An overview of Quality of Service in networking

Séminaire à Mont-de-Marsan
Département informatique
6 décembre 2013

C. Pham
Université de Pau et des Pays de l'Adour
http://www.univ-pau.fr/~cpham
Congduc.Pham@univ-pau.fr

Outline

- What is Quality of Service?
- How QoS is realized in current Internet
- How it can be optimally realized
- How it can be realistically realized
- 2 leading technologies for QoS
  - DiffServ
  - MPLS

The big-bang of the Internet

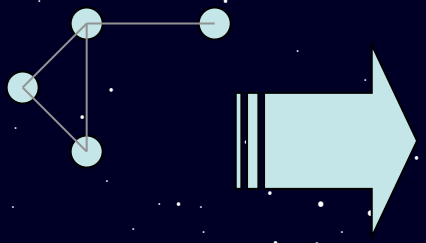

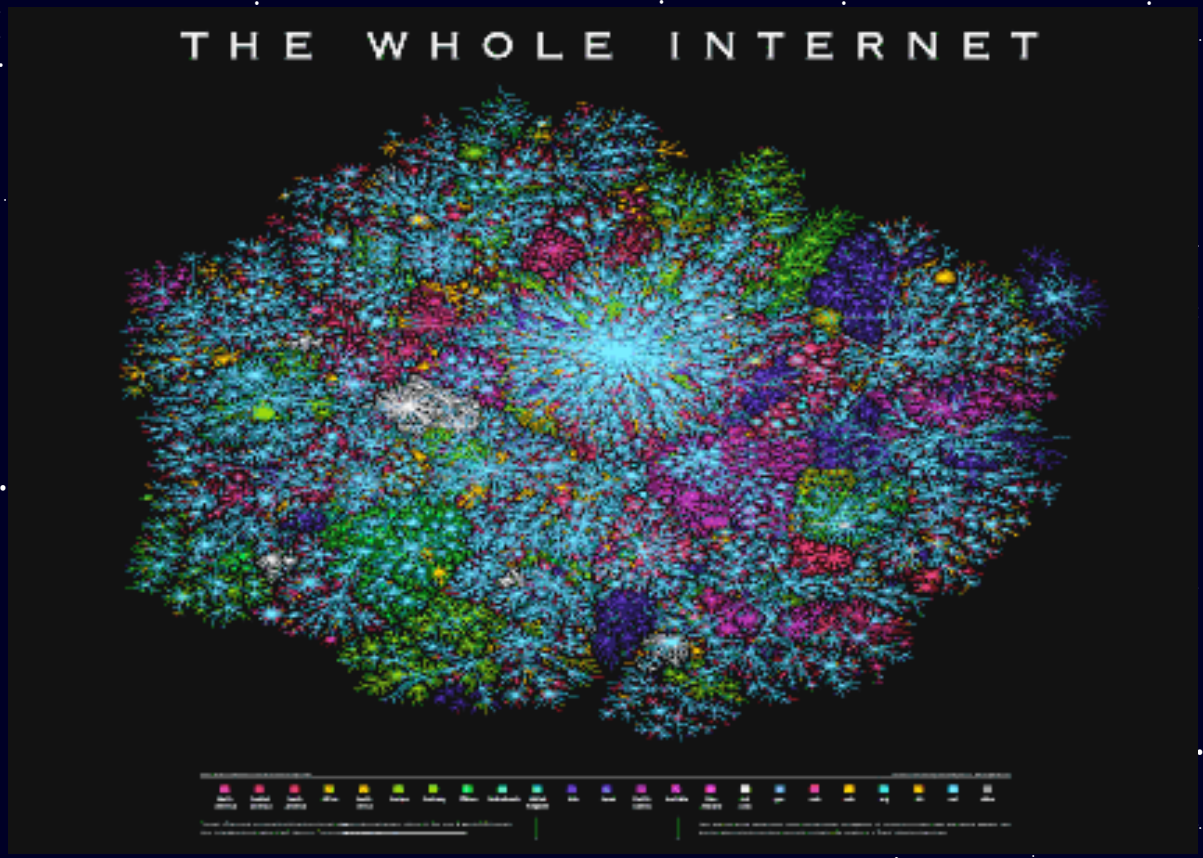

Internet host

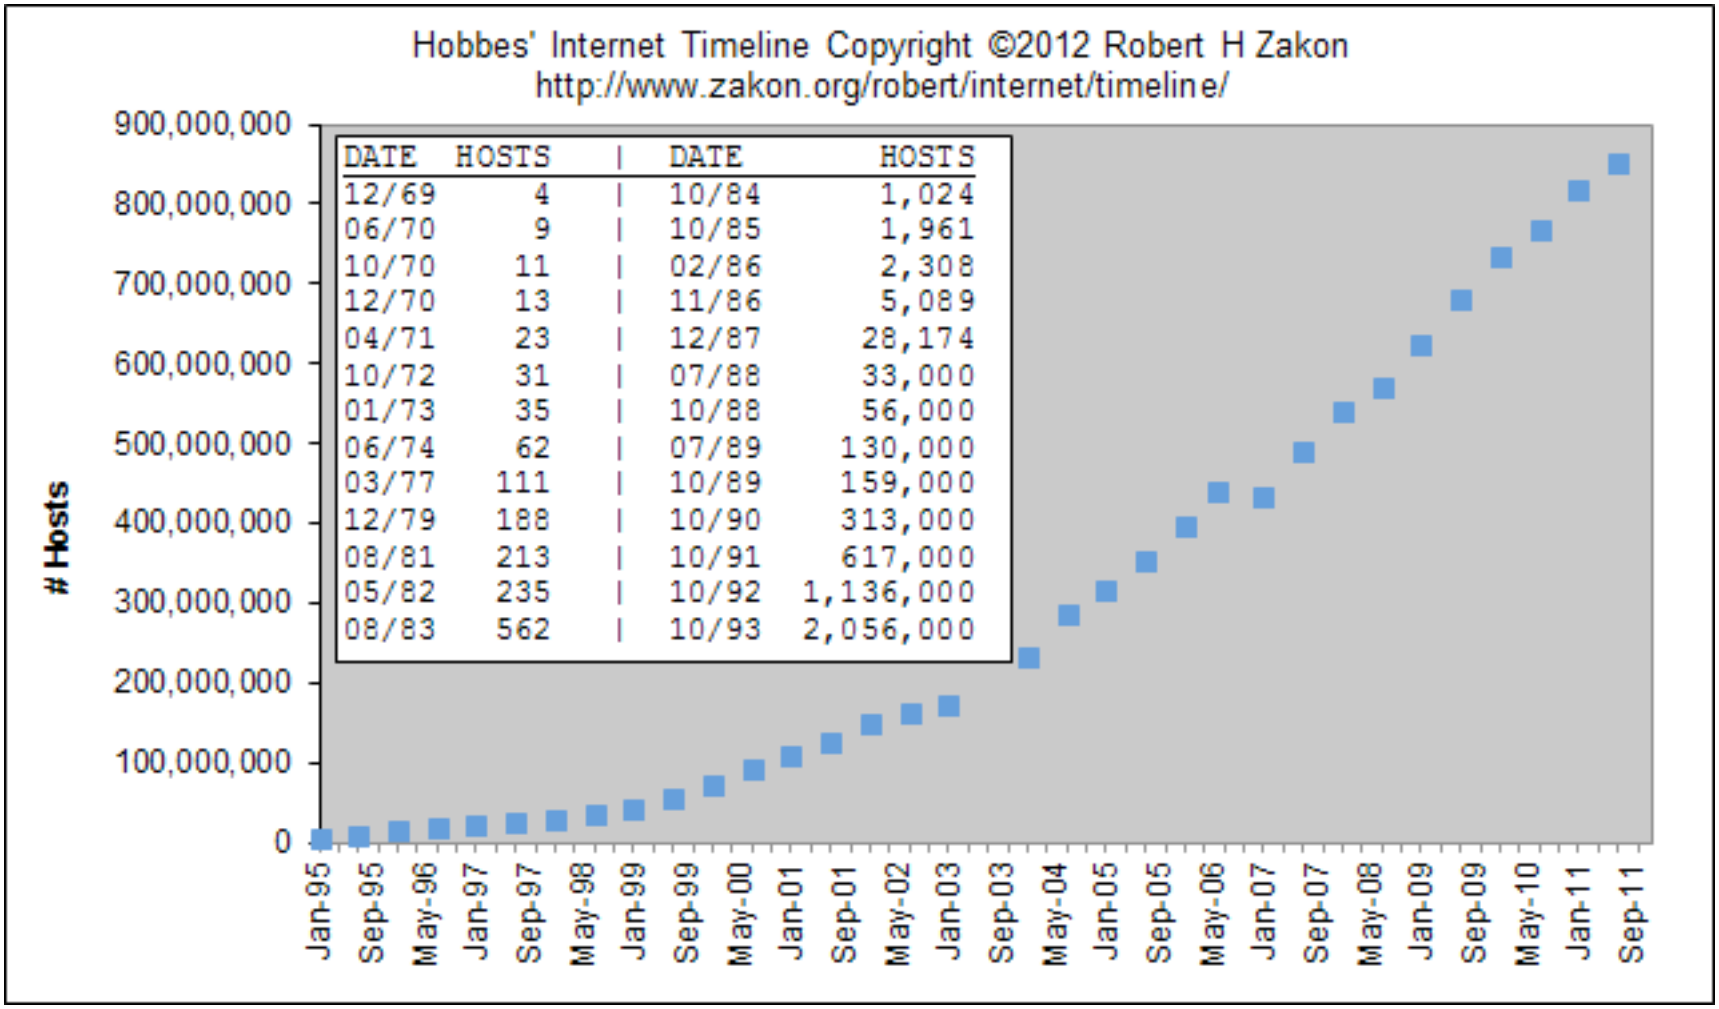

DATE	HOSTS	DATE	HOSTS
12/69	4	10/84	1,024
06/70	9	10/85	1,961
10/70	11	02/86	2,308
12/70	13	11/86	5,089
04/71	23	12/87	28,174
10/72	31	07/88	33,000
01/73	35	10/88	56,000
06/74	62	07/89	130,000
03/77	111	10/89	159,000
12/79	188	10/90	313,000
08/81	213	10/91	617,000
05/82	235	10/92	1,136,000
08/83	562	10/93	2,056,000

of www sites

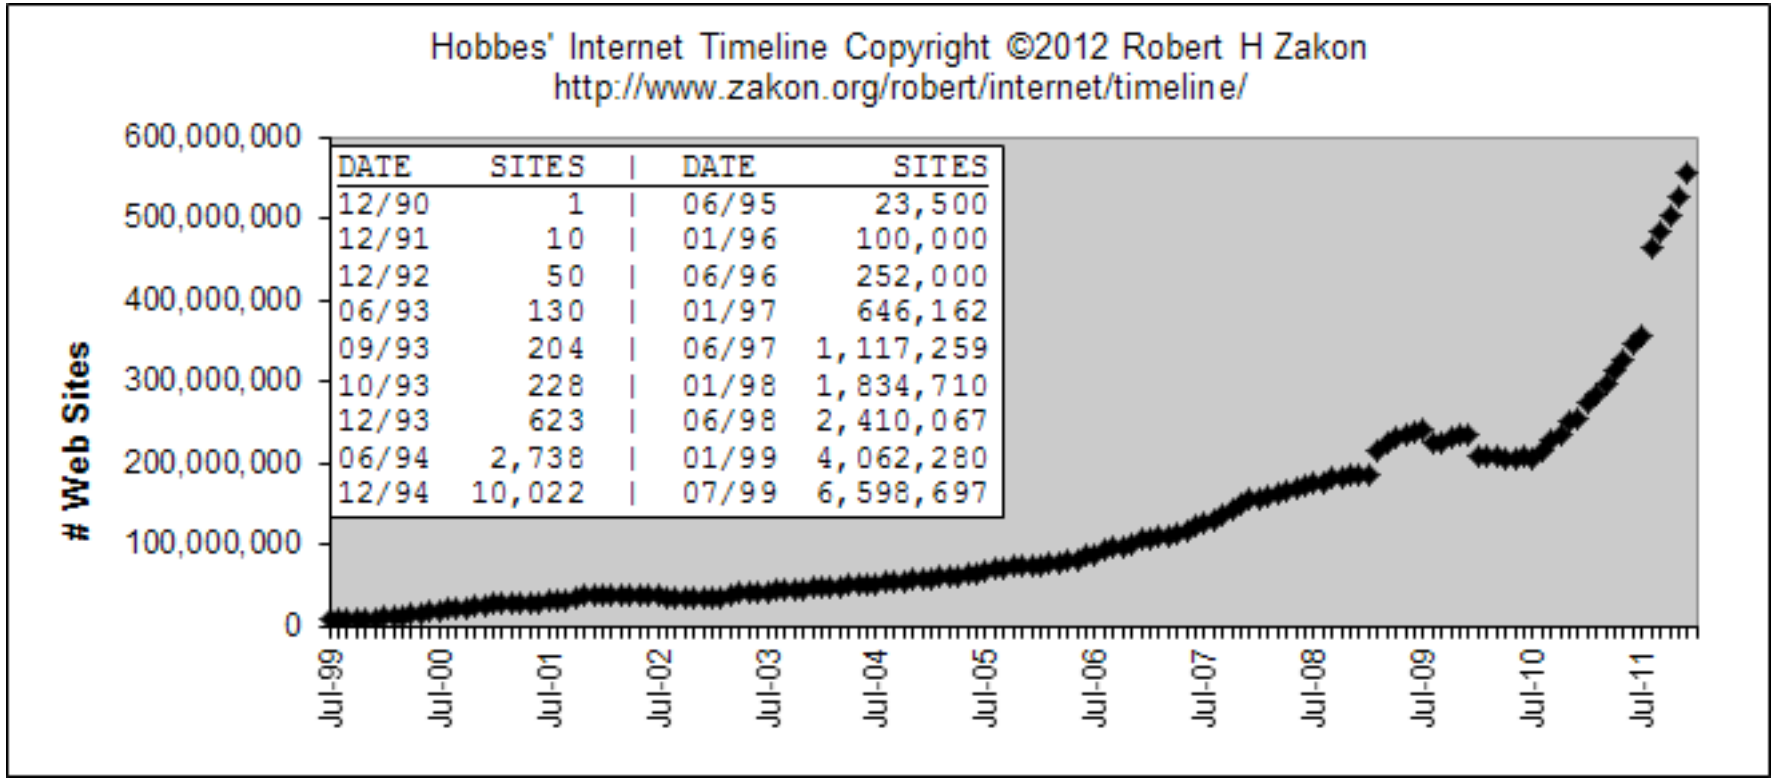

DATE	SITES	DATE	SITES
12/90	1	06/95	23,500
12/91	10	01/96	100,000
12/92	50	06/96	252,000
06/93	130	01/97	646,162
09/93	204	06/97	1,117,259
10/93	228	01/98	1,834,710
12/93	623	06/98	2,410,067
06/94	2,738	01/99	4,062,280
12/94	10,022	07/99	6,598,697

of facebook account

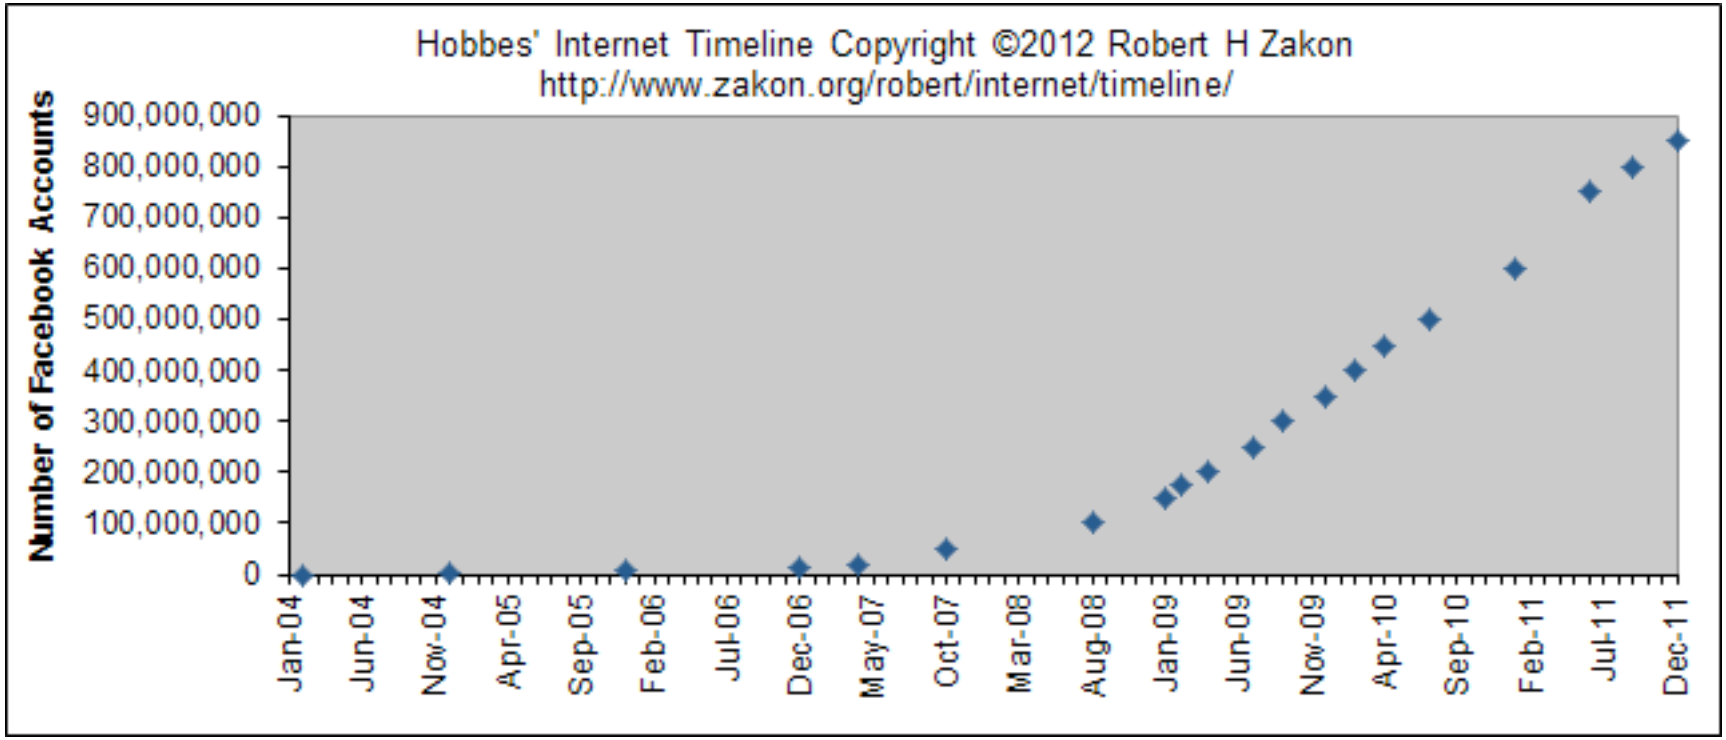

1st revolution: going optical

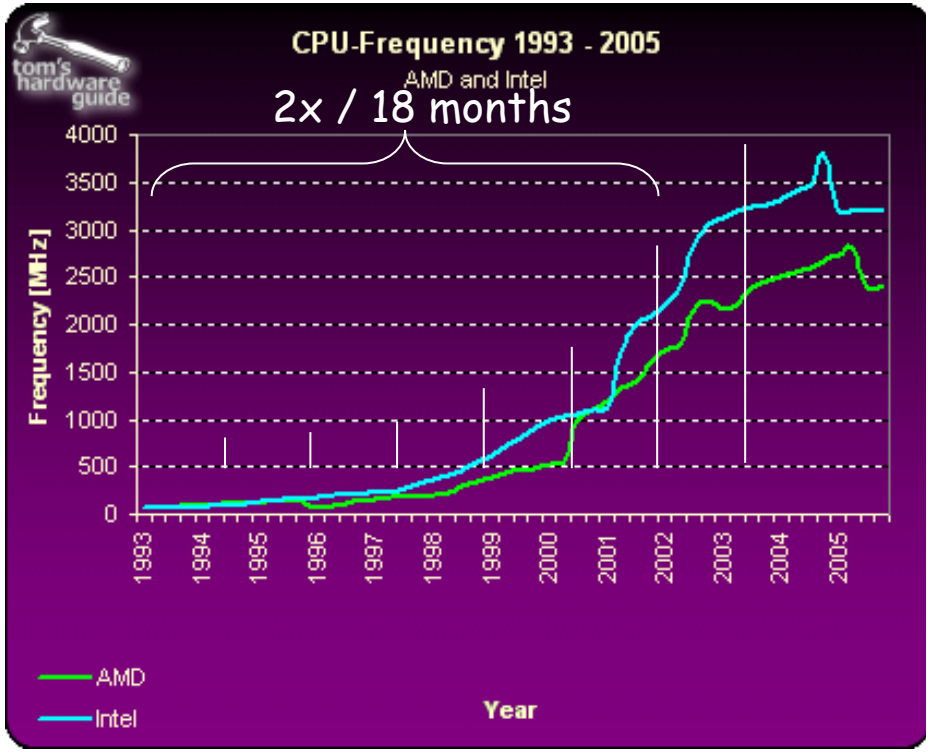

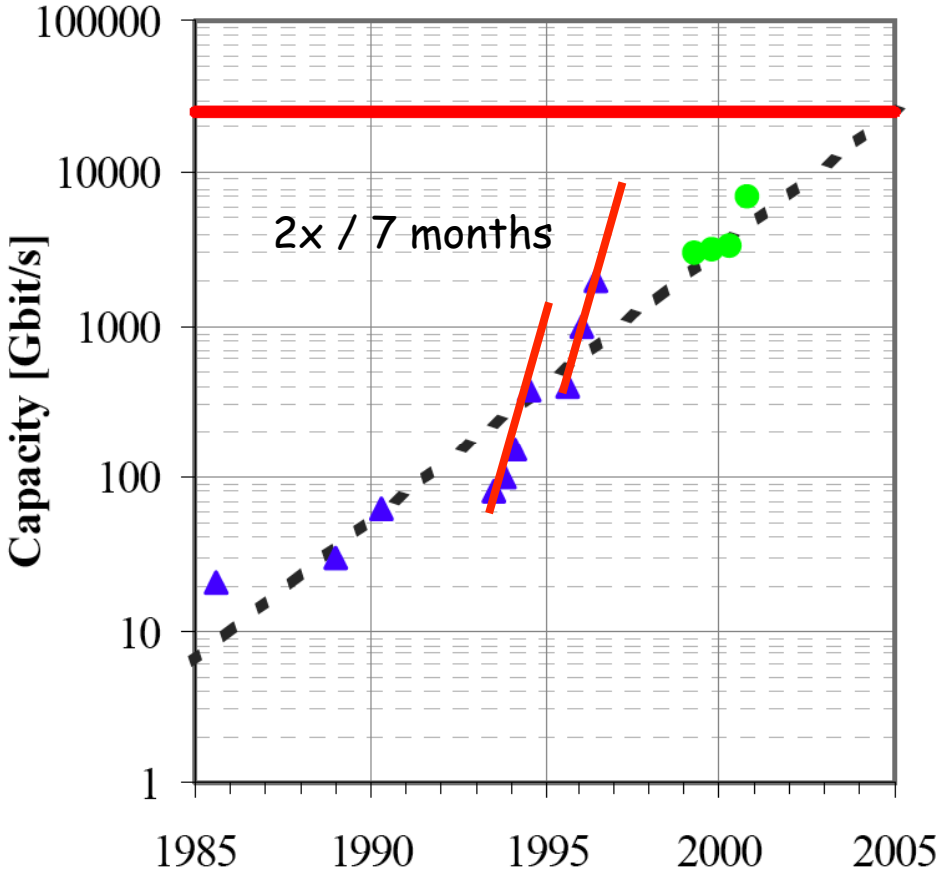

Source « Optical fibers for Ultra-Large Capacity Transmission » by J. Grochocinski

DWDM, bandwidth for free?

