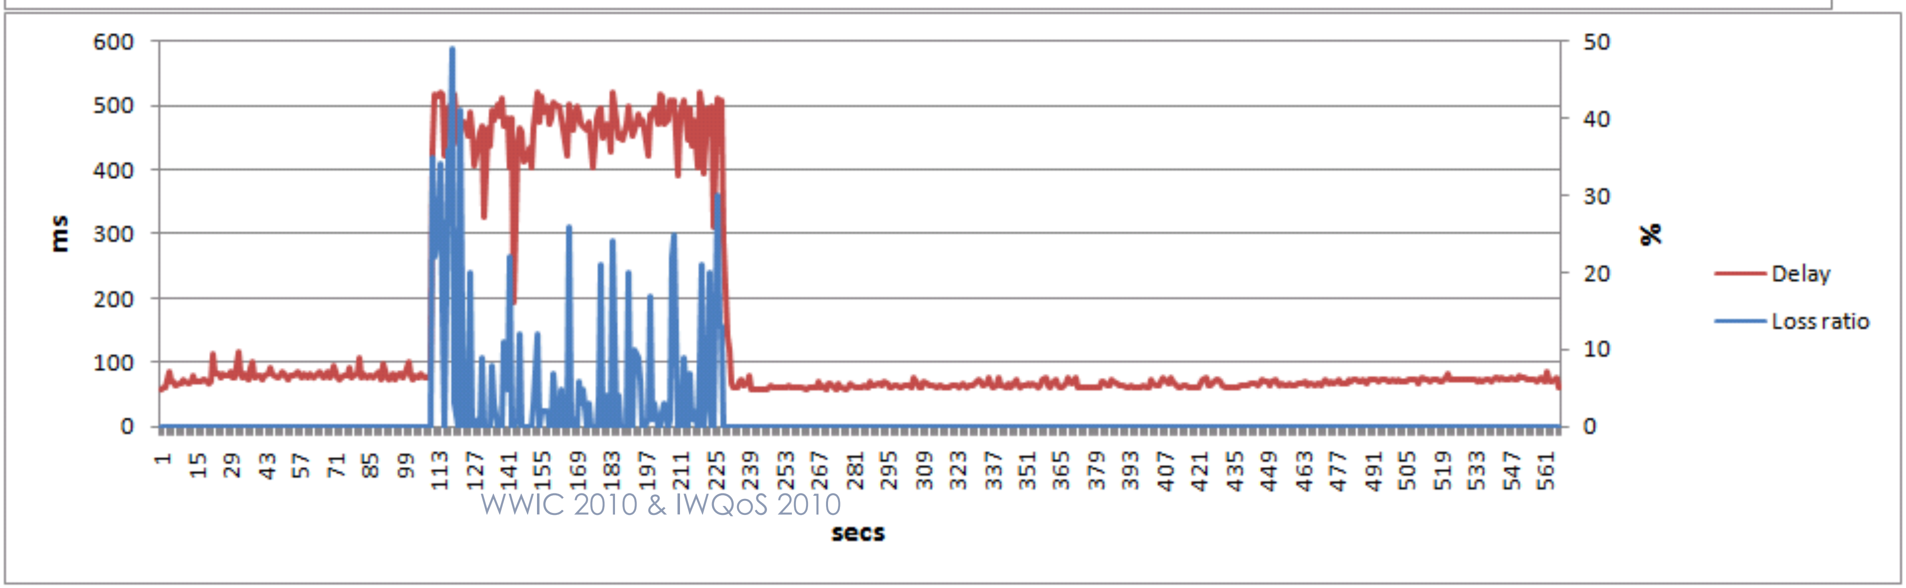

Eyebeam

Bittorrent

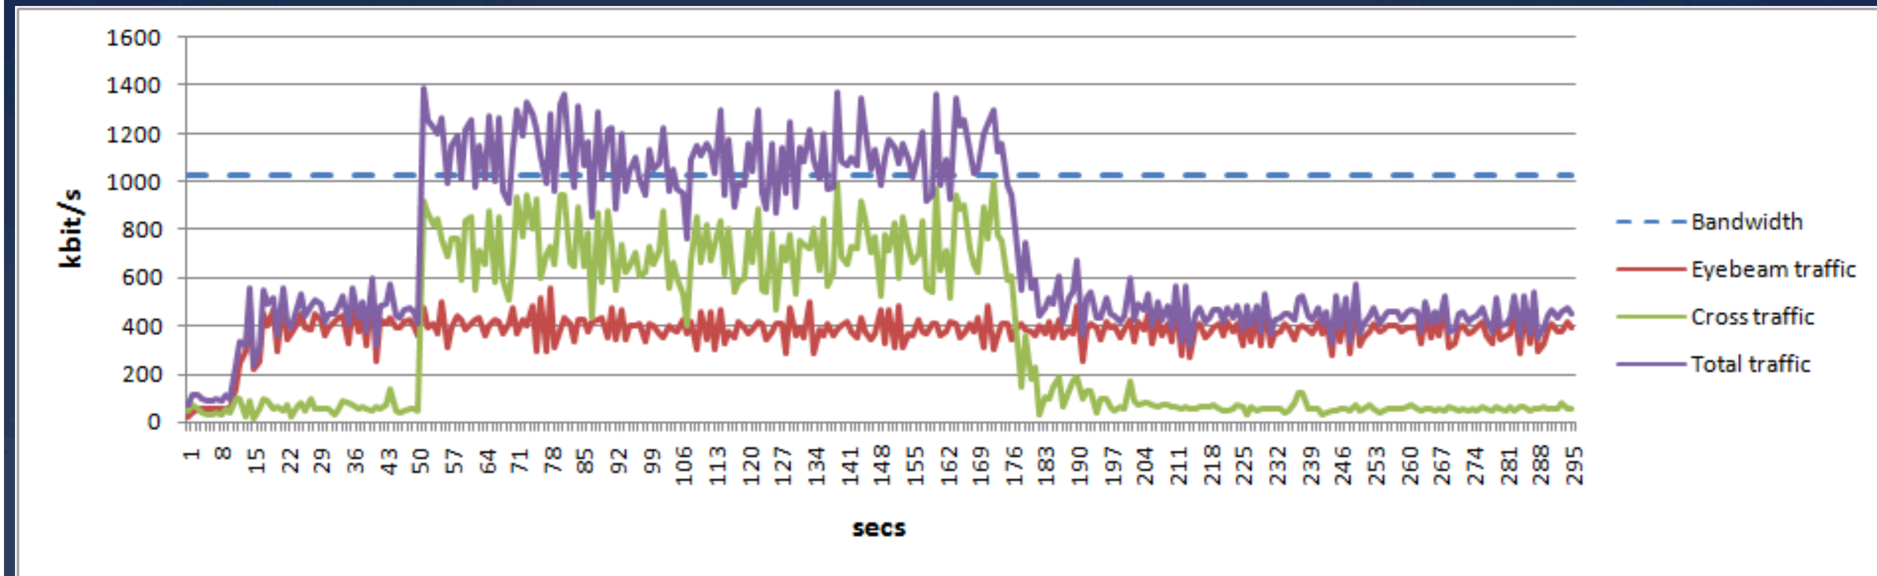

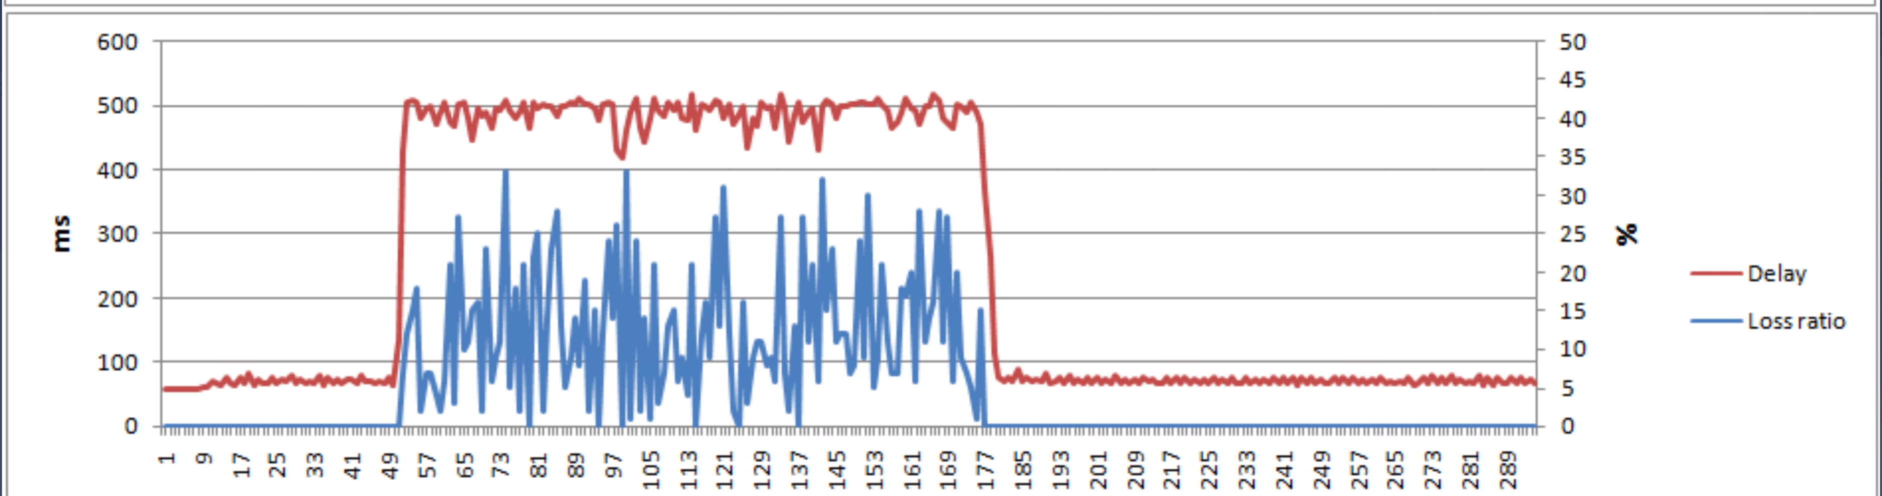

Summary of results

- Skype, Live Messenger, X-Lite and Eyebeam.
  - Skype best:
    - by adapting its codec parameters not only on packet loss but also on RTT and jitter.
    - follow the changes in bandwidth without causing packet loss
  - Eyebeam worst:
    - high fluctuations
    - poor adaptation to bandwidth fluctuations
- Due to limited upstream bandwidth, video clients must have bandwidth adaptation mechanisms and must be able to differentiate between wireless losses and congestion losses