DWDM: Dense Wavelength Division Multiplexing

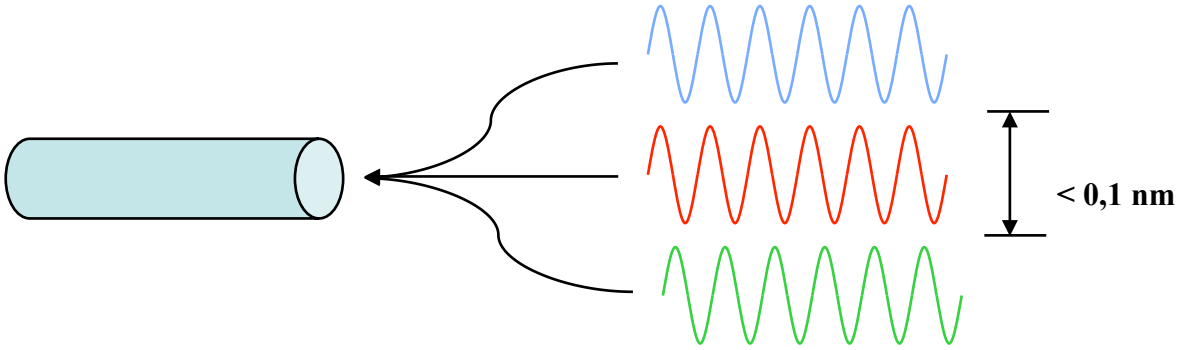

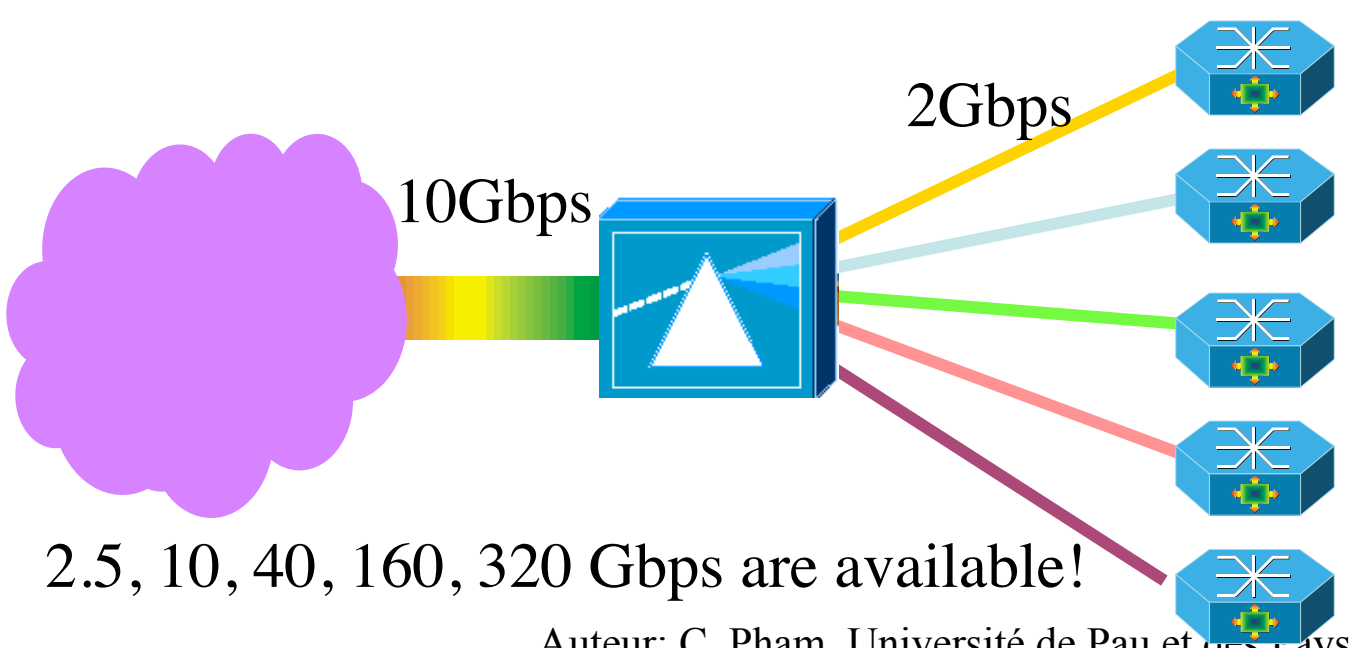

2.5, 10, 40, 160, 320 Gbps are available!

Internet core is 100% optical

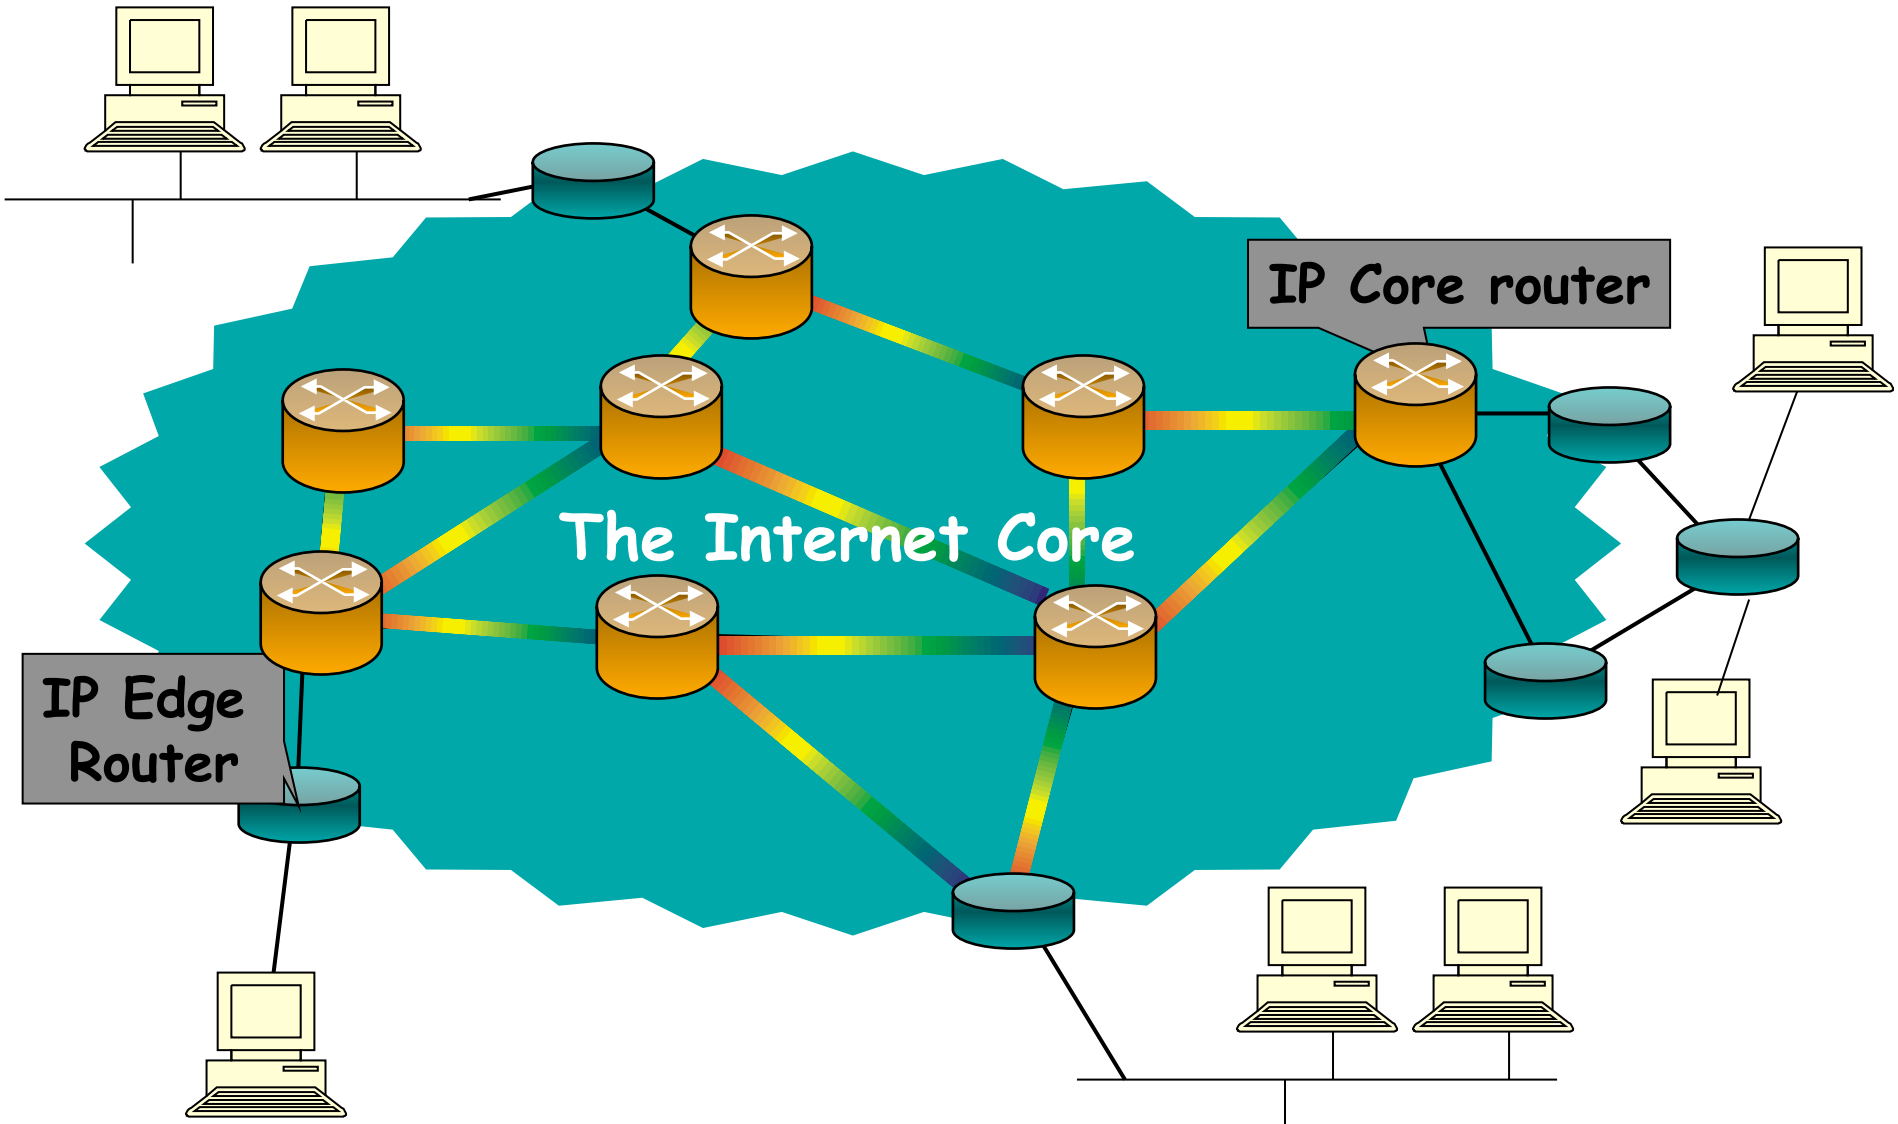

Fibers everywhere?

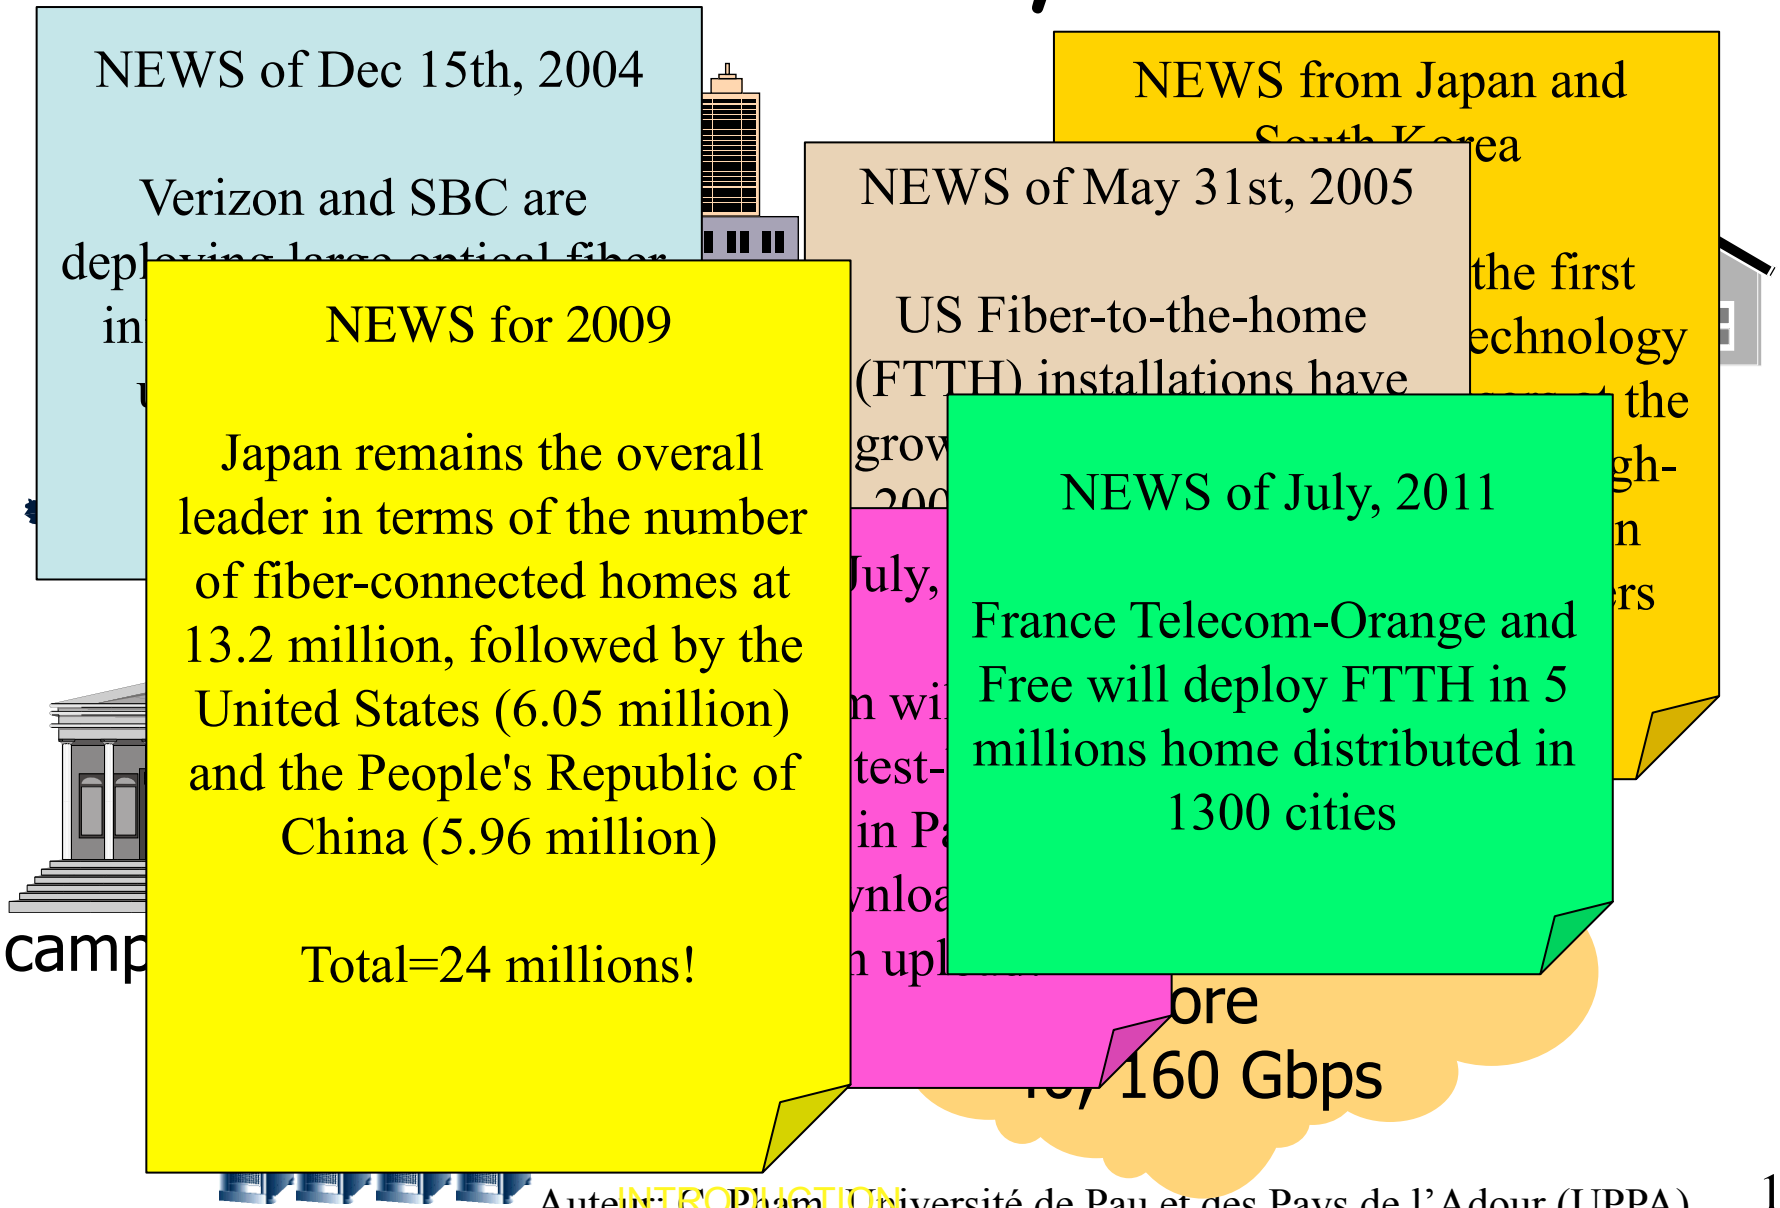

Handle big data!

Towards small, smart devices!

MCU sales

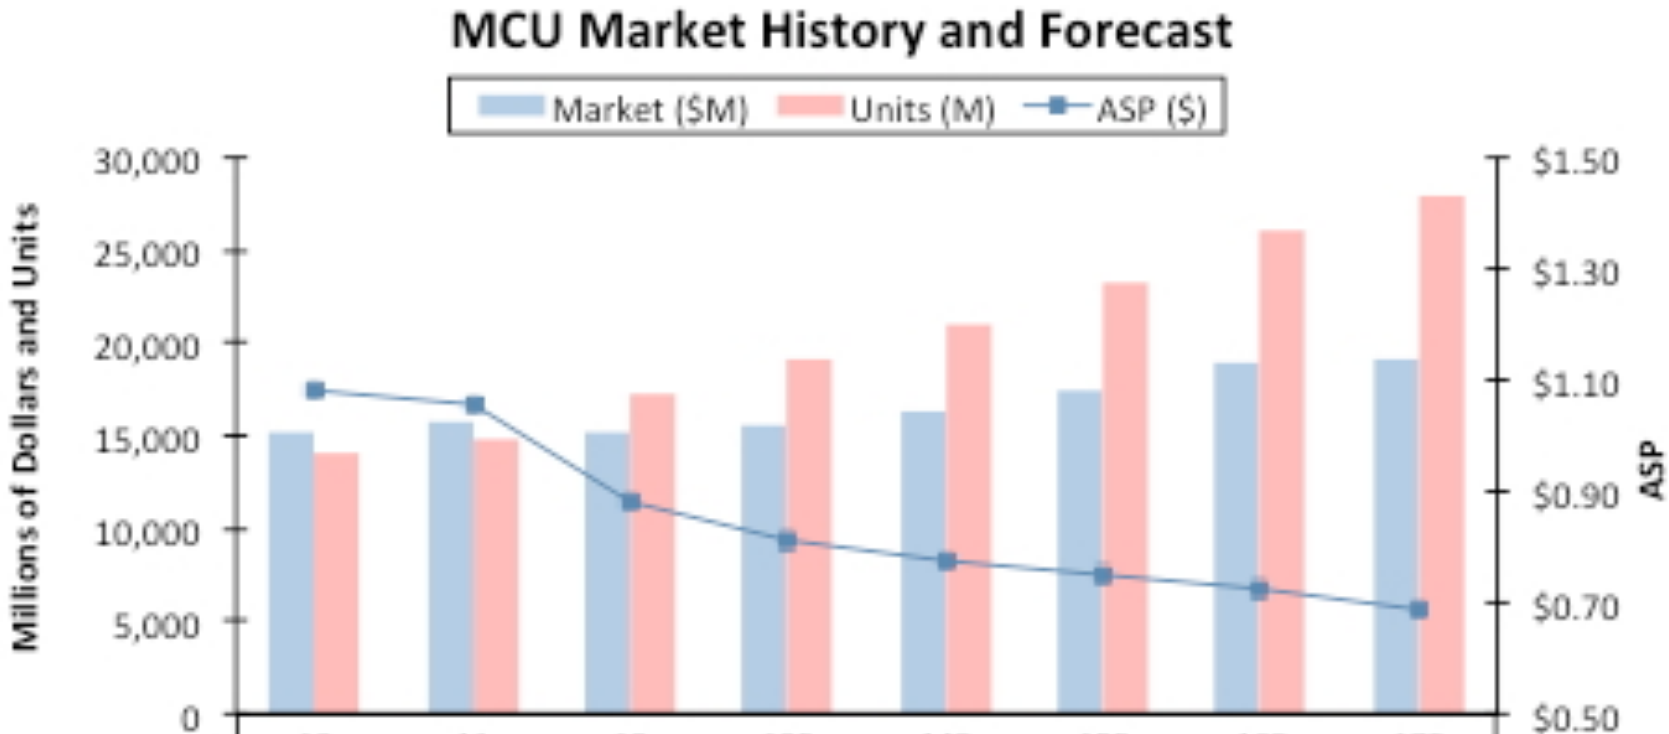

	10	11	12	13F	14F	15F	16F	17F
Market ($M)	15,166	15,716	15,168	15,518	16,281	17,451	18,970	19,208
Units (M)	14,029	14,891	17,271	19,071	21,080	23,314	26,104	27,970
ASP ($)	1.08	1.06	0.88	0.81	0.77	0.75	0.73	0.69

Source: IC Insights

2nd revolution: Wireless Networks

- WiFi, WiMax
- BlueTooth, ZigBee, IrDA…
- GSM, GPRS, EDGE, UMTS, 3G, 4G,…

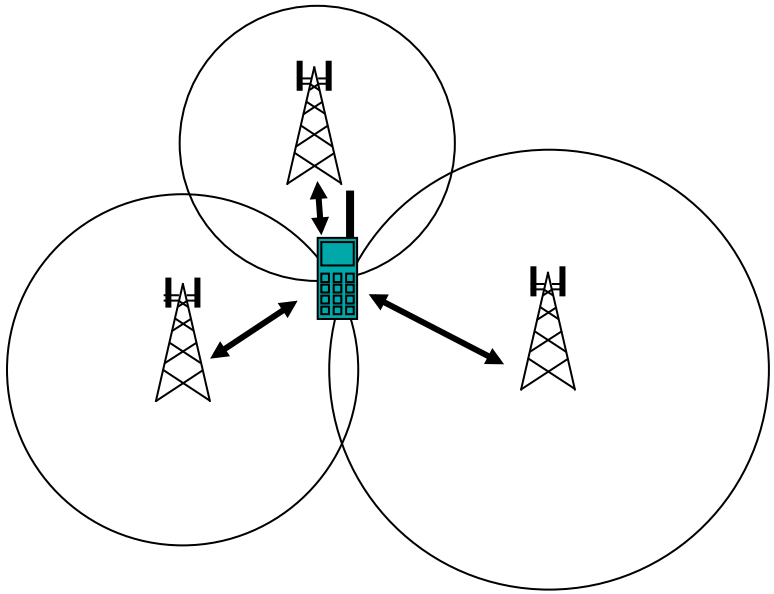

Wireless communication made easy!