Distributed diagnostics of QoS (and other) problems

Circle of blame

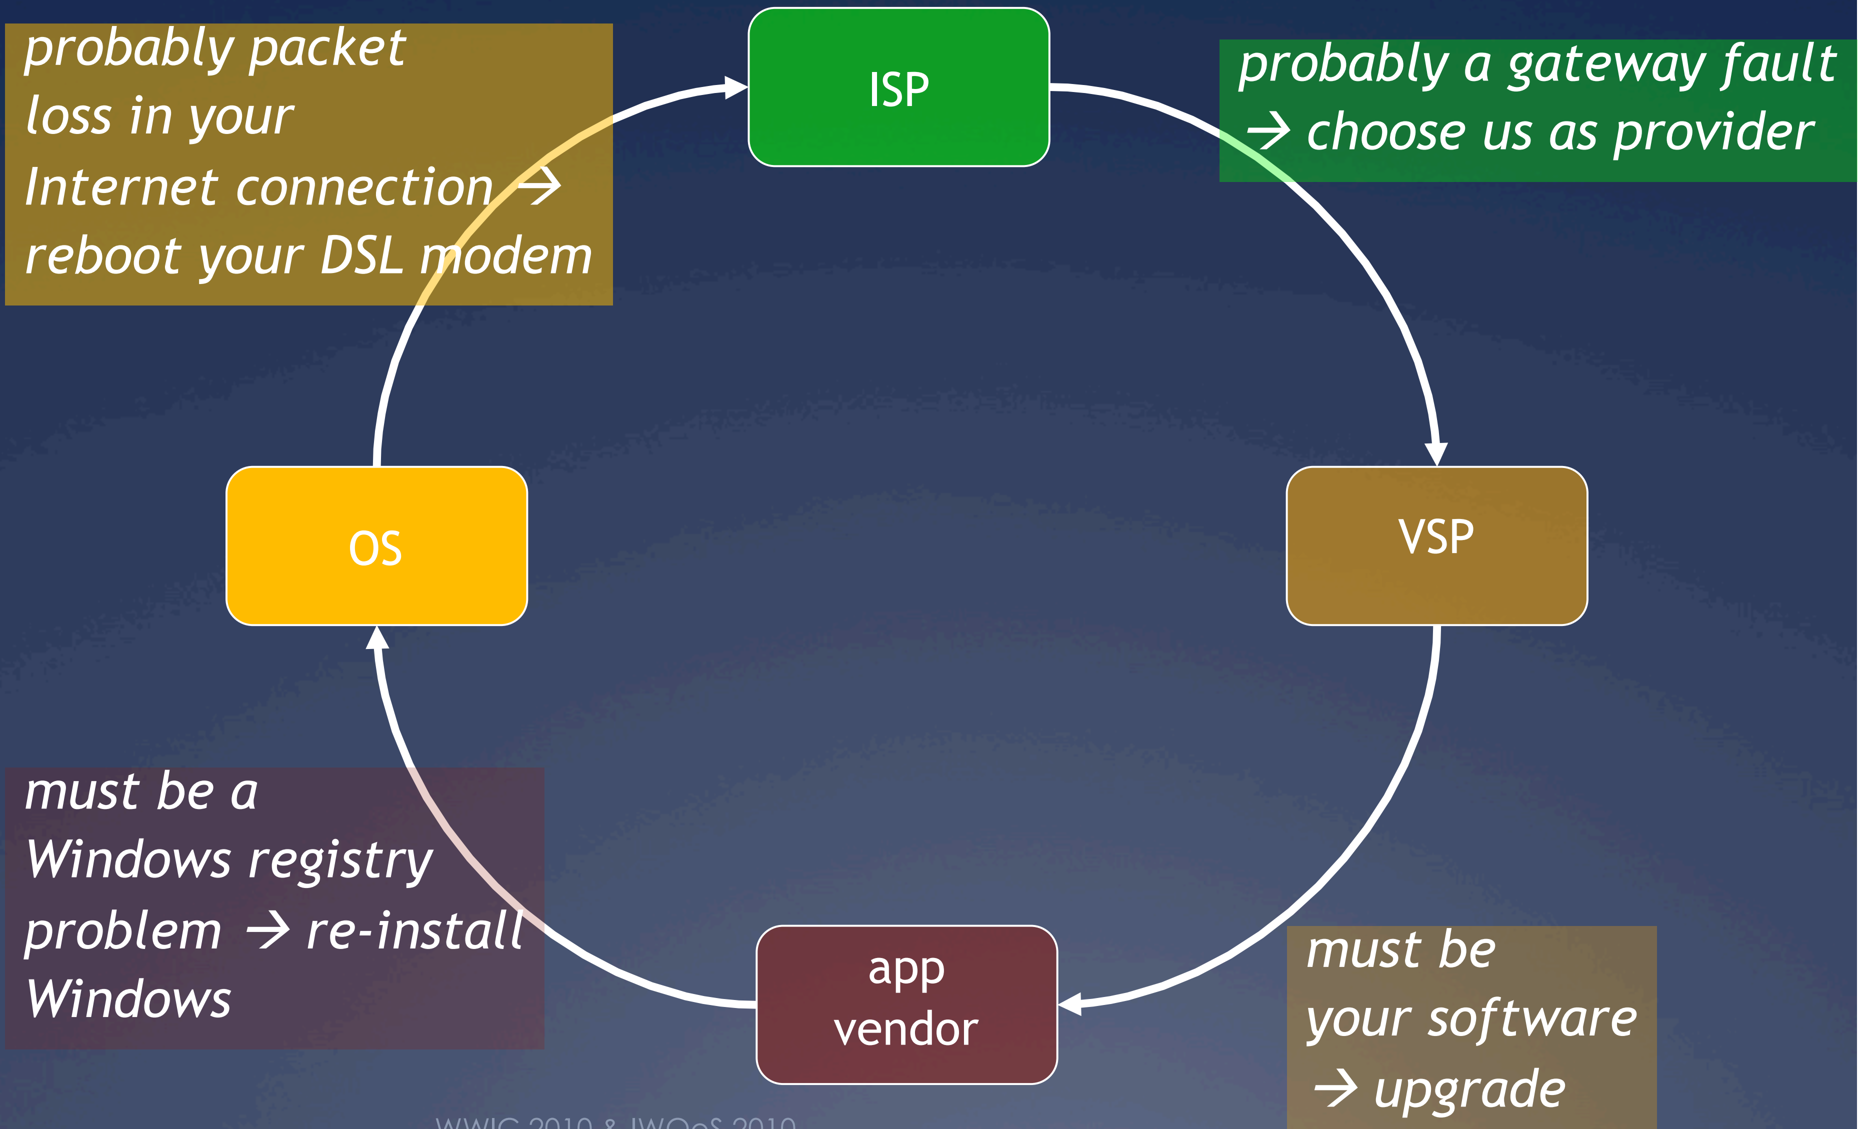

Problems in VoIP systems

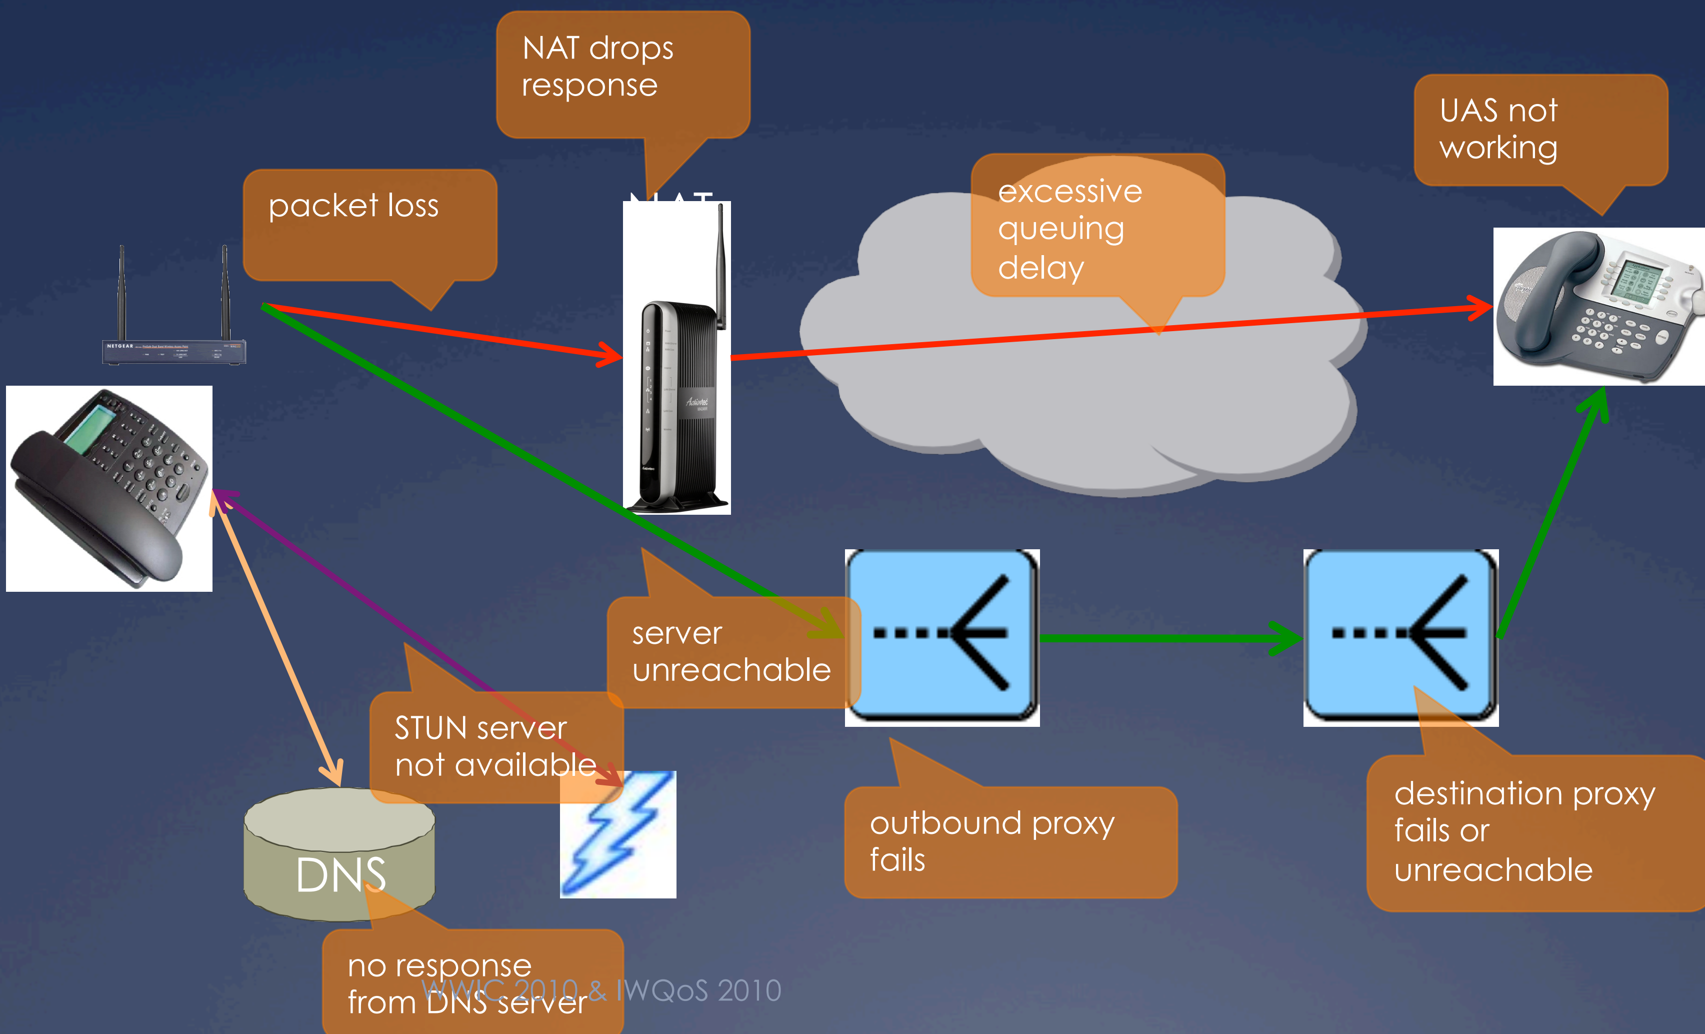

DYSWIS

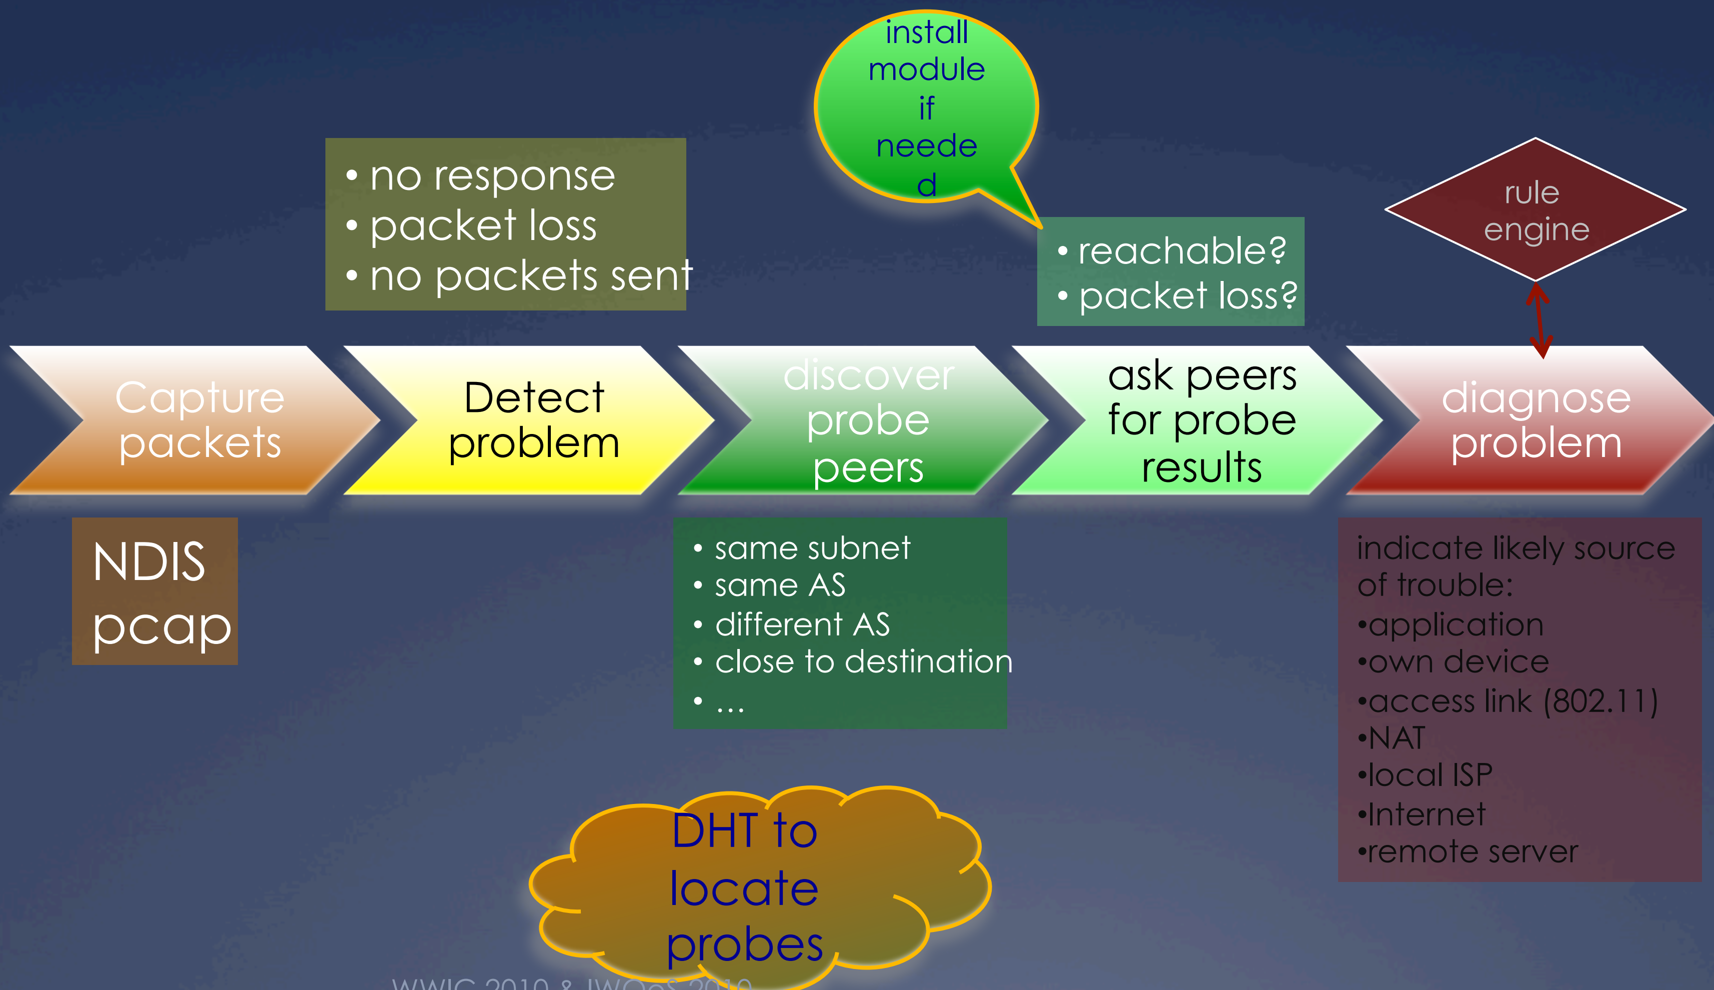
install module if neede d
• no response
• packet loss
• no packets sent
• reachable?
• packet loss?
rule engine
Capture packets
Detect problem
discover probe peers
ask peers for probe results
diagnose problem
NDIS pcap
• same subnet
• same AS
• different AS
• close to destination
• …
indicate likely source of trouble:
•application
•own device
•access link (802.11)
•NAT
•local ISP
•Internet
•remote server
DHT to locate probes