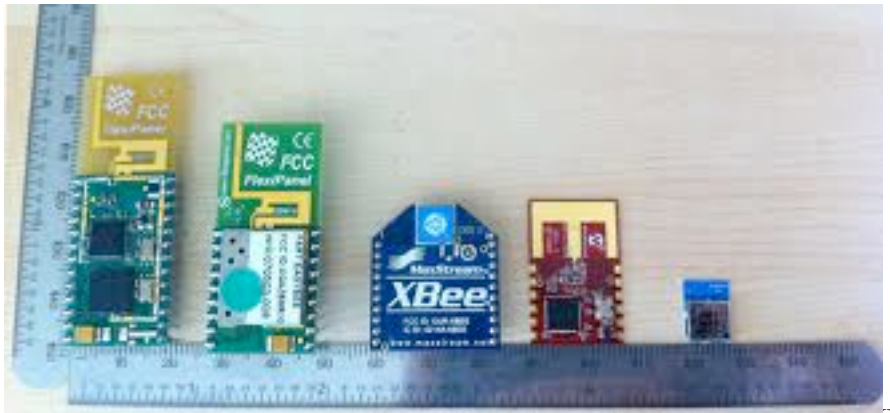

Internet of Things/ M2M

- Native communication:

- Added communication
  - Active communication

  - Passive communication

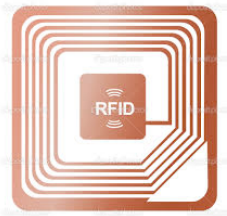

Digital (Ubiquitous) World

Wi Fi

wimax

ZigBee® Member

3G

4G

100mbps-1GBps

Lte™

WiMedia ALLIANCE

INTERNET

MOBILE INTERNET

PERVASIVE SYSTEMS

SENSORS

INTERNET OF THINGS

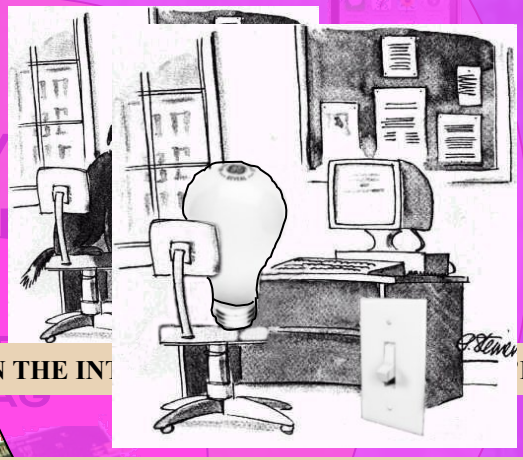

ON THE INTERNET NOBODY KNOWS YOU'RE A DOG!

ON THE INTERNET NOBODY KNOWS YOU'RE A LIGHT BULB!

Towards all IP

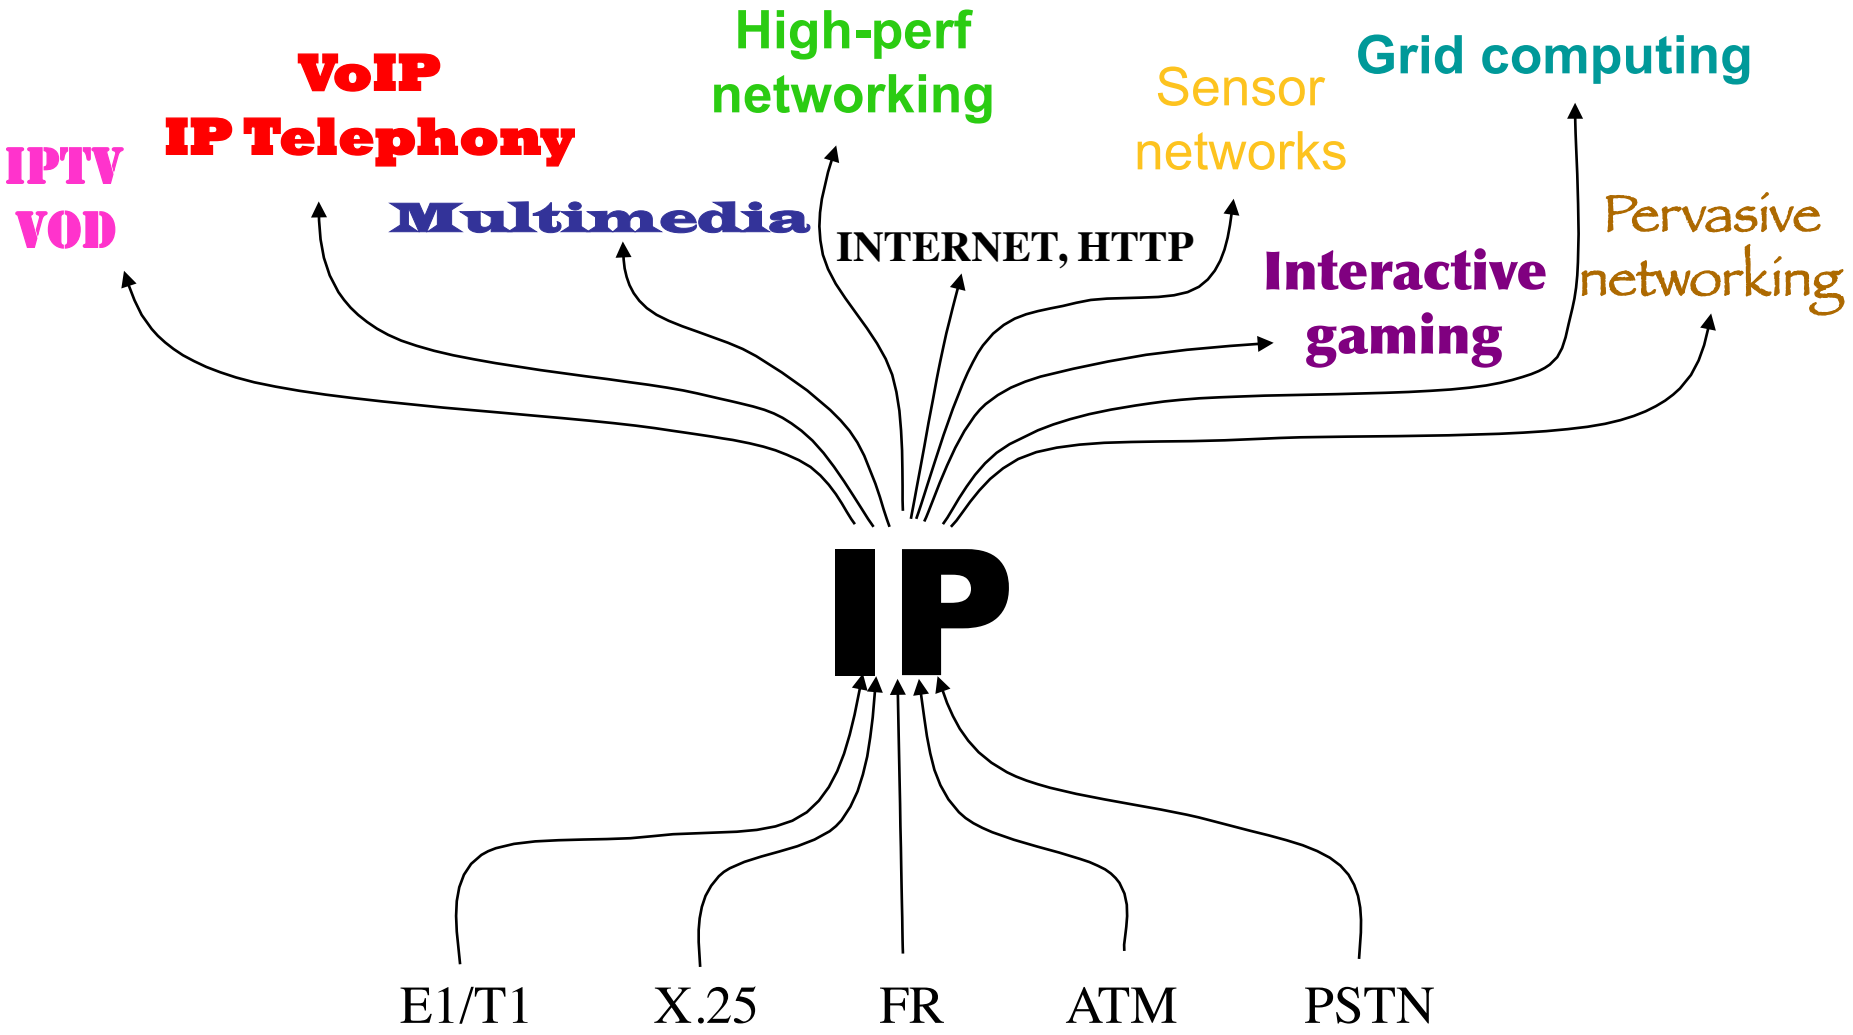

IP for IoT

- IPv6 gives plenty of addresses
- 6LowPan adapts IPv6 to resource-constrained devices

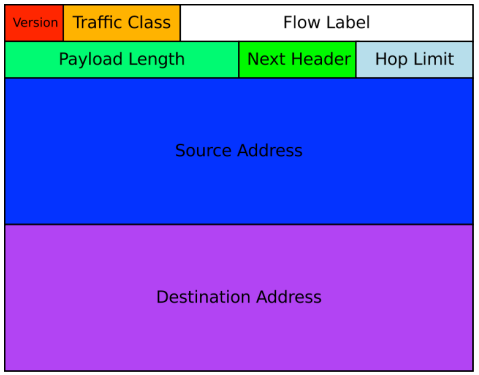

40 bytes

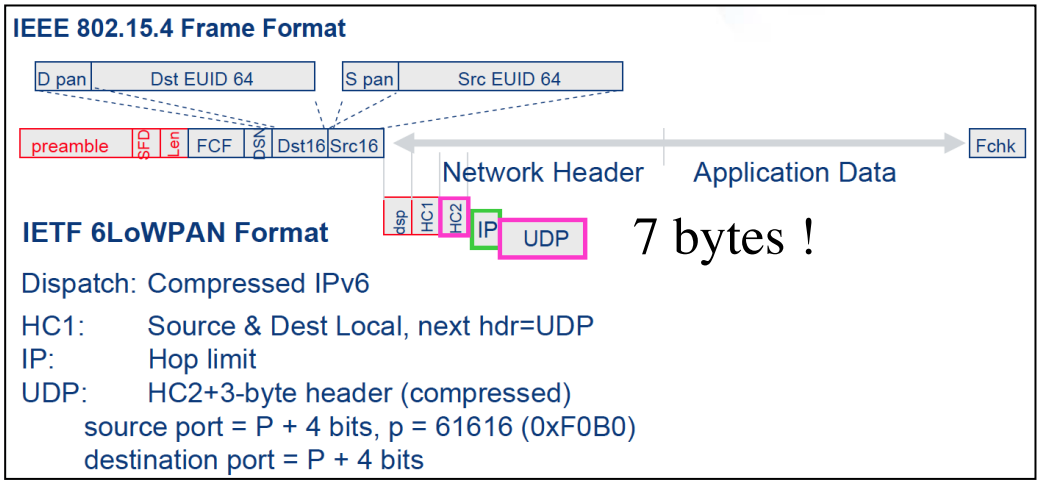

- CoAP provides interoperability

IP connectivity can take many (unexpected) forms!

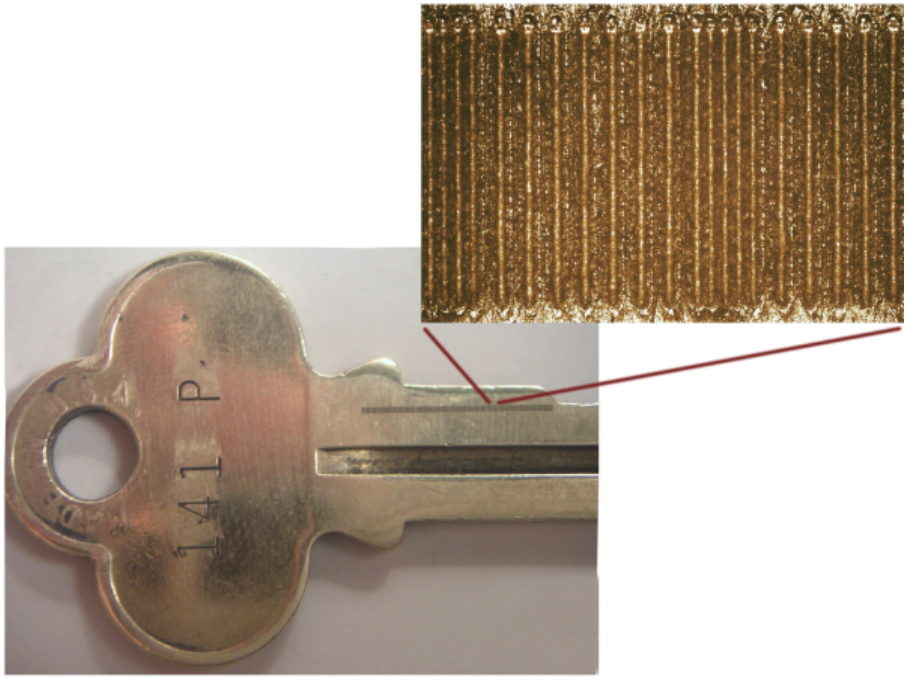

Be Smart* !

- Smart…
  - City, Building, Road, Traffic
  - Agriculture
  - Farming
  - Environment: Water, Forest
  - Energy, Electricity Grid
  - Vehicule & Transportation
  - Transport & Logistic
  - Surveillance, security, safety
  - …

Cities

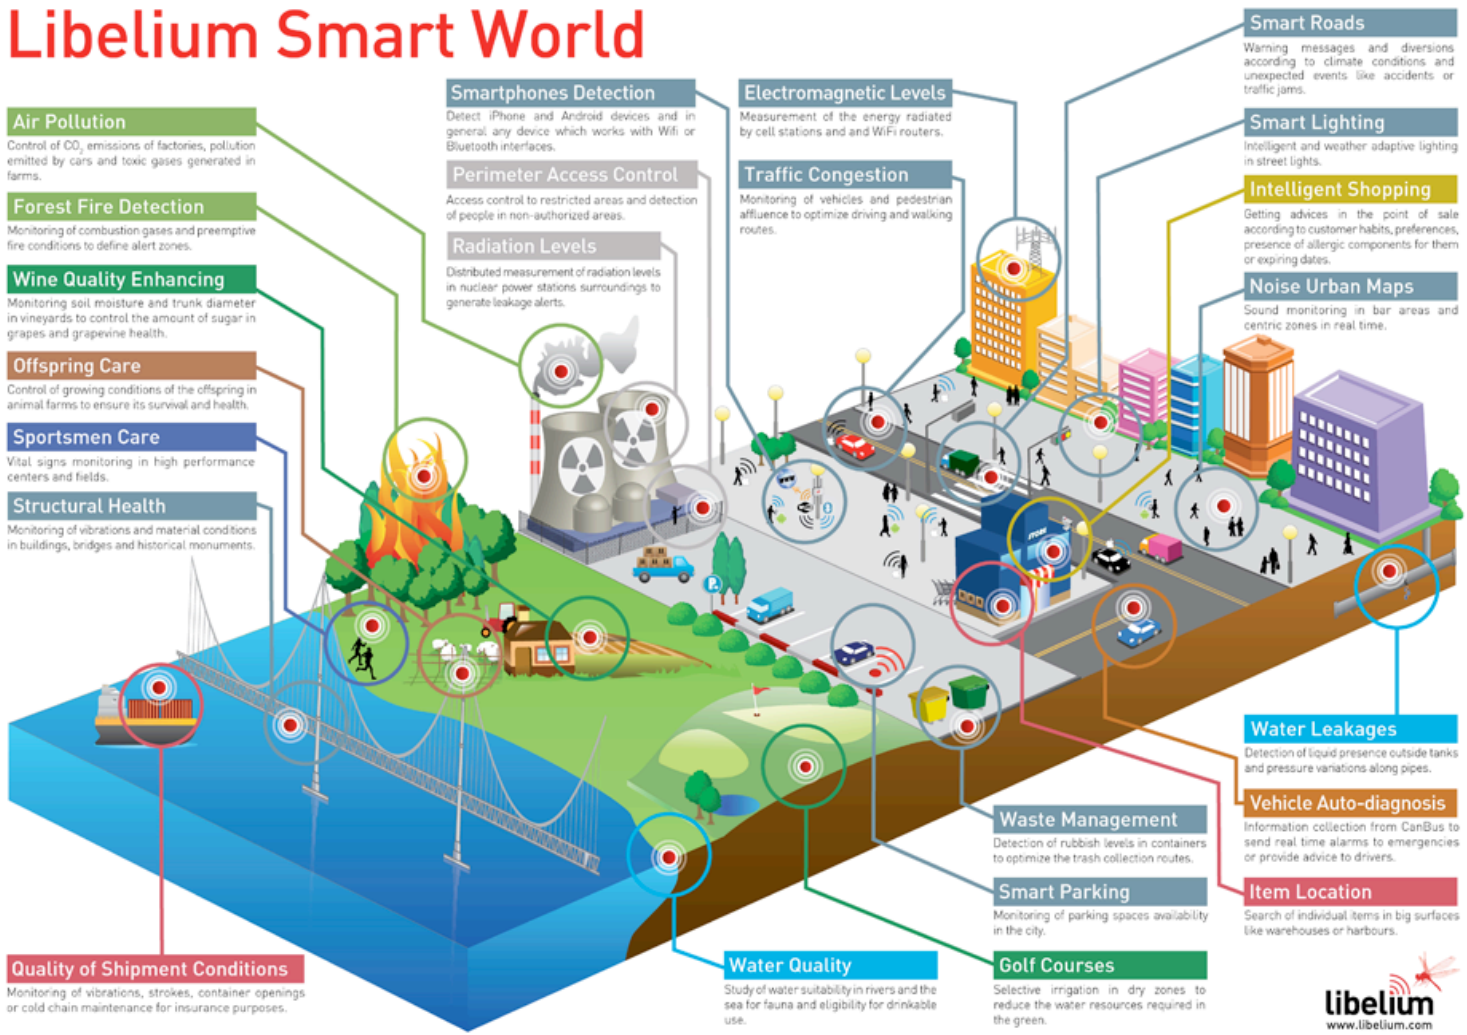

HTTP://WWW.LIBELIUM.COM/TOP_50_IOT_SENSOR_APPLICATIONS_RANKING/#SHOW_INFOGRAPHIC

SmartSantander

www.smartsantander.eu

SmartSantander test-bed

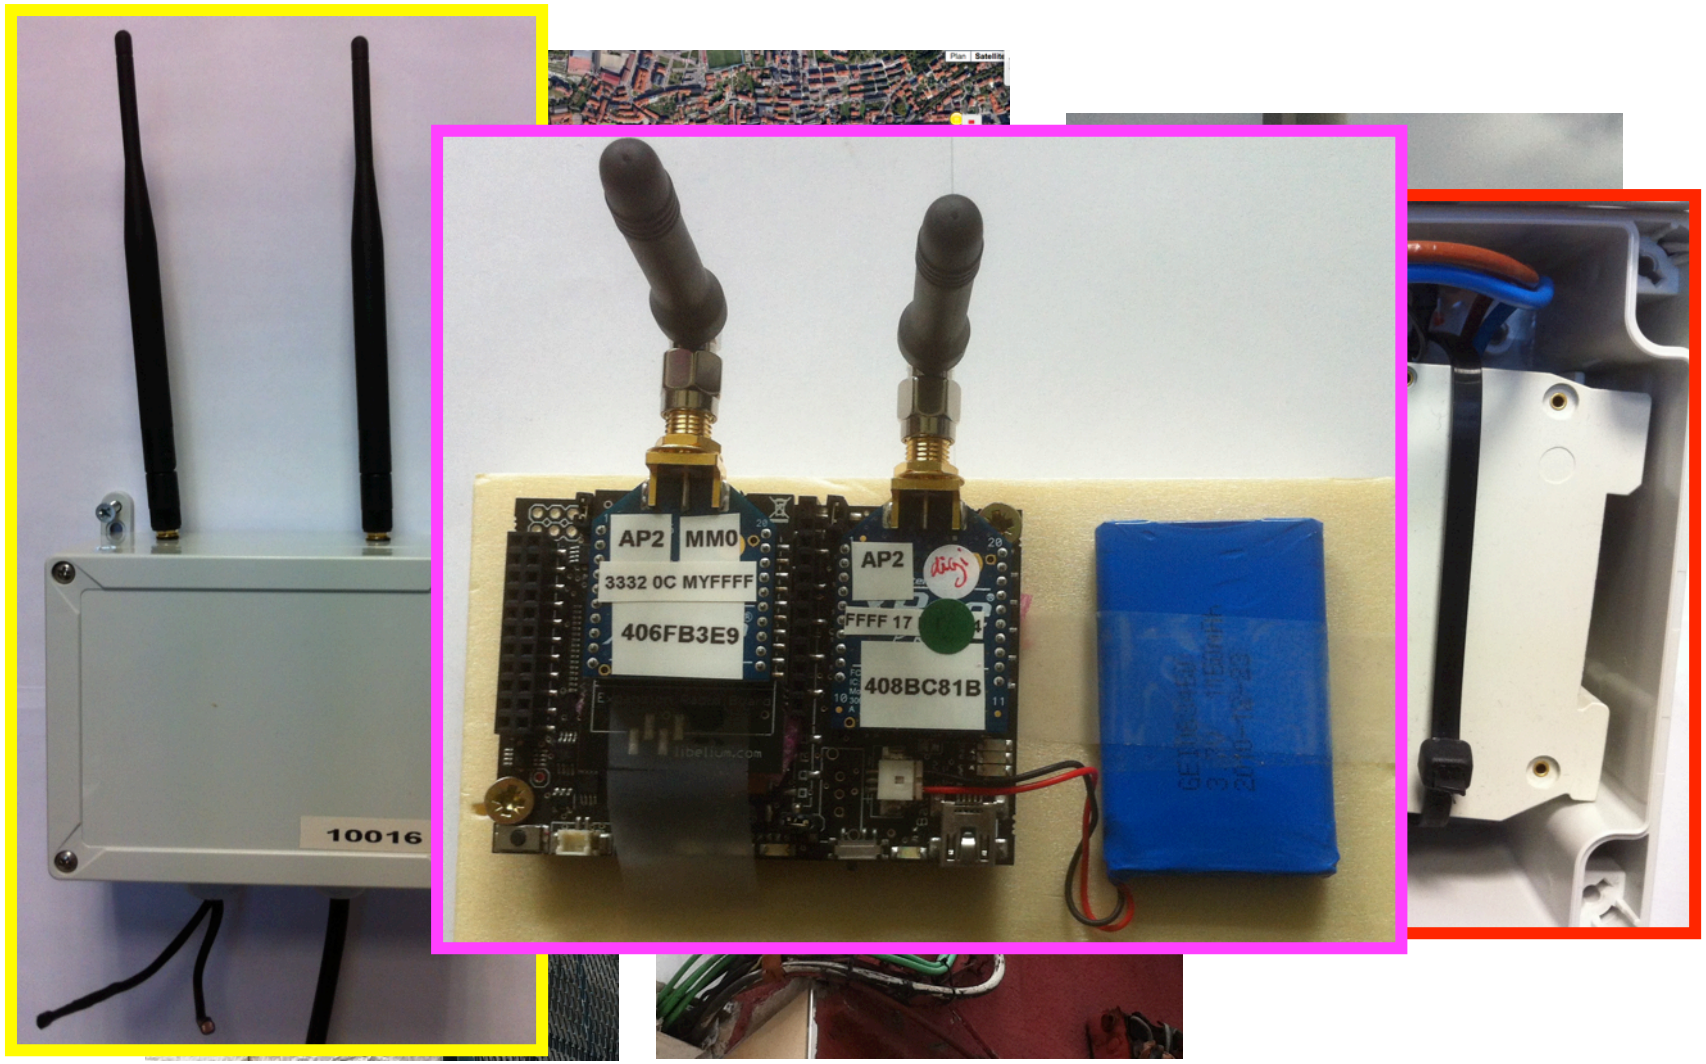

Test campaign – April 9th-10th 2013

Control & Instrument

The network behind the apps

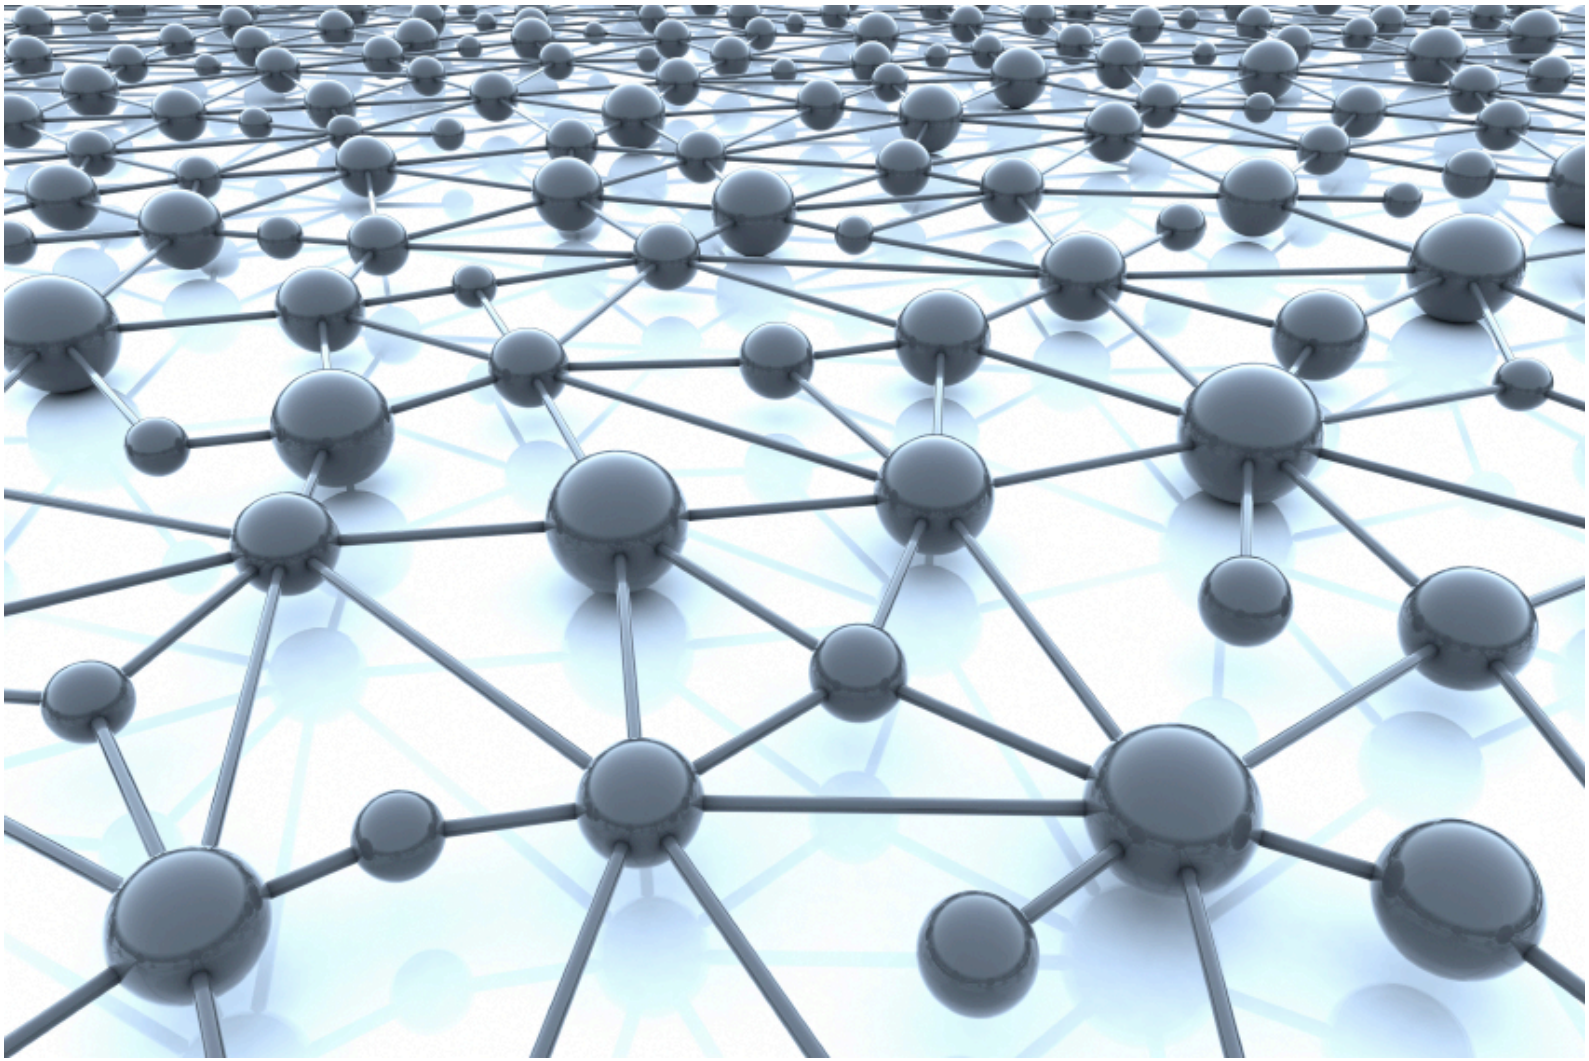

Operators and ISPs: they rule the Internet

- « 1st tier ISP » own their lines.
- Interconnections happen mostly at private peering points

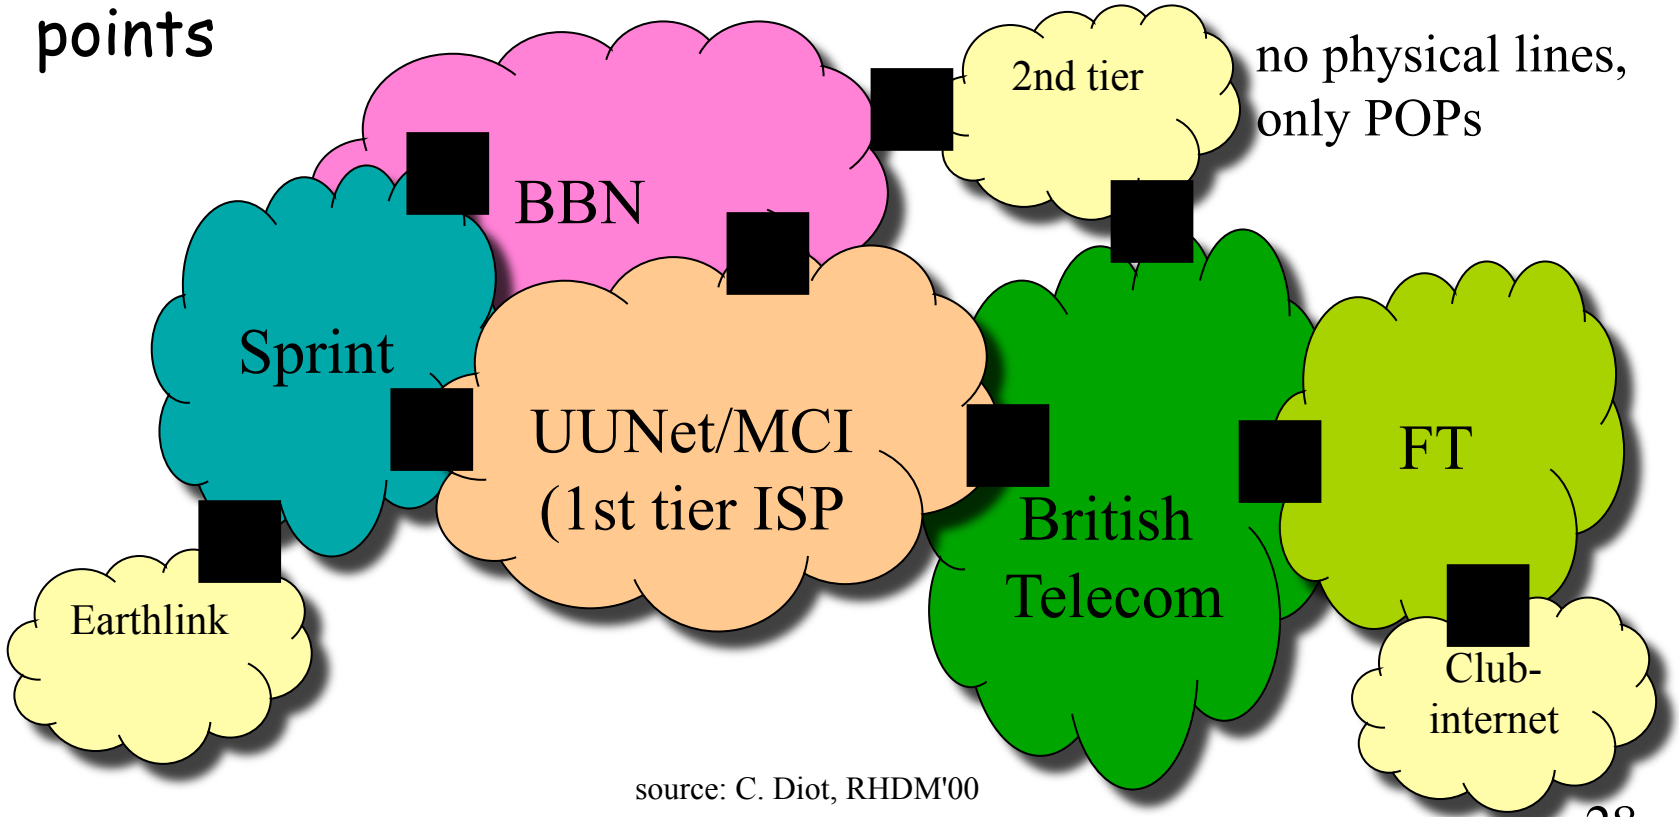

source: C. Diot, RHDM'00

Back in time:
The telephone system & network

First automatic Branch Exchange Almond B. Strowger, 1891…

Signaling replaces the operator

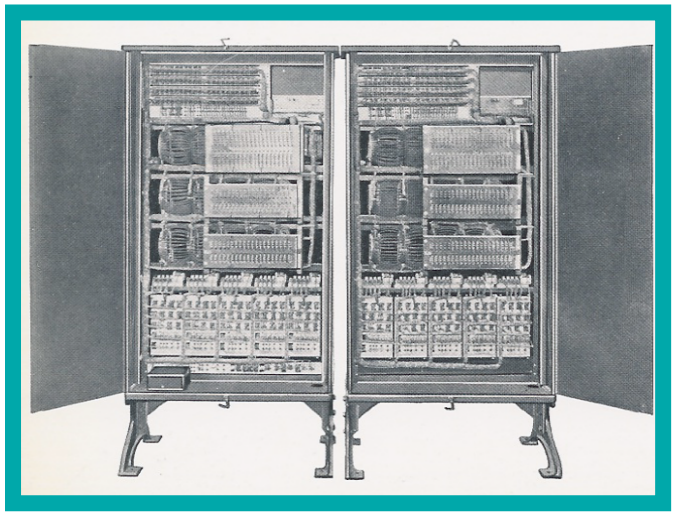

Source J. Tiberghien, VUB

The telephone circuit view

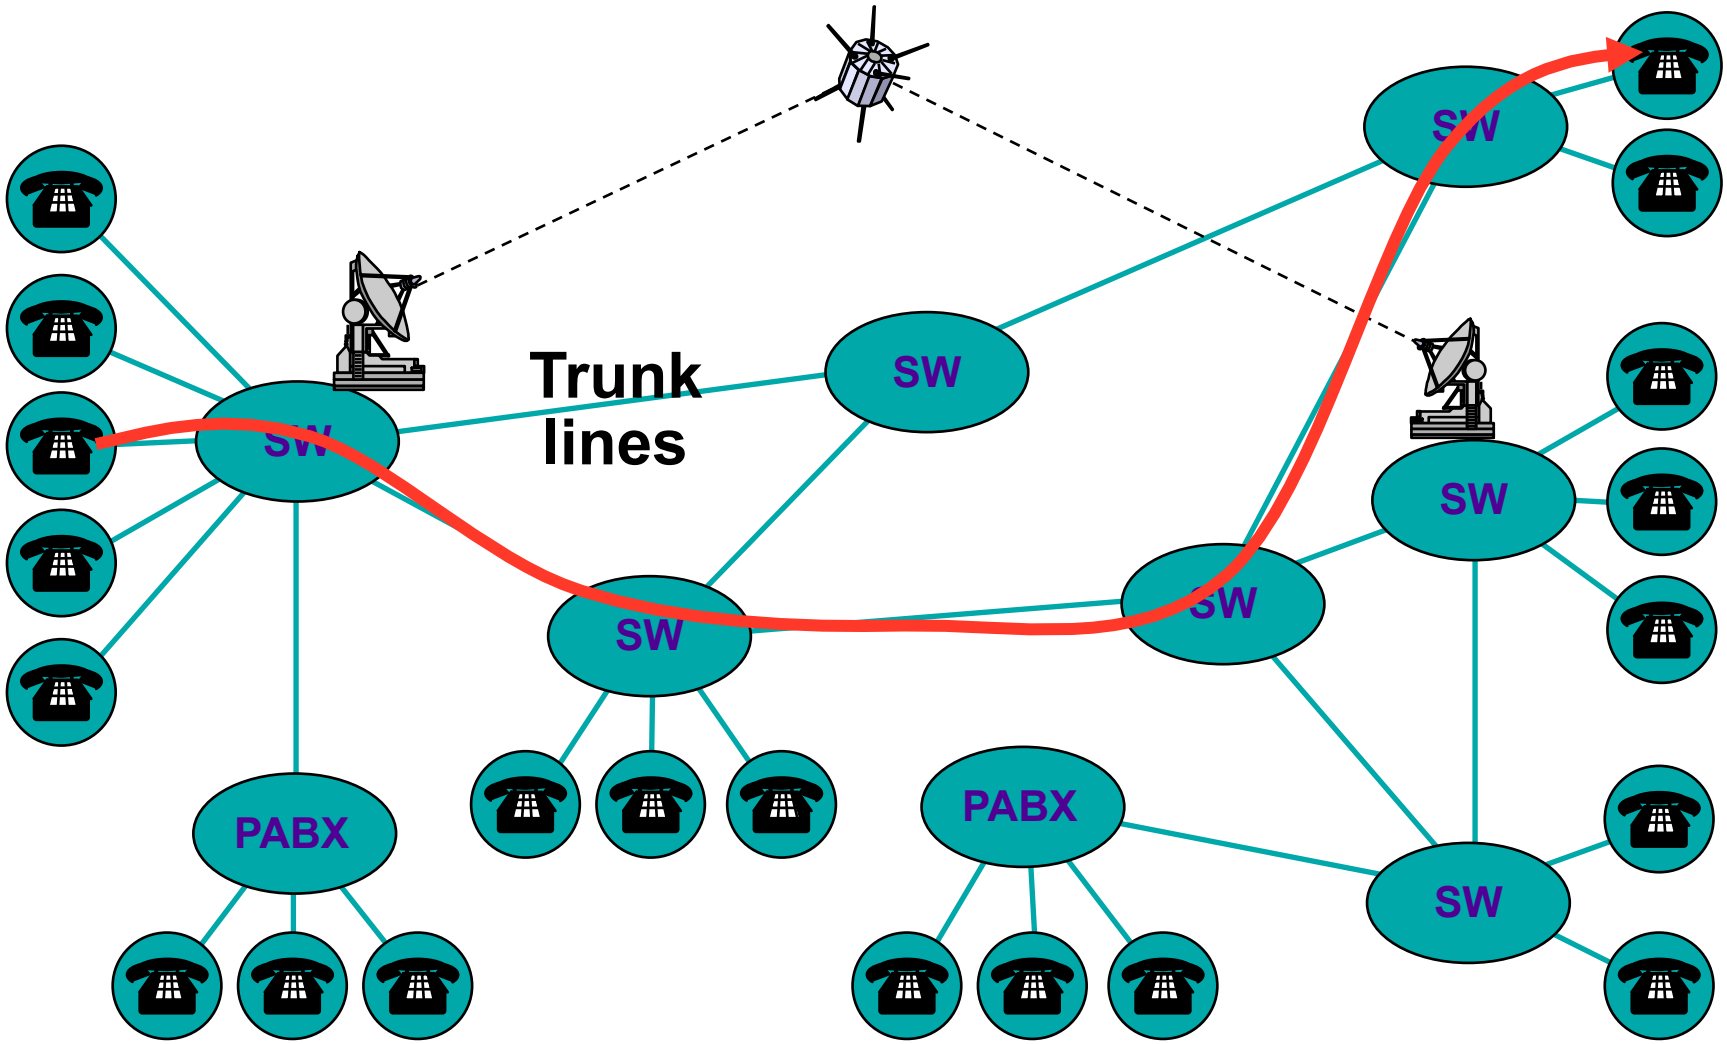

SONET/SDH transport network infrastructure

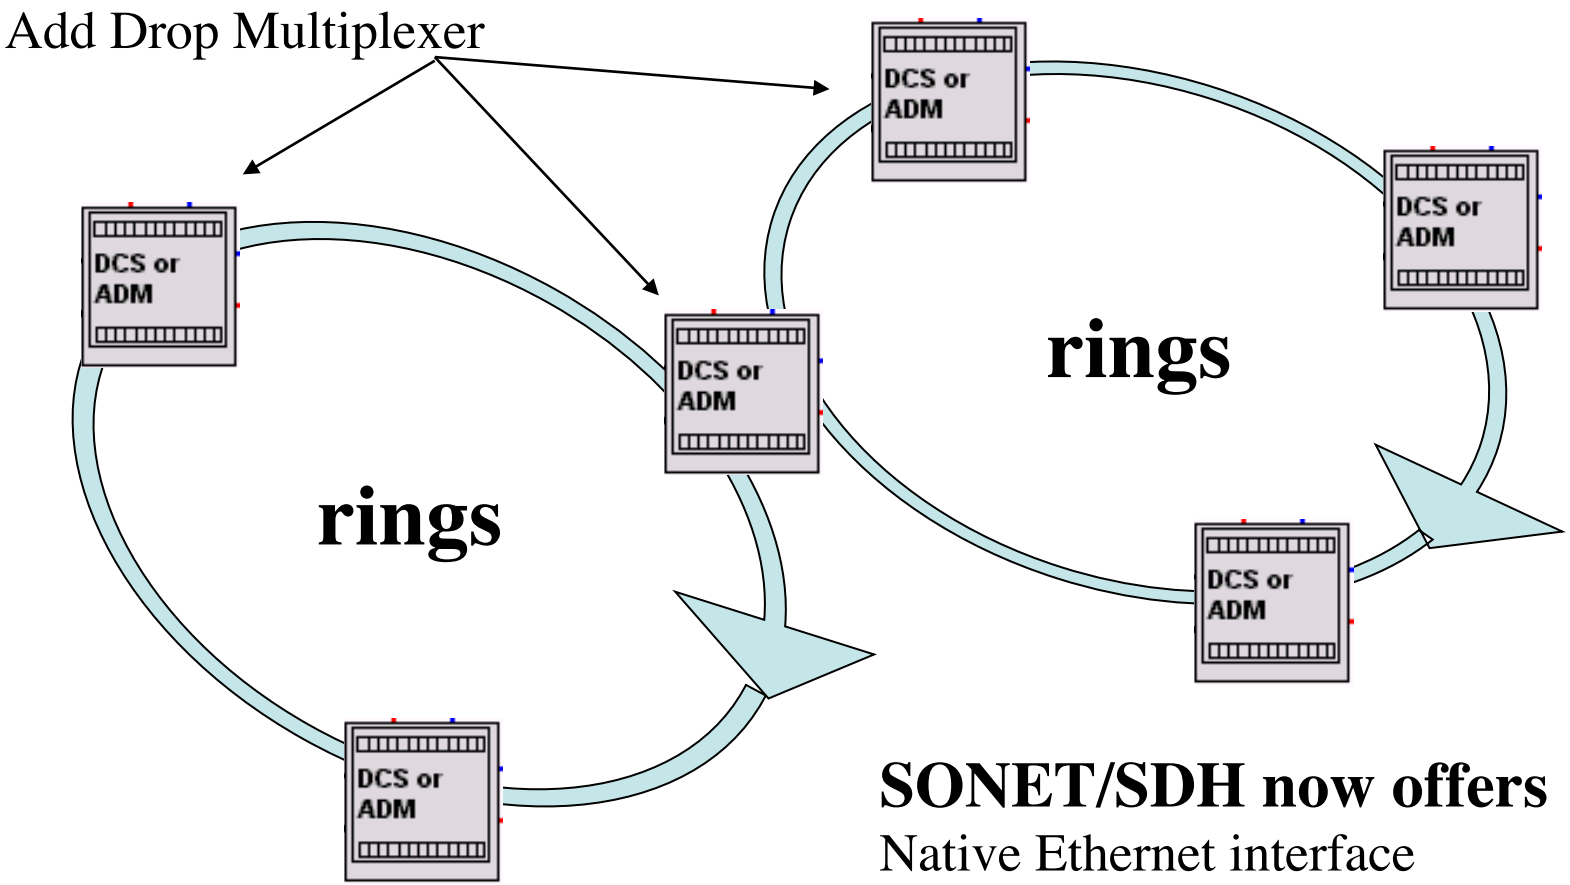

SONET/SDH now offers

Native Ethernet interface

Generic Framing Procedure

Virtual Concatenation

SONET/SDH and resiliency

- SONET/SDH has built-in fault-tolerant features with multiple rings
- Ex: simple case

DCS
(Digital Cross-Connects)