Implementation: system tray

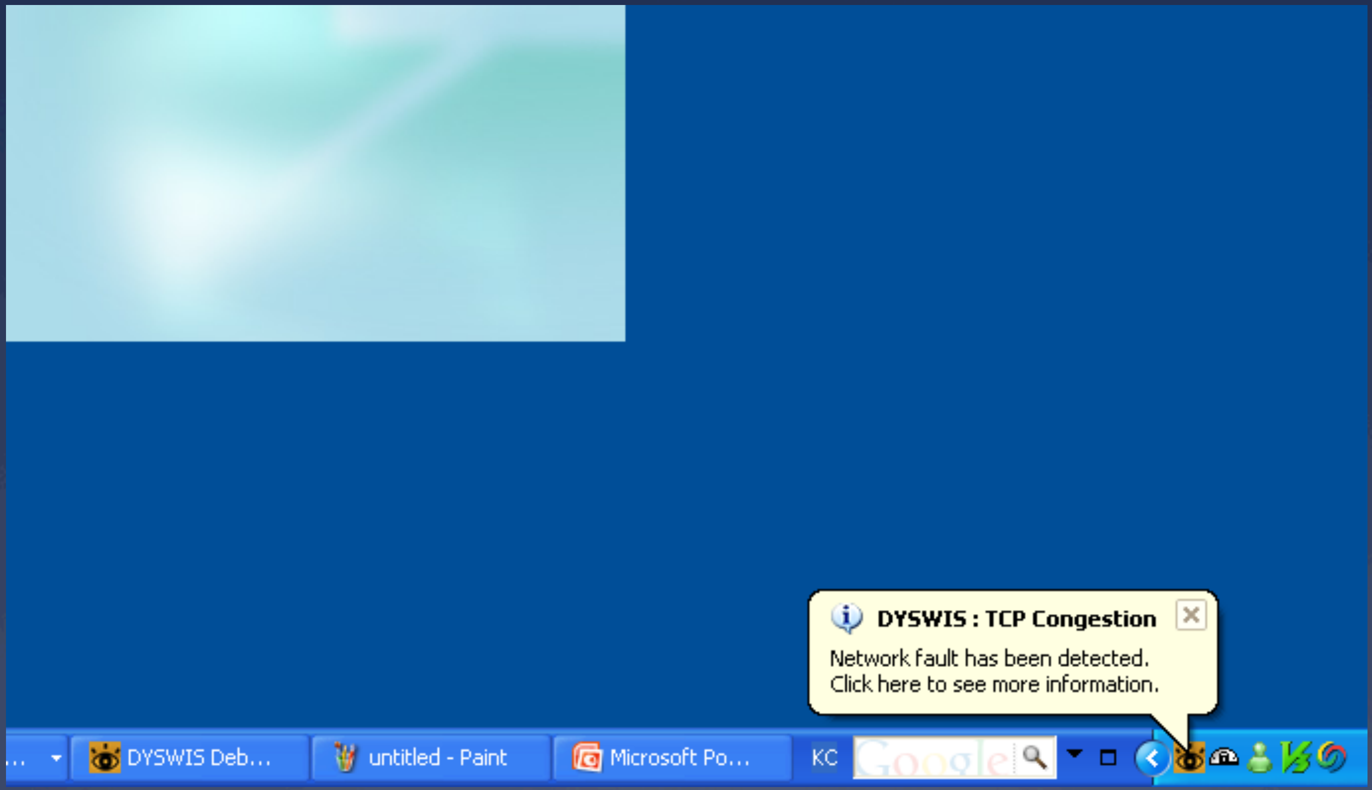

Summary

* QoS = our community's longest running network research topic
  * transition of field from classical performance and queuing theory to security and Internet systems
* Reflect on research role and outcomes
  * are we distilling results or just adding to them?
* How can we identify topics that
  * matter to real users & operators

    AND
  * are amenable to research?