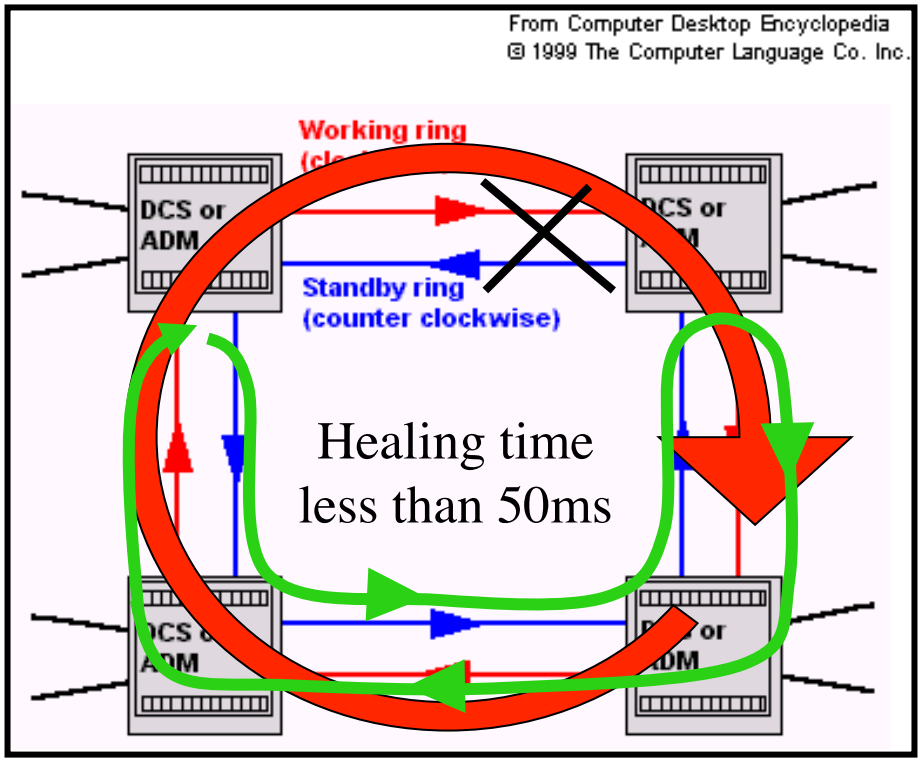

SONET/SDH and resiliancy

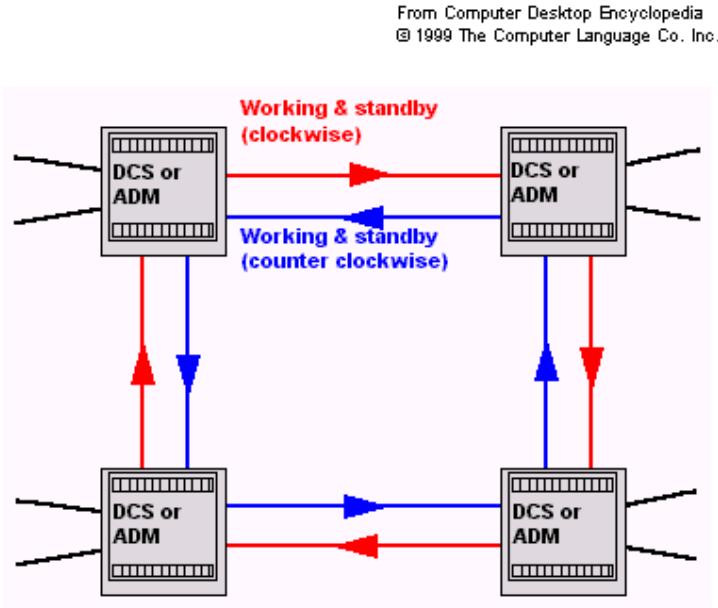

bi-directional

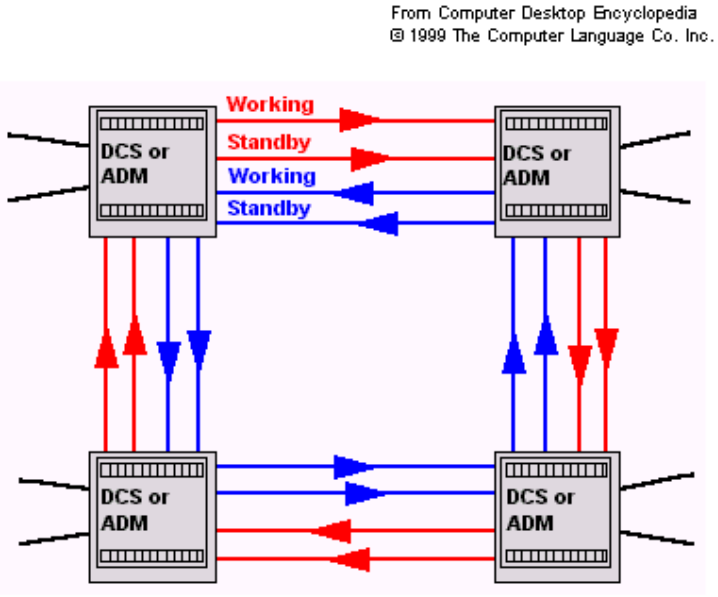

Found in most operators

General Purpose SDH Networks

NG SONET/SDH now offers

Native Ethernet interface

Generic Framing Procedure

Virtual Concatenation

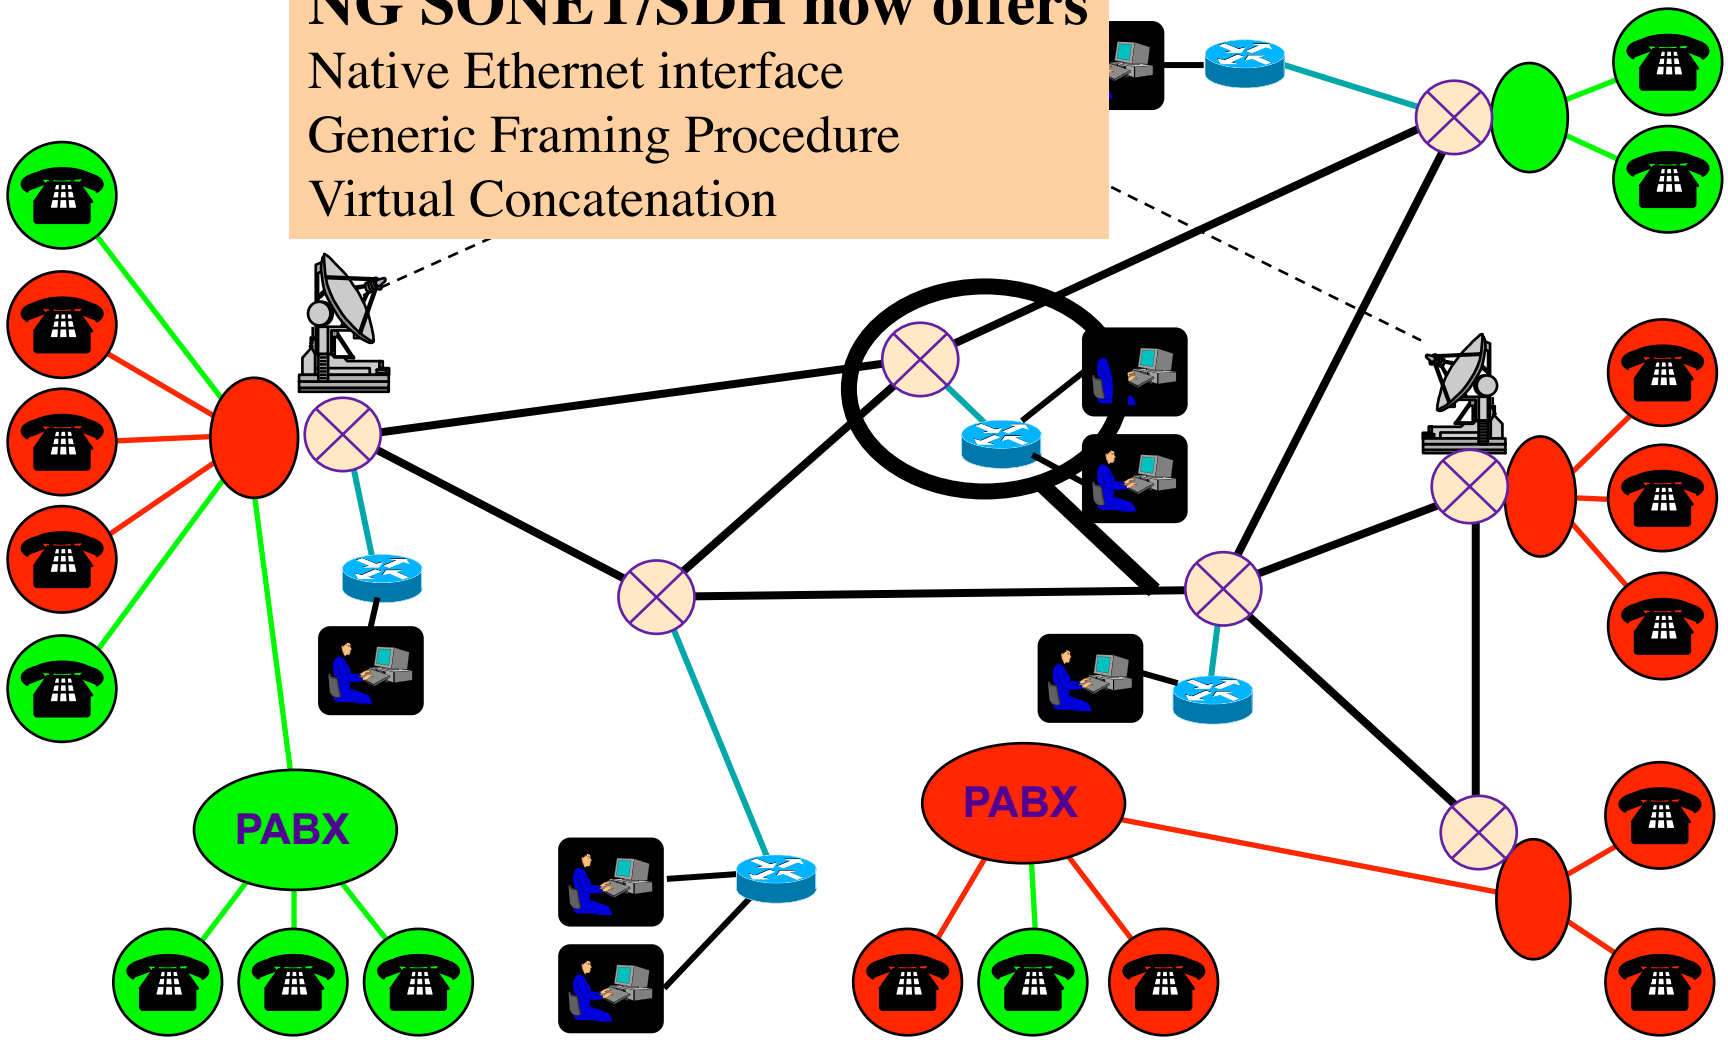

What is Quality of Service?

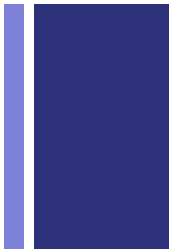

- Quality of service is the ability to provide different priority to different applications, users, or data flows, or to guarantee a certain level of performance
- QoS criteria are numerous and is highly dependant of the app.
  - Throughput, Delay, jitter, Loss rate, Availability, Uptime, ...
- ... Or driven by the end-user
  - Image resolution, sound quality, appropriate language, ...

What is QoS (contd) ?

- These parameters can be measured at several granularities:
  - "micro" flow, aggregate flow, population.
- QoS considered "better" if
  - more parameters can be specified
  - QoS can be specified at a fine-granularity.
- QoS spectrum:

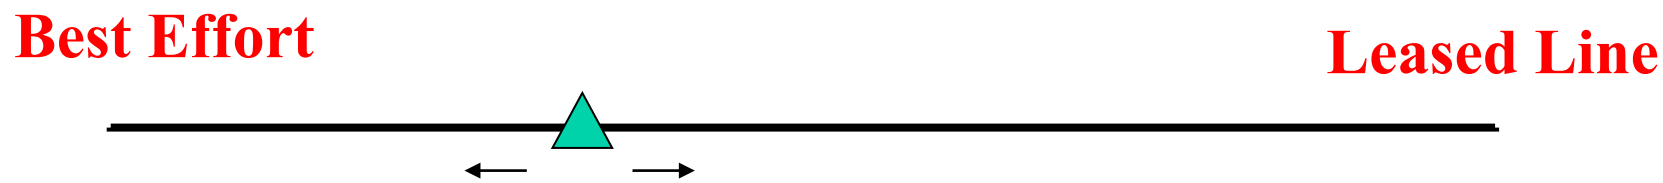

Where to put QoS?

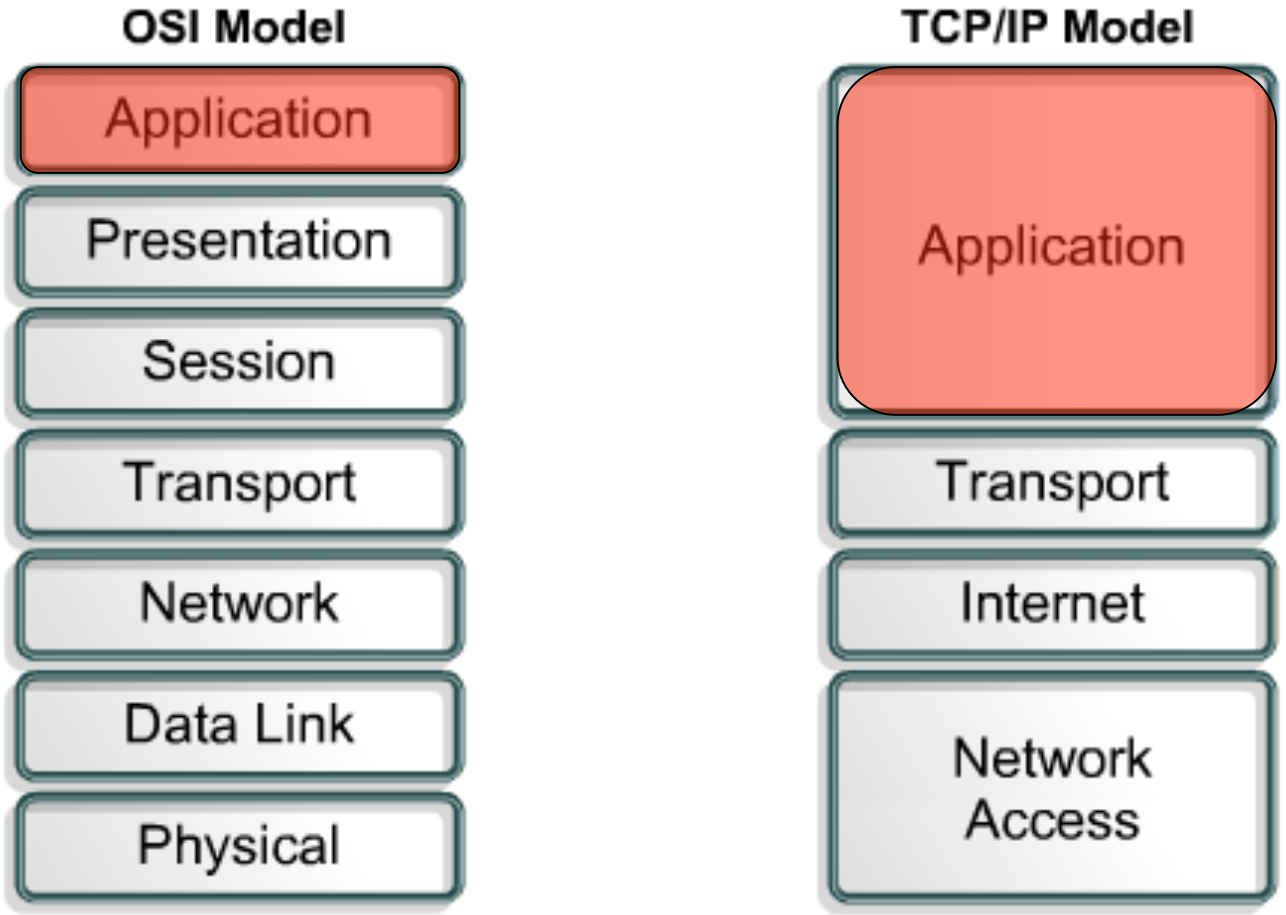

Application layer=network as a black box

Dealing with packet jitter
Fixed playout delay

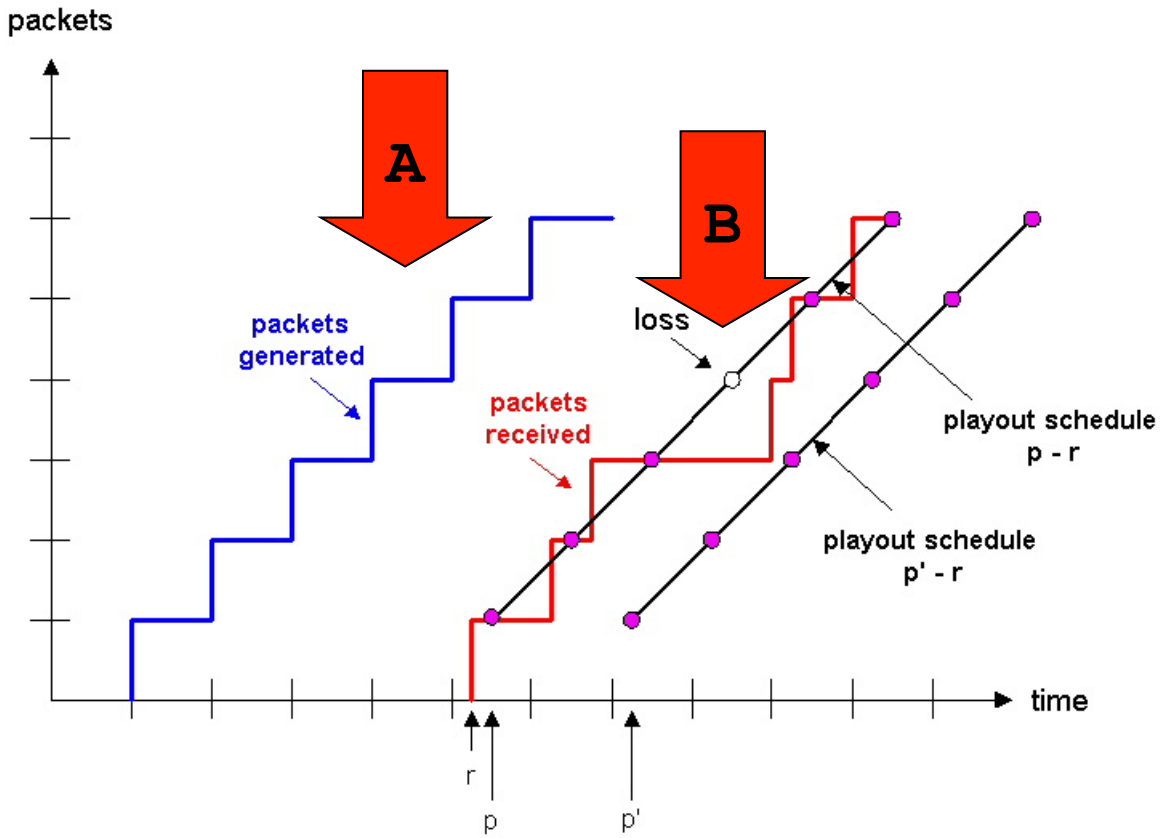

From Xavier Appé, modified by C. Pham for educational purpose only

Recovering from packet loss Interleaving

- Divide 20 msec of audio data into smaller units of 5 msec each and interleave
- Upon loss, have a set of partially filled chunks

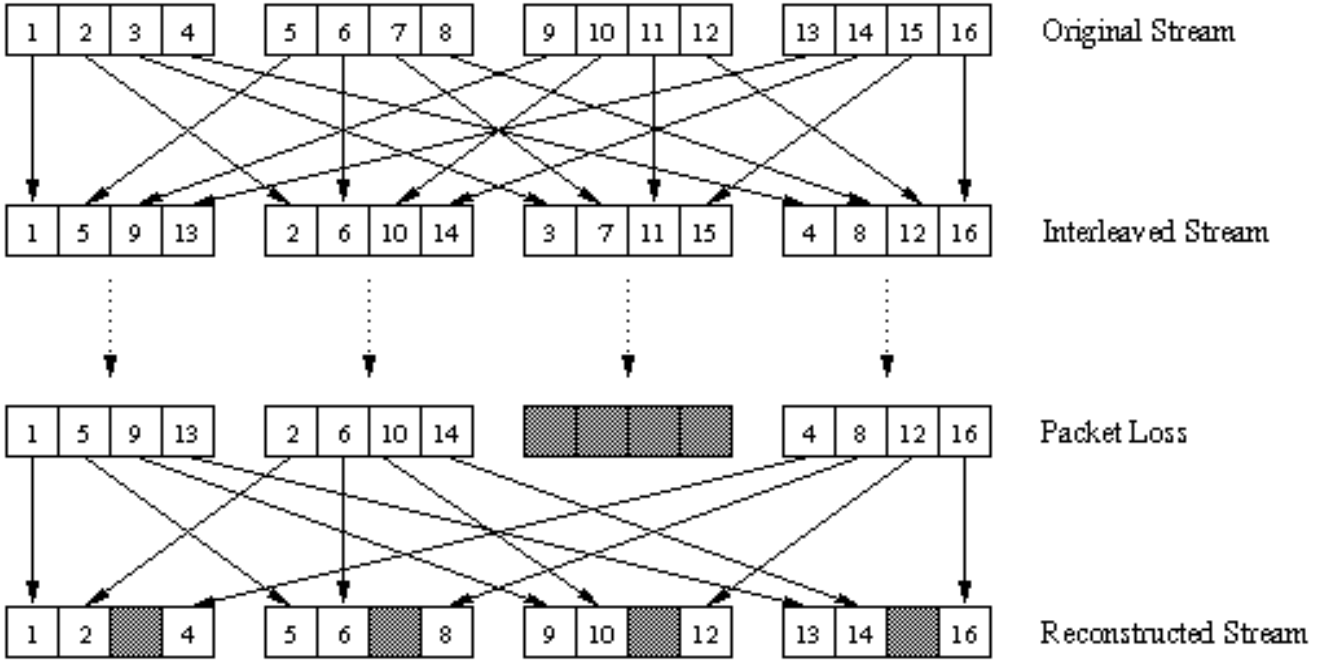

From Xavier Appé, modified by C. Pham for educational purpose only

Large block FEC codes...

- an example: LDPC code
  - based on XOR operations (⊕)
  - uses bipartite graphs between source and FEC symbols
  - iterative decoding

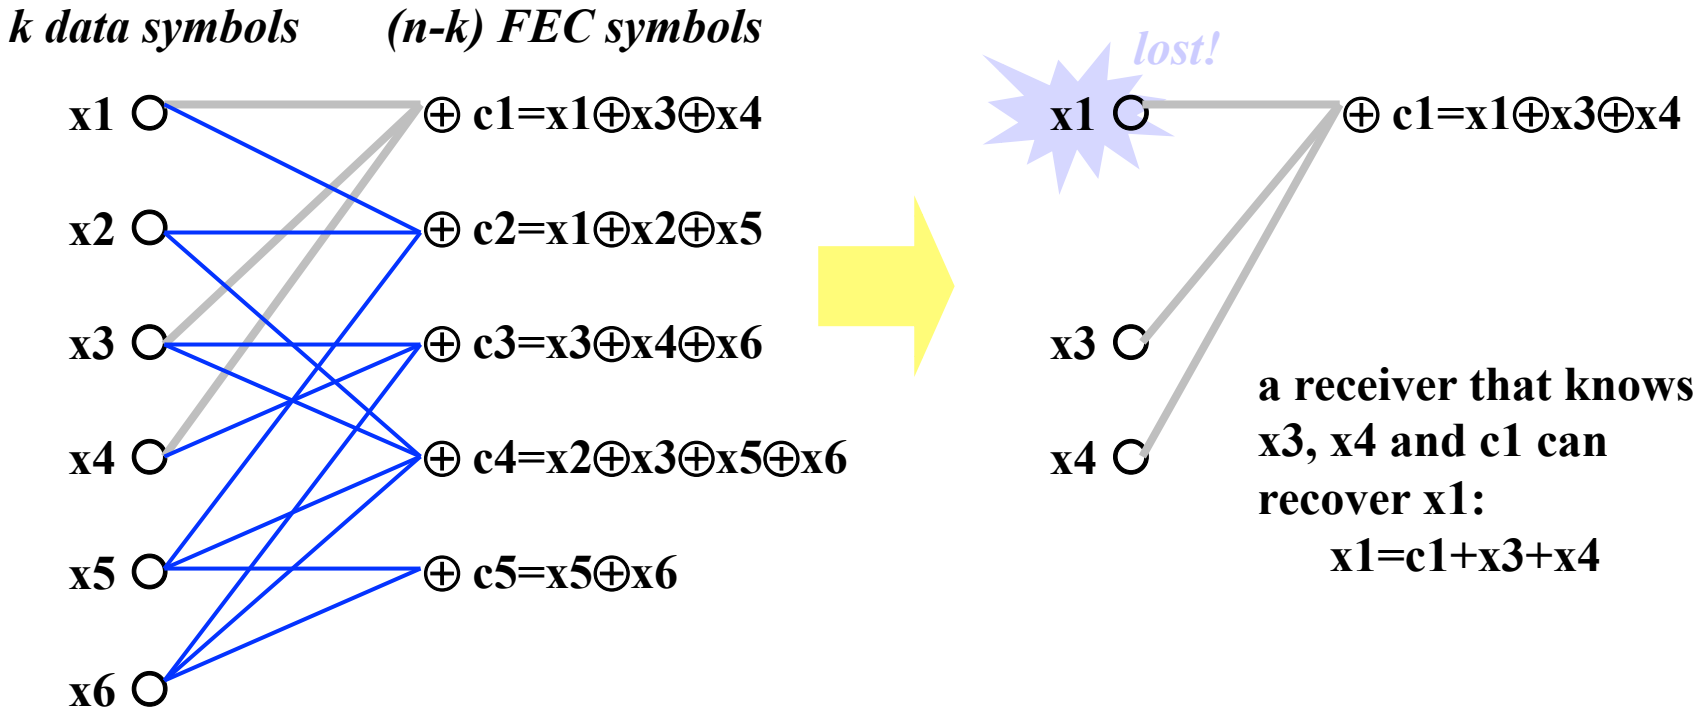

From Xavier Appé, modified by C. Pham for educational purpose only

MUST maintain IP philosophy

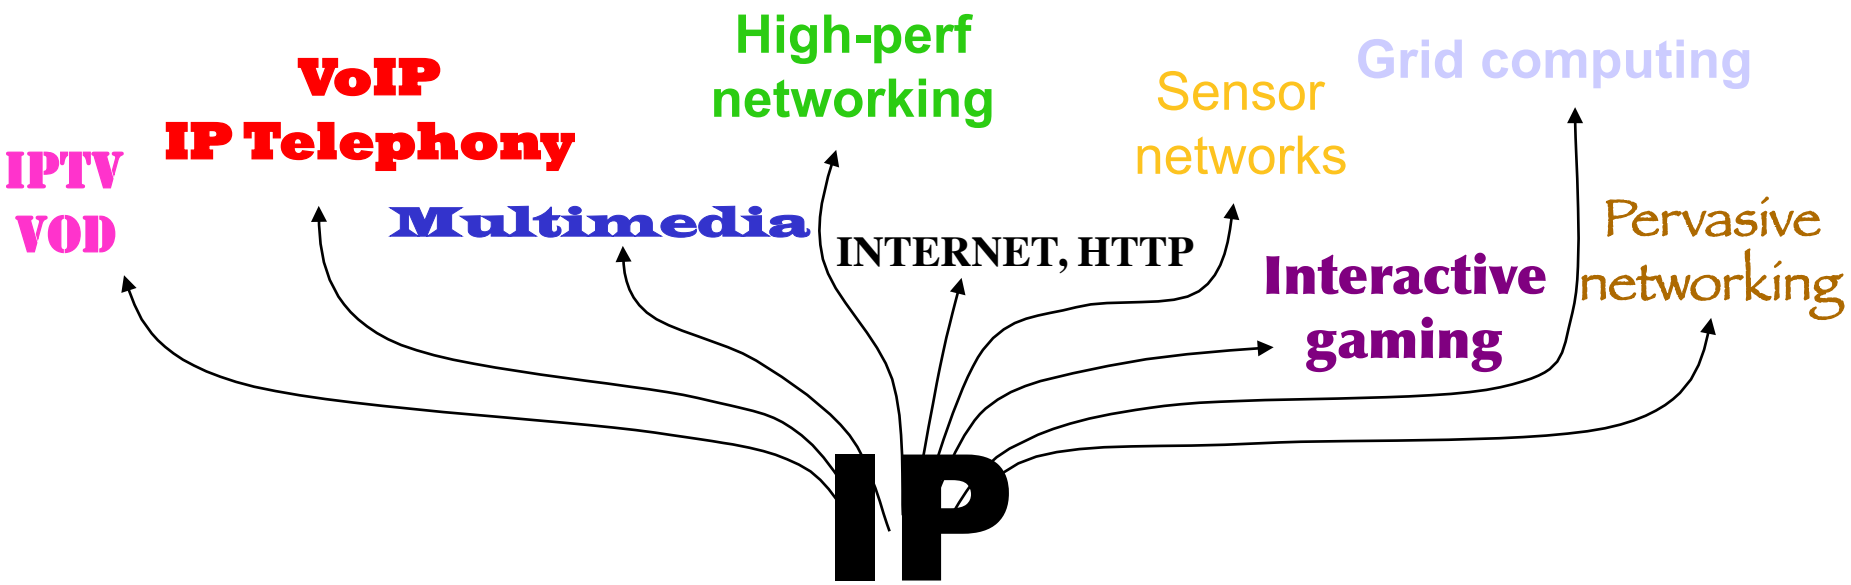

Overprovisioning in the core

- Most operators are overprovisioning bandwidth with DWDM fibers
- 10Gbps, 40Gbps, 160 GBps, 320 Gbps, much mote to come!

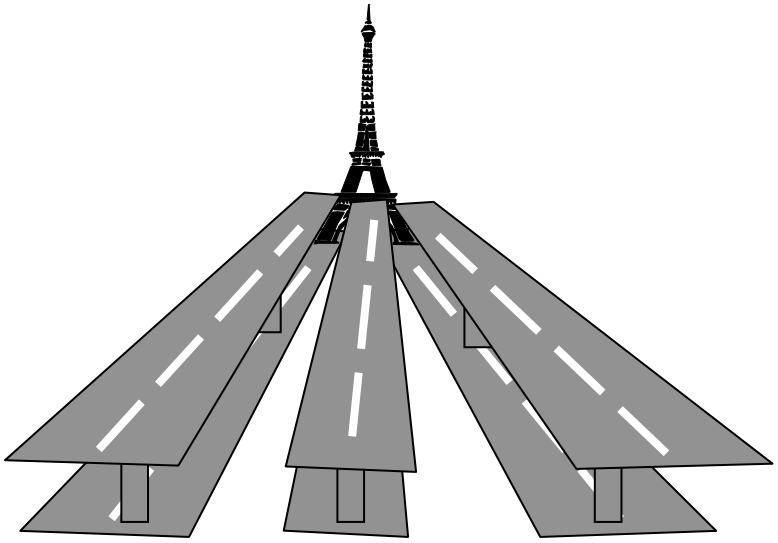

IP desired service

- Isolation: my traffic is not impacted at all by yours
- Protection: my transmission path is backed up to the nth degree by failover paths
- Throughput: I get the capacity I pay for
- Delay: Whatever pattern of packets timing I send with is preserved at the far-end

30 years of INTERNET QoS…

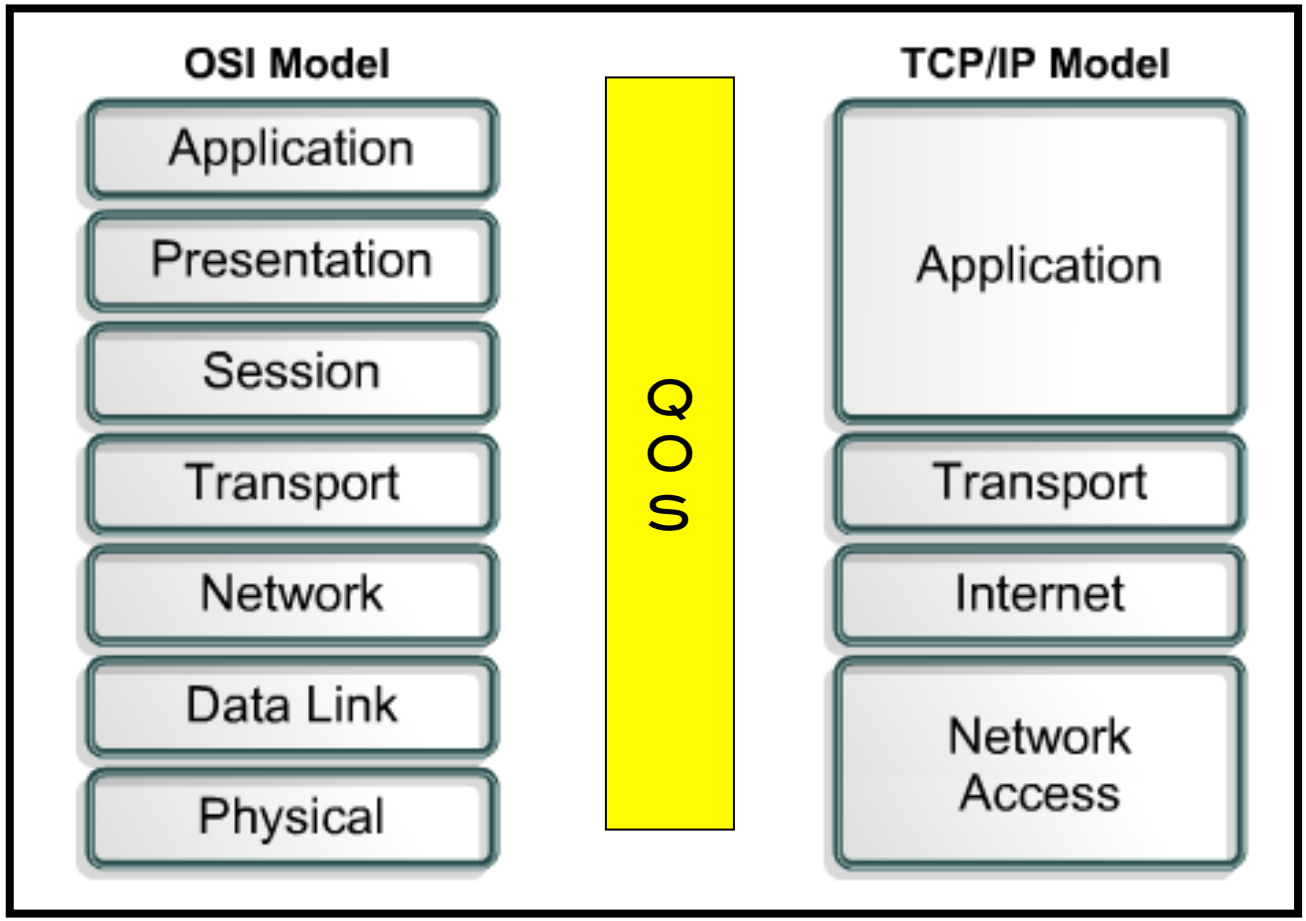

…have shown the power of selfishness!

WHY SHOULD I BOTHER WITH QoS WHEN…

…OTHERS DON'T DO IT?

It's not my fault!

« environmental problems often have impacts beyond borders »

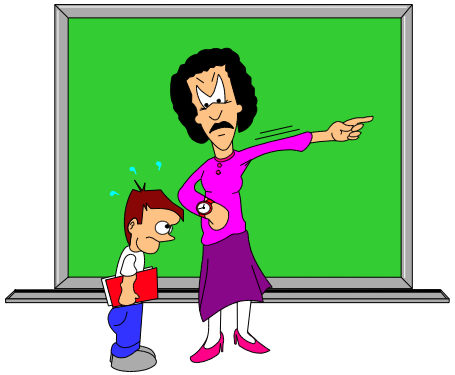

- What's the point of deploying QoS if others don't?

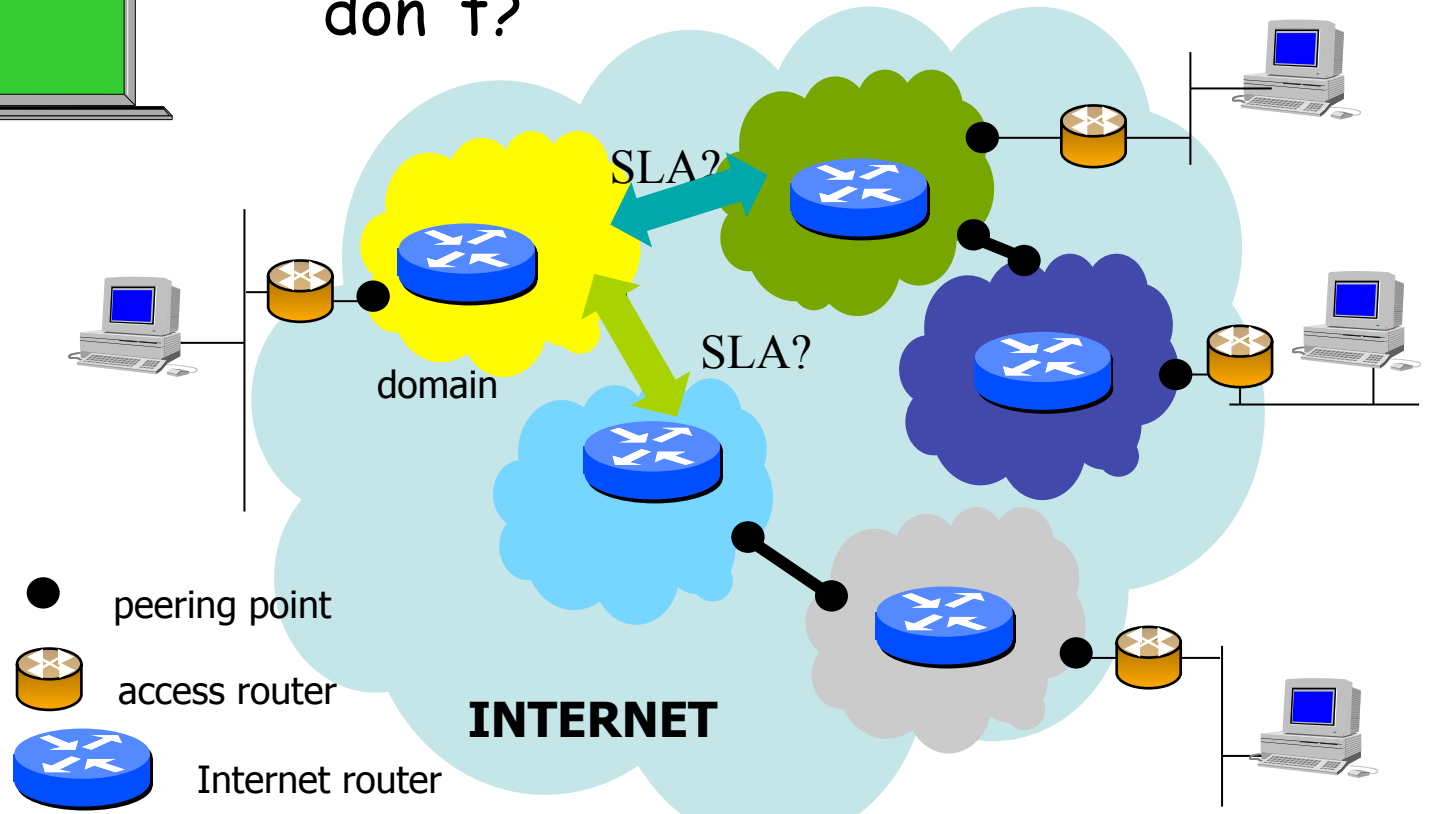

Current Internet's QoS

So why change?

Sustainable development

- "meets the needs of the present without compromising the ability of future generations to meet their own needs" [Brundtland Report, 1987]
- Trade-off between performance and needs: « why are we producing? »
- Use the right ressource, at the right place, at the right time

a new dimension of global responsibility—
not only to planetary resources but also to planetary fairness

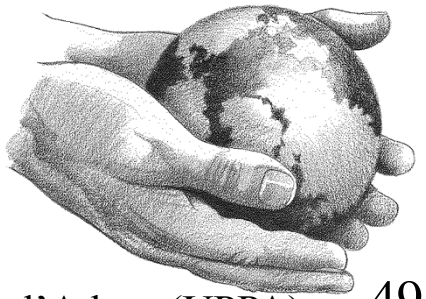

Is overprovisioning harmful?

- NO: o ly.
  Addin
  Custor n on
  invest
- YES: w ive
  solutic that
  relying makes
  upgrac)



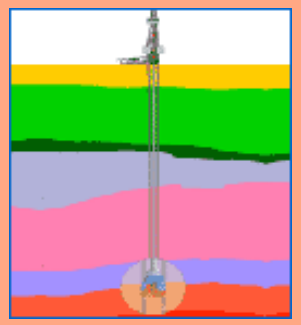

Each new oilfield discovery delays research and development of alternative energies

Lessons learned from sustainable development

- Limit globalization
- Limit the pursuit of continued economic prosperity
- Redistribute labour, wages,…
- Promote the use of local resources
- Change mentality

Community networks?

CONNECTED COMMUNITIES

CONFINE

Limitations of the current Internet

- Bandwidth
  - Raw bandwidth is not a problem: DWDM
  - Provisioning bandwidth on demand is more problematic
- Latency
  - Mean latencies on Internet is about 80-160ms
  - Bounding latencies or ensuring lower latencies is a problem
- Loss rate
  - Loss rate in backbone is very low
  - End-to-End loss rates, at the edge of access networks are much higher
- Communication models
  - Only unicast communications are well-defined: UDP, TCP
  - Multi-parties communication models are slow to be deployed

Net Neutrality or Not?

- NN or NNN? That's the question!
- NN = dumb network!
- Internet's success is in a large part debtful to what's called Net Neutrality (IP neutrality)

Some NN success stories

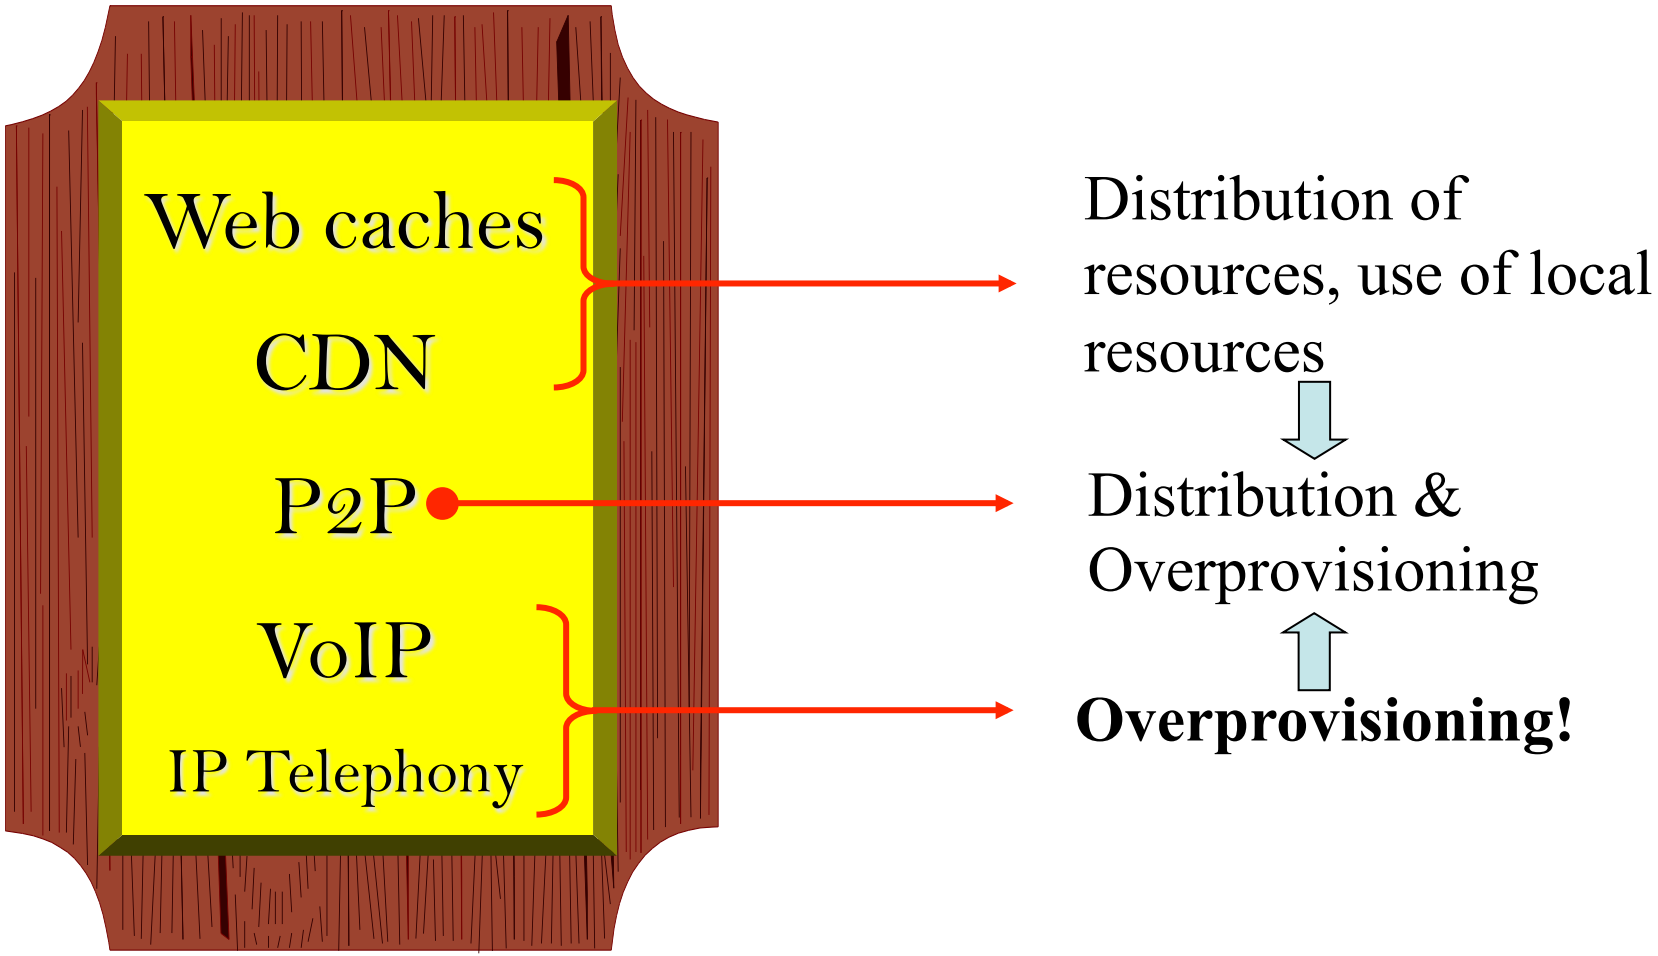

Where to put QoS?

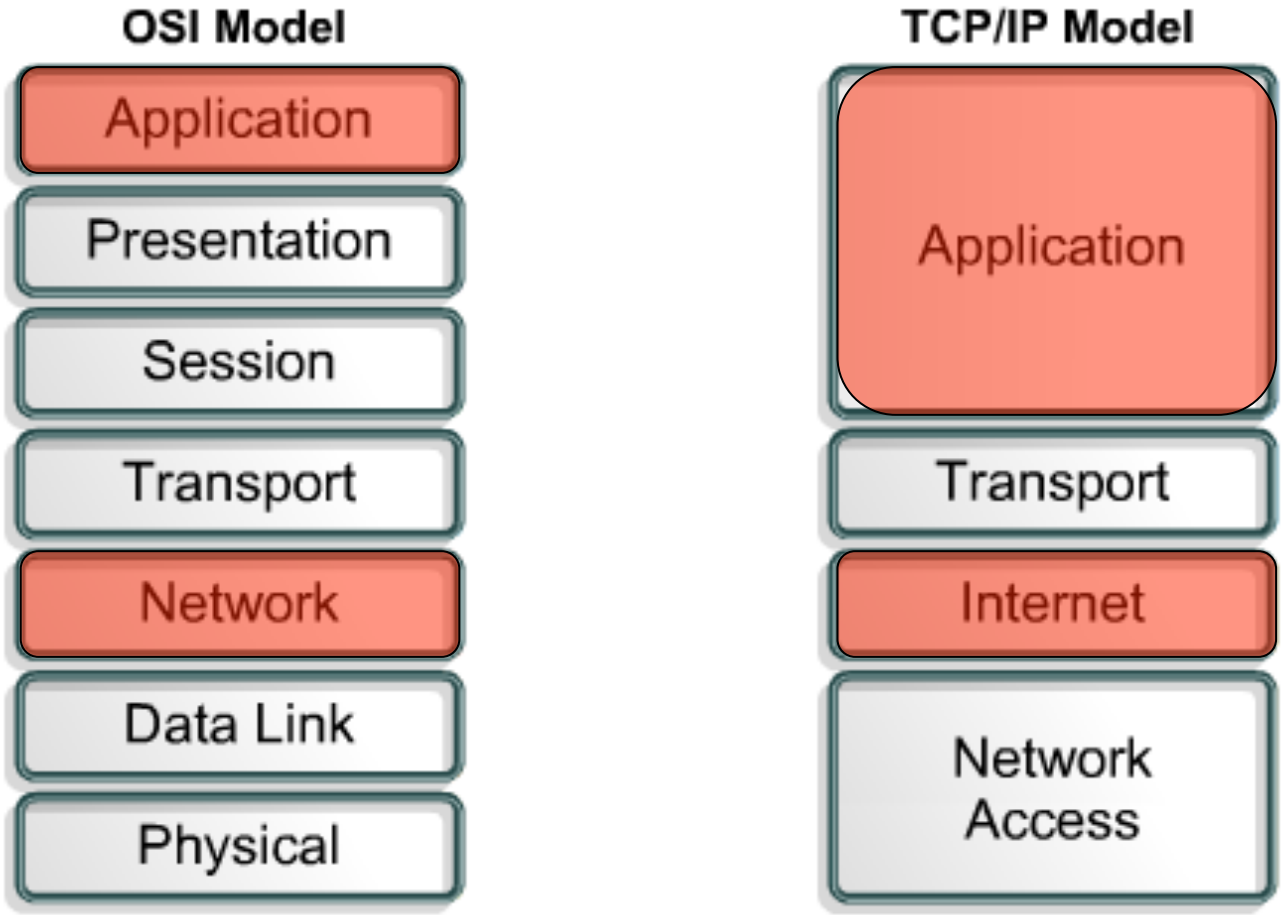

Internet Routers

©Juniper

©Procket Networks

©Alcatel

©Nortel Networks

©Lucent

and more…

If no NN then give more power to routers!

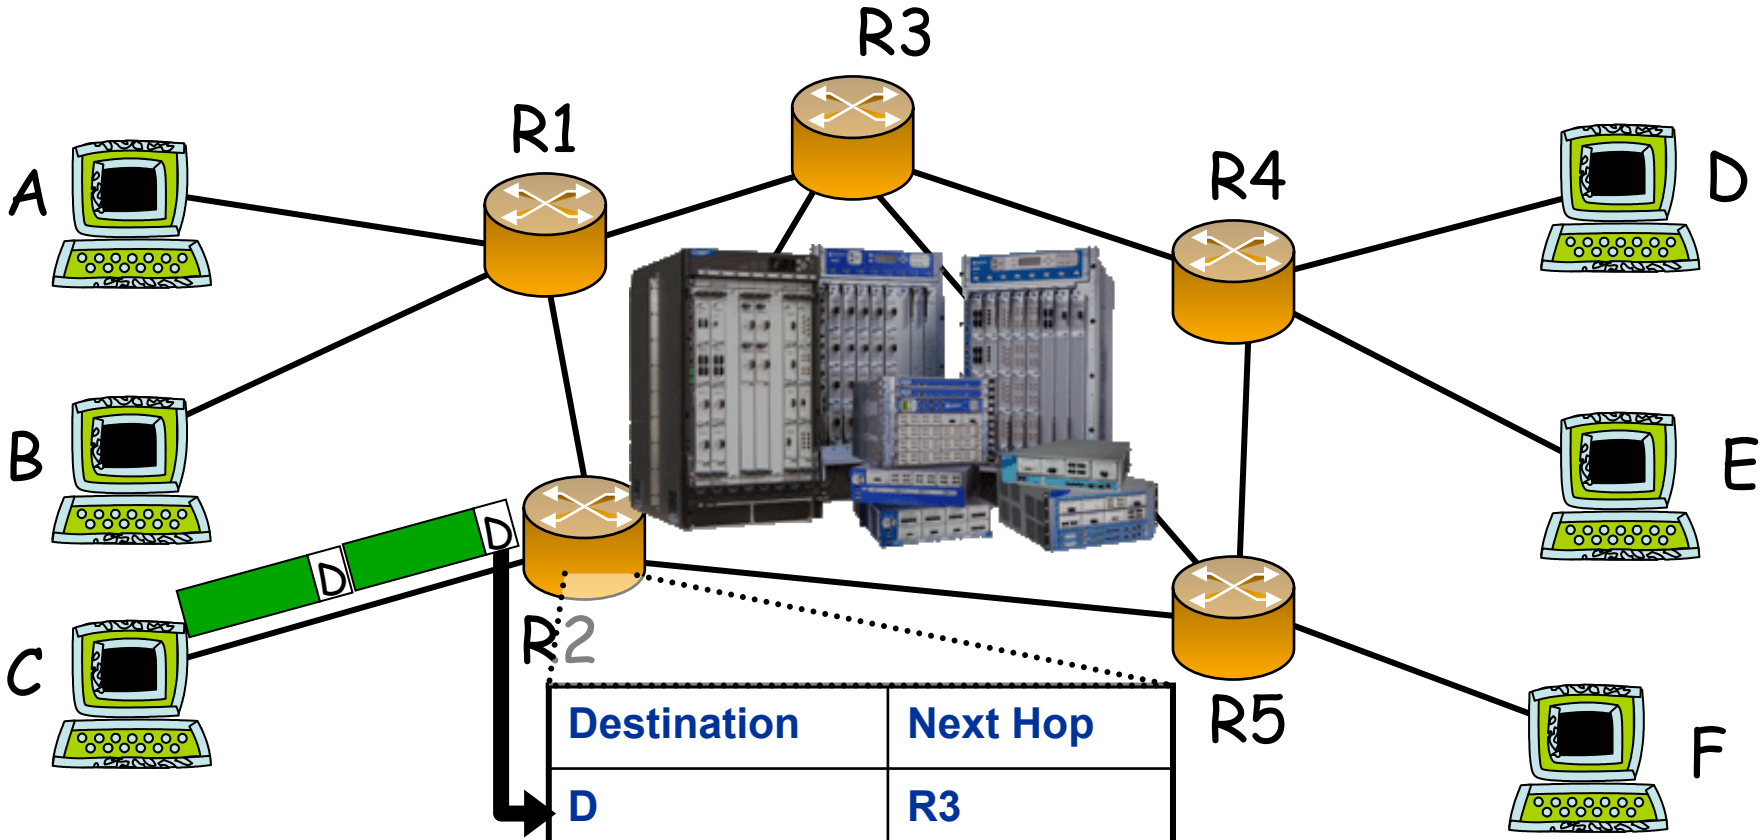

Destination	Next Hop
D	R3
E	R3
F	R5

General architecture of an IP router

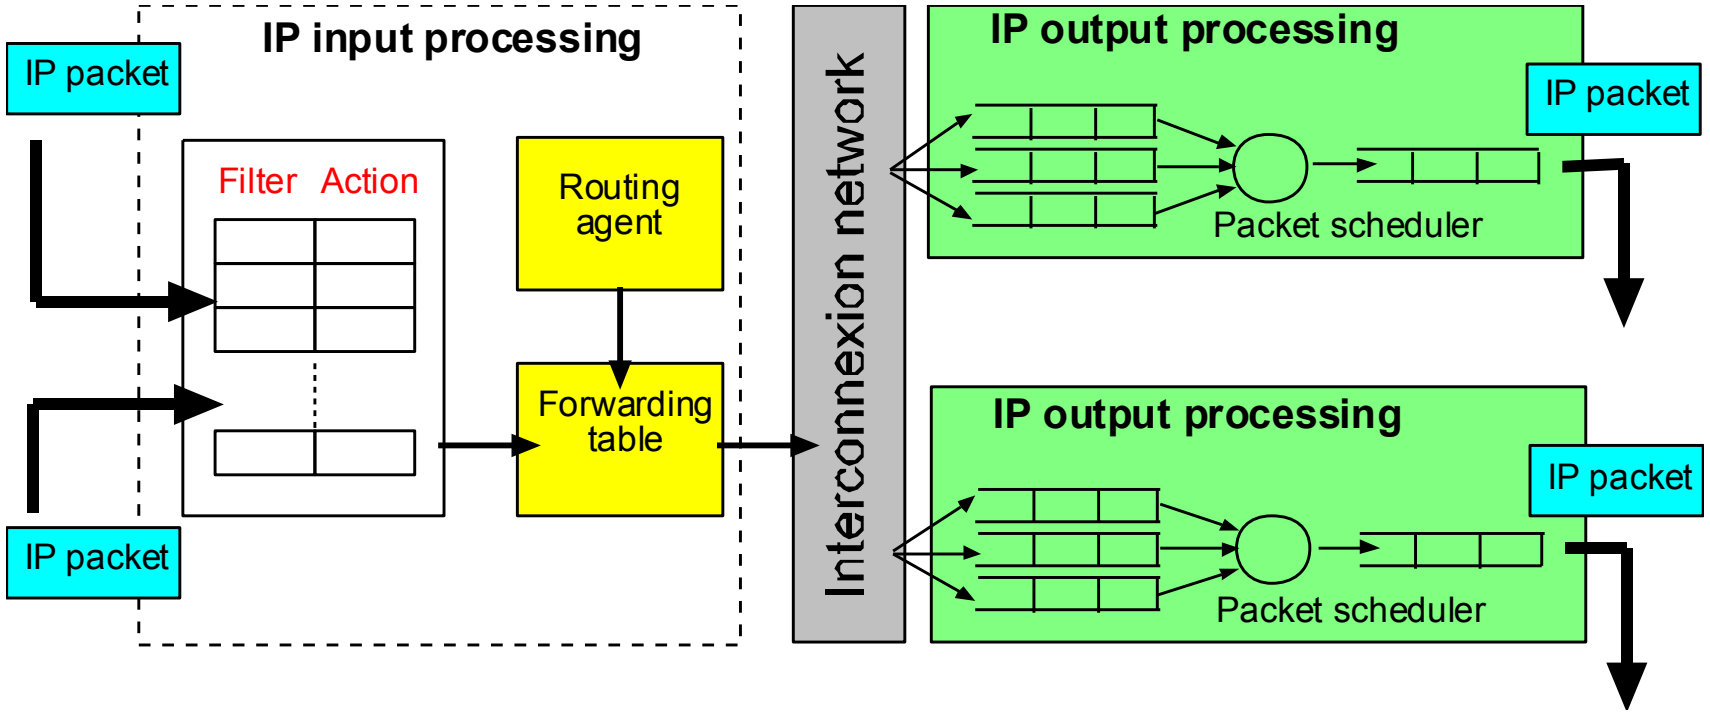

- receives input packets,
- sends packets to output buffers,
- transmits packets.

In 2000, I had a dream: active networking!

- Programmable nodes/routers
- Customized computations on packets
- Standardized execution environment and programming interface
- No killer applications, only a different way to offer high-value services, in an elegant manner
- However, adds extra processing cost

Motivations behind Active Networking

- From the user perspective
  - applications can specify, implement, and deploy (on-the-fly) customized services and protocols
- From the operator perspective
  - reduce the latency/cost for new services deployment/management
- From the network perspective
  - globally better performances by reducing the amount of traffic

Active networks implementations

- Discrete approach (operator's approach)
  - Adds dynamic deployment features in nodes/ routers
  - New services can be downloaded into router's kernel
- Integrated approach
  - Adds executable code to data packets
  - Capsule = data + code
  - Granularity set to the packets

The discrete approach

- Separates the injection of programs from the processing of packets

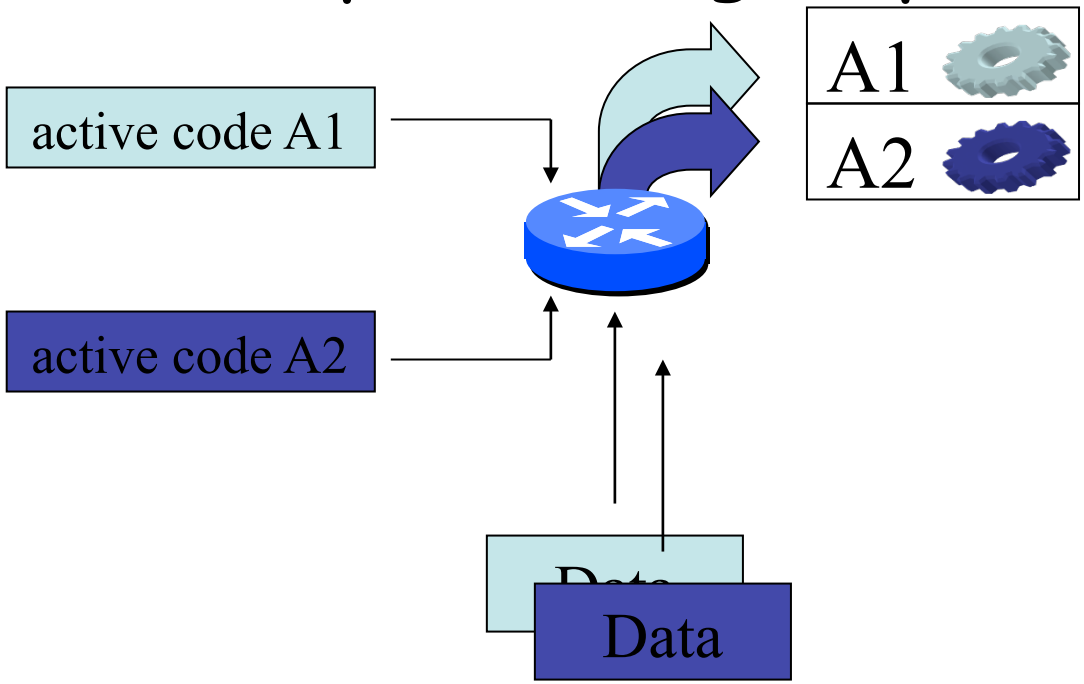

The integrated approach

- User packets carry code to be applied on the data part of the packet

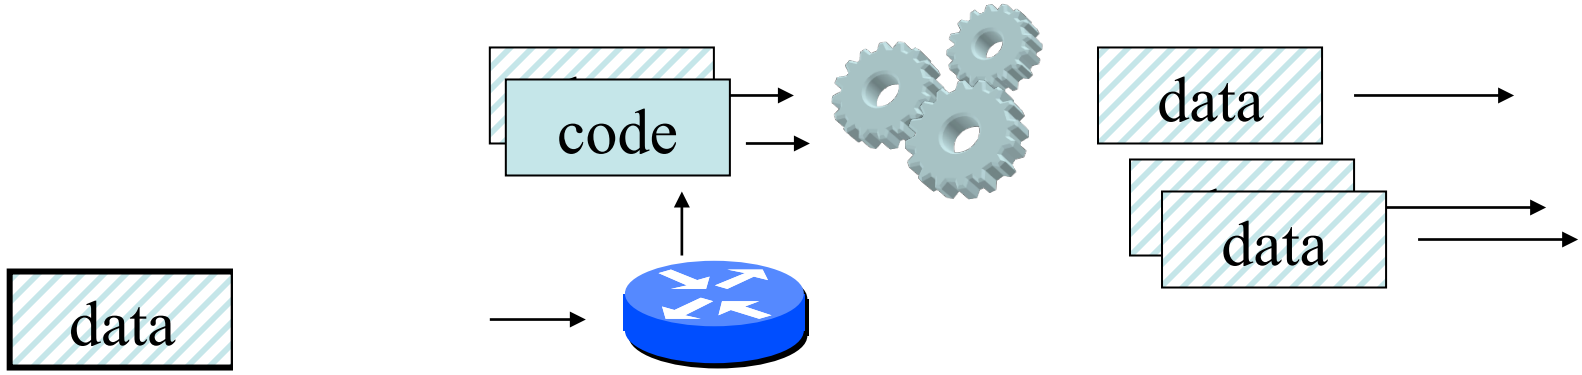

- High flexibility to define new services

An active router

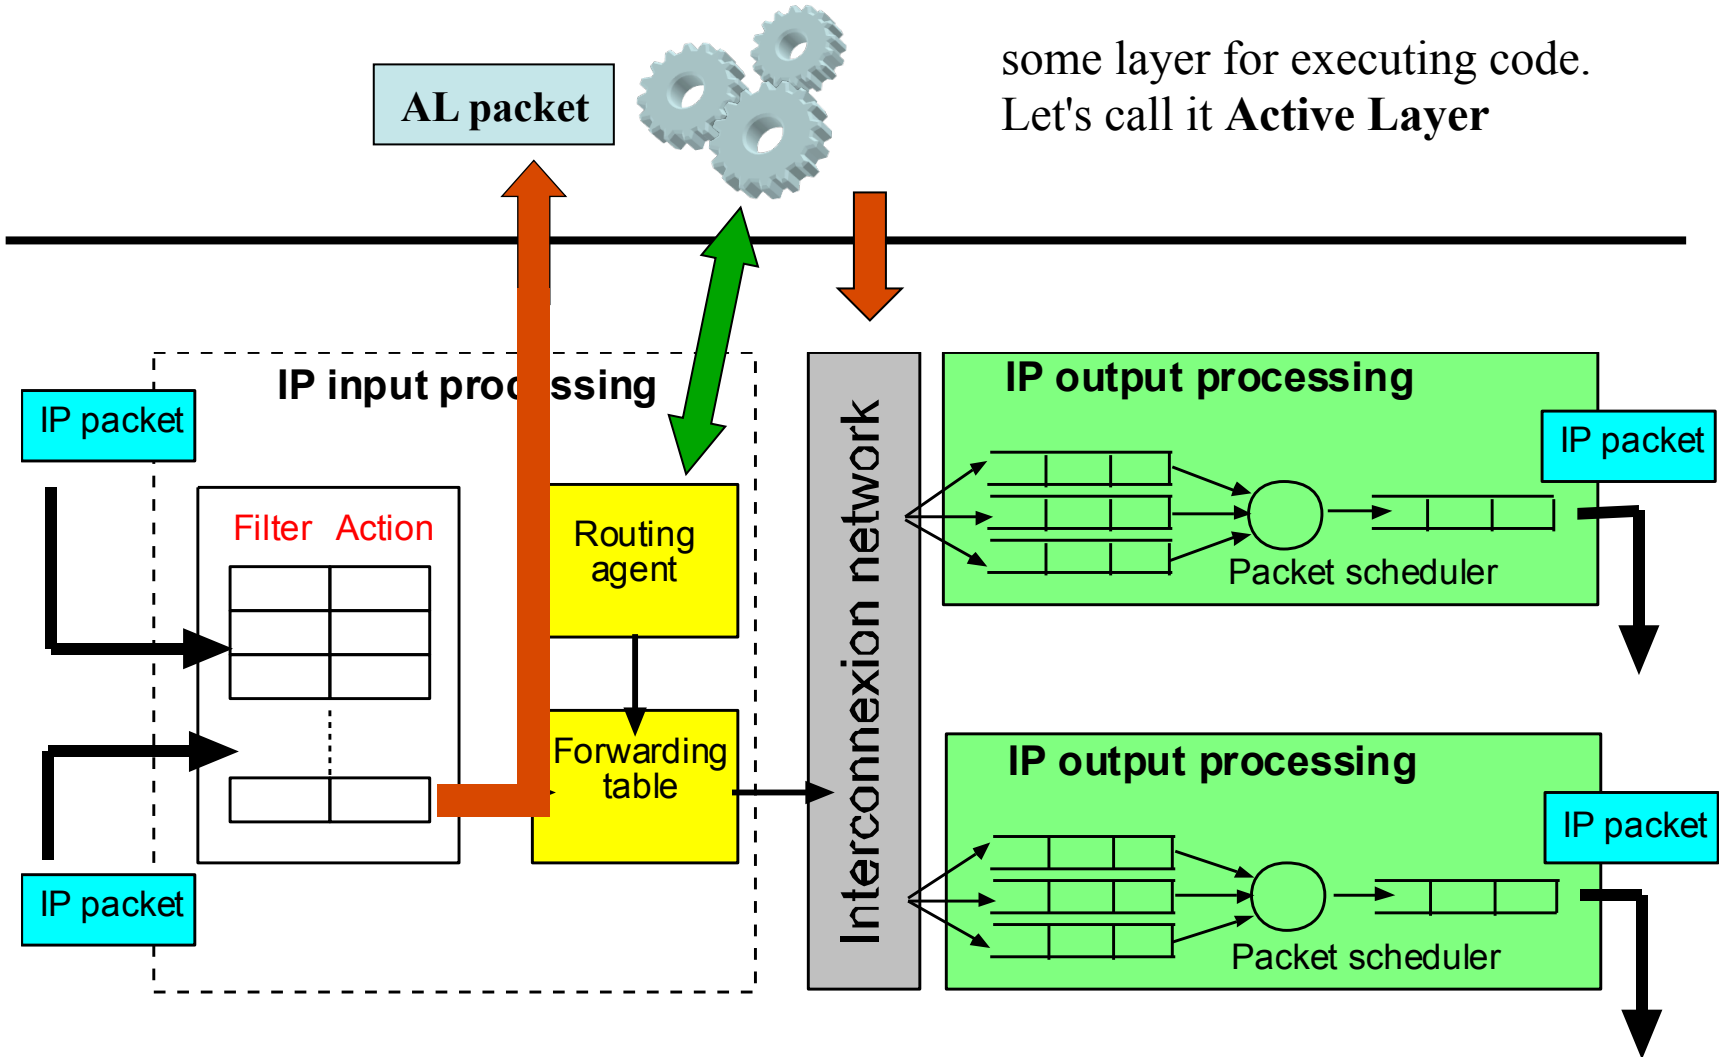

Interoperability with legacy routers

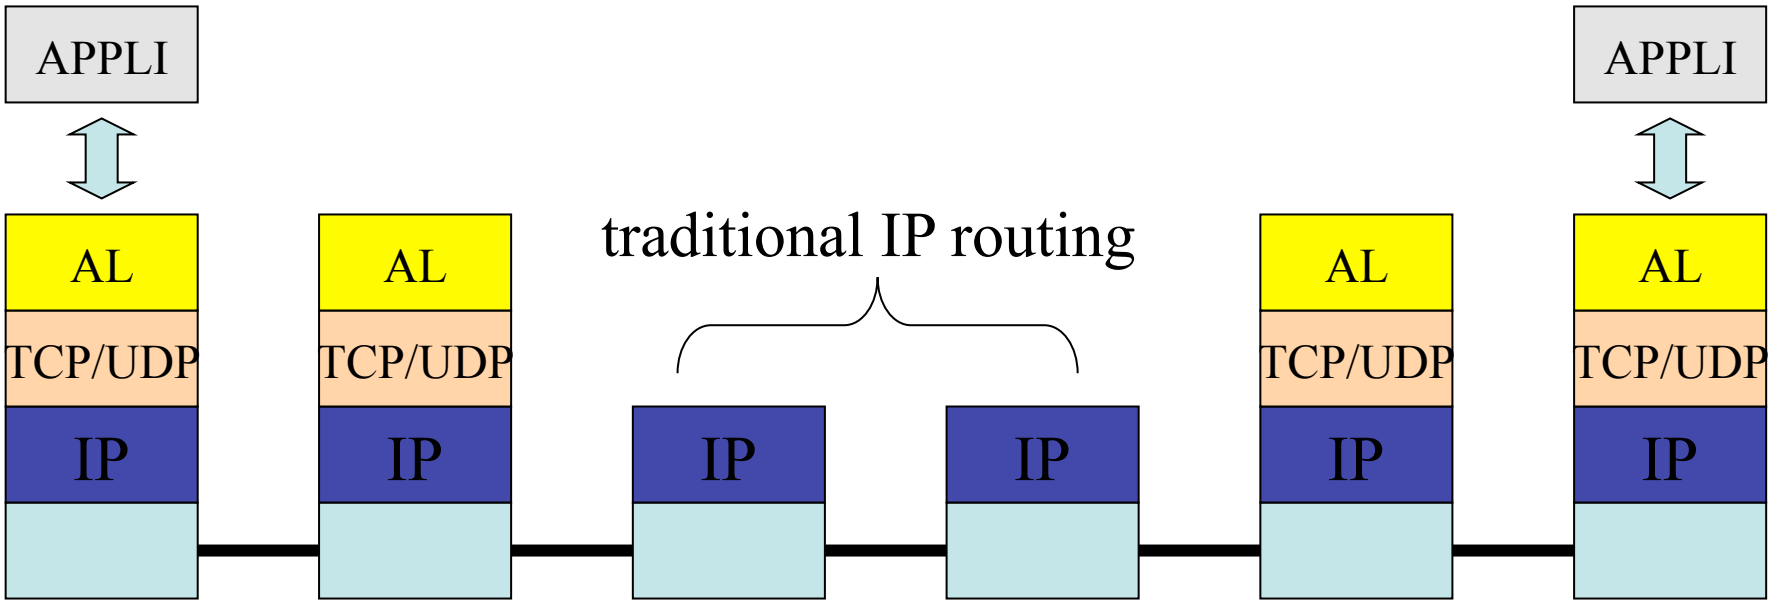

Active network revisited

- Software Defined Networking (SDN) wants to decouple the control plane from the data plane
- Somehow similar to discrete active/ programmable network concepts
- Better perception from the user because controlled by operators and hardware vendors

How to upgrade the Internet for QoS?

- Approach: de-couple end-system evolution from network evolution
- End-to-end protocols: TCP, RTP, H.323, etc to spur the growth of adaptive multimedia applications
  - Assume best-effort or better-than-best-effort clouds
- Network protocols: IntServ, DiffServ, RSVP, MPLS, COPS …
  - To support better-than-best-effort capabilities at the network (IP) level

Principles for QOS Guarantees

- Consider a phone application at 1Mbps and an FTP application sharing a 1.5 Mbps link.
  - bursts of FTP can congest the router and cause audio packets to be dropped.
  - want to give priority to audio over FTP
- **PRINCIPLE 1: Marking of packets is needed for router to distinguish between different classes; and new router policy to treat packets accordingly**

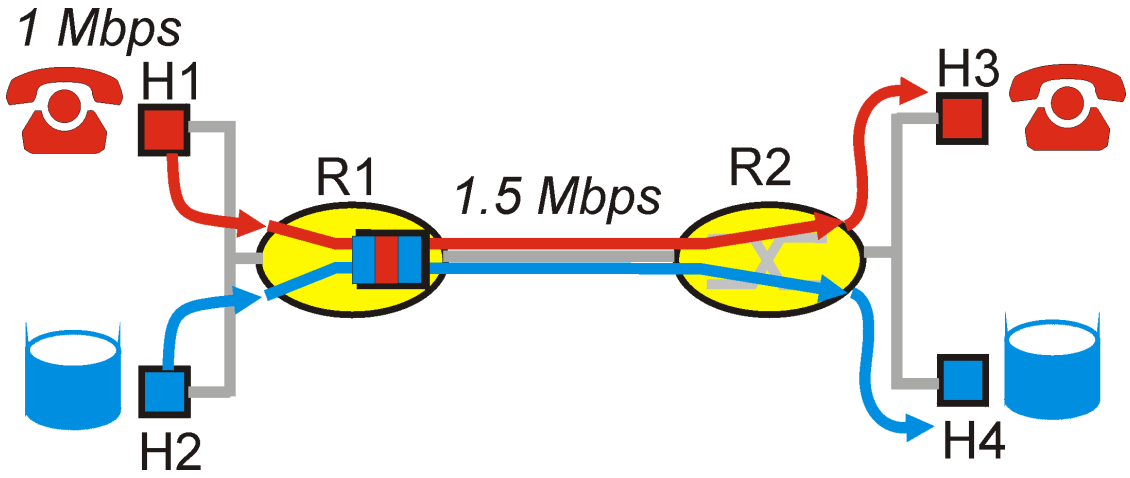

Principles for QOS Guarantees (more)

- Applications misbehave (audio sends packets at a rate higher than 1Mbps assumed above);
- **PRINCIPLE 2: provide protection (isolation) for one class from other classes**
- Require Policing Mechanisms to ensure sources adhere to bandwidth requirements; Marking and Policing need to be done at the edges:

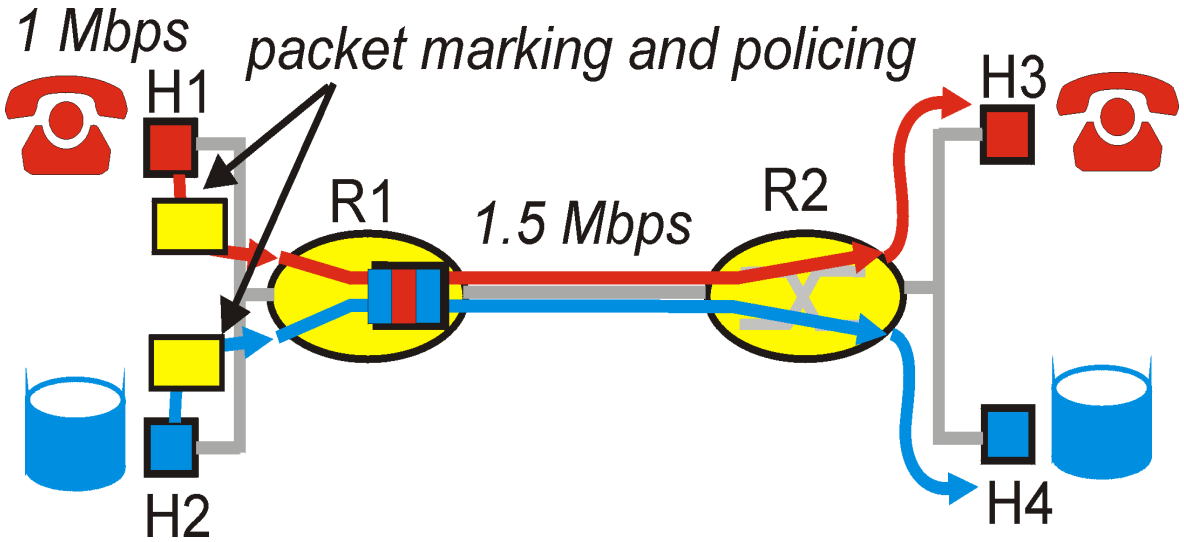

Principles for QOS Guarantees (more)

- Alternative to Marking and Policing: allocate a set portion of bandwidth to each application flow; can lead to inefficient use of bandwidth if one of the flows does not use its allocation
- **PRINCIPLE 3: While providing isolation, it is desirable to use resources as efficiently as possible**

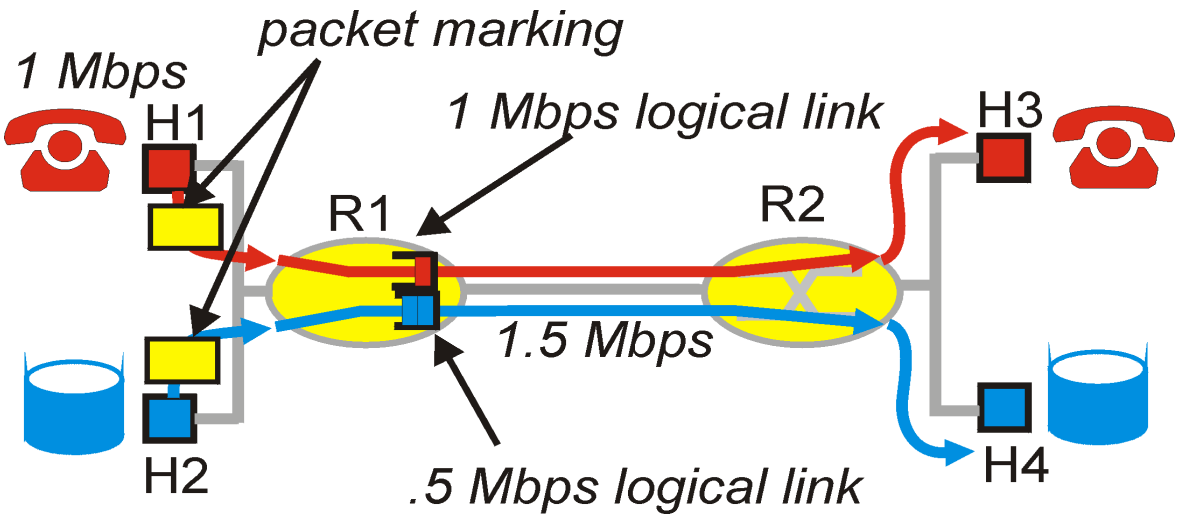

Generic Router

- Route packet and store in output buffer
- Decide when and what packet to send on output link

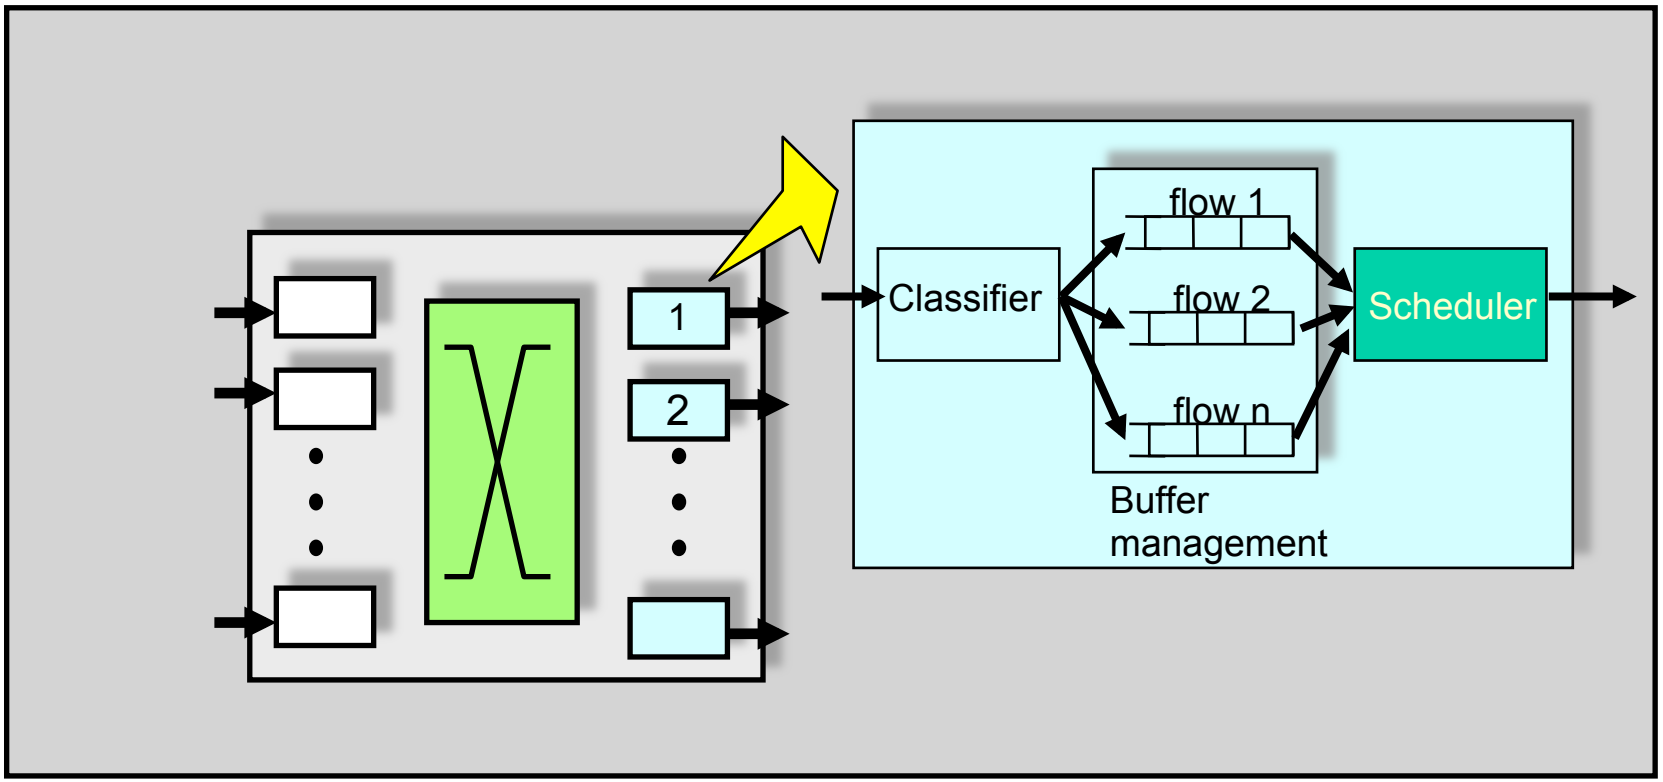

SCHEDULING

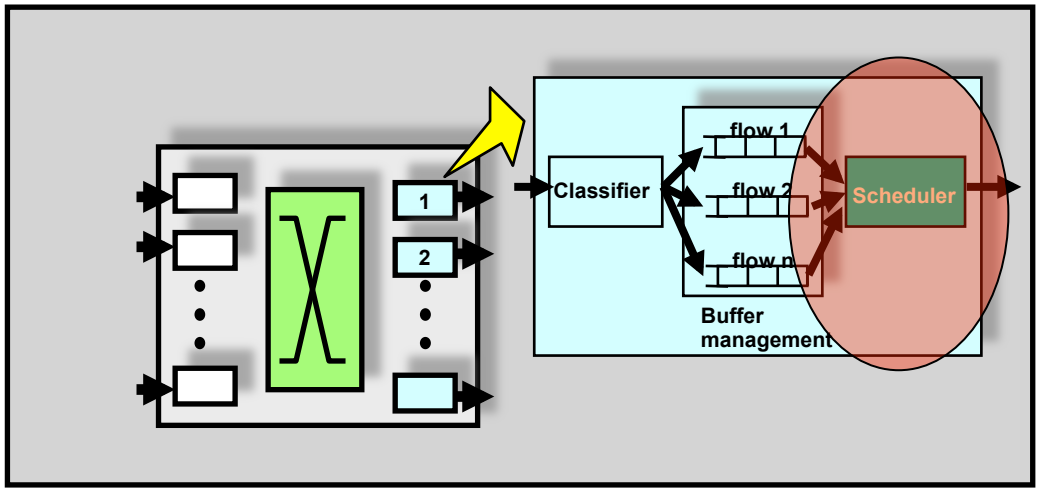
1
2
Classifier
flow 1
flow 2
flow n
Scheduler
Buffer
management

Round Robin (RR)

- Round Robin: scan class queues serving one from each class that has a non-empty queue

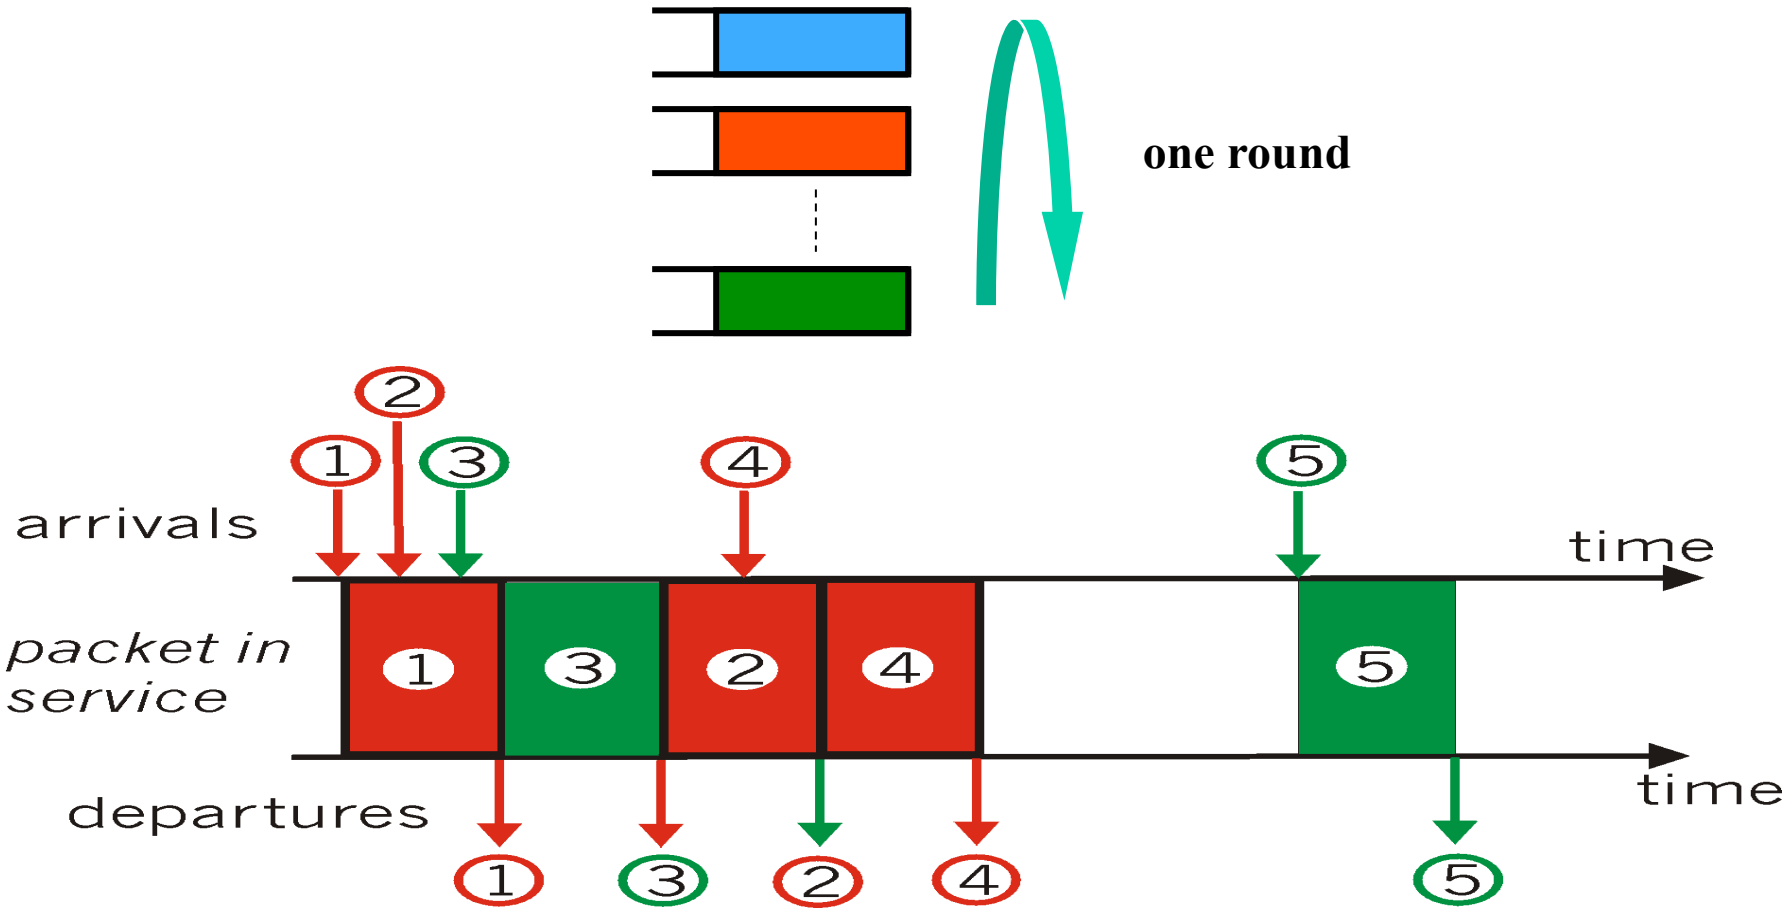

BUFFER MANAGEMENT

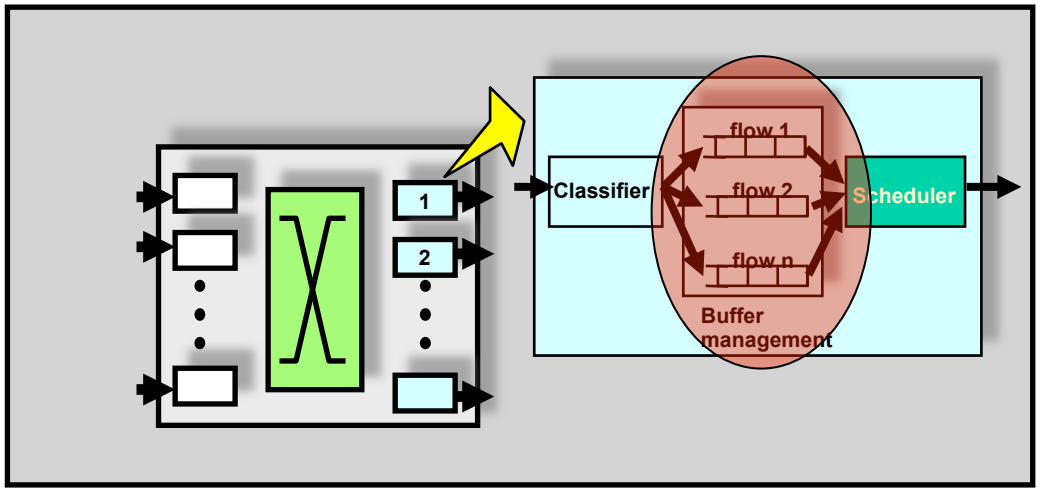

Typical Internet Queuing

- FIFO + drop-tail
  - Simplest choice
  - Used widely in the Internet
- FIFO (first-in-first-out)
  - Implies single class of traffic
- Drop-tail
  - Arriving packets get dropped when queue is full regardless of flow or importance
- FIFO Issues:
  - No isolation between flows: full burden on e2e control
  - No policing: send more packets → get more service
- Drop-tail issues:
  - Synchronization: end hosts react to same events because packets tend to be lost in bursts (see TCP!!)

Random Early Detection (RED)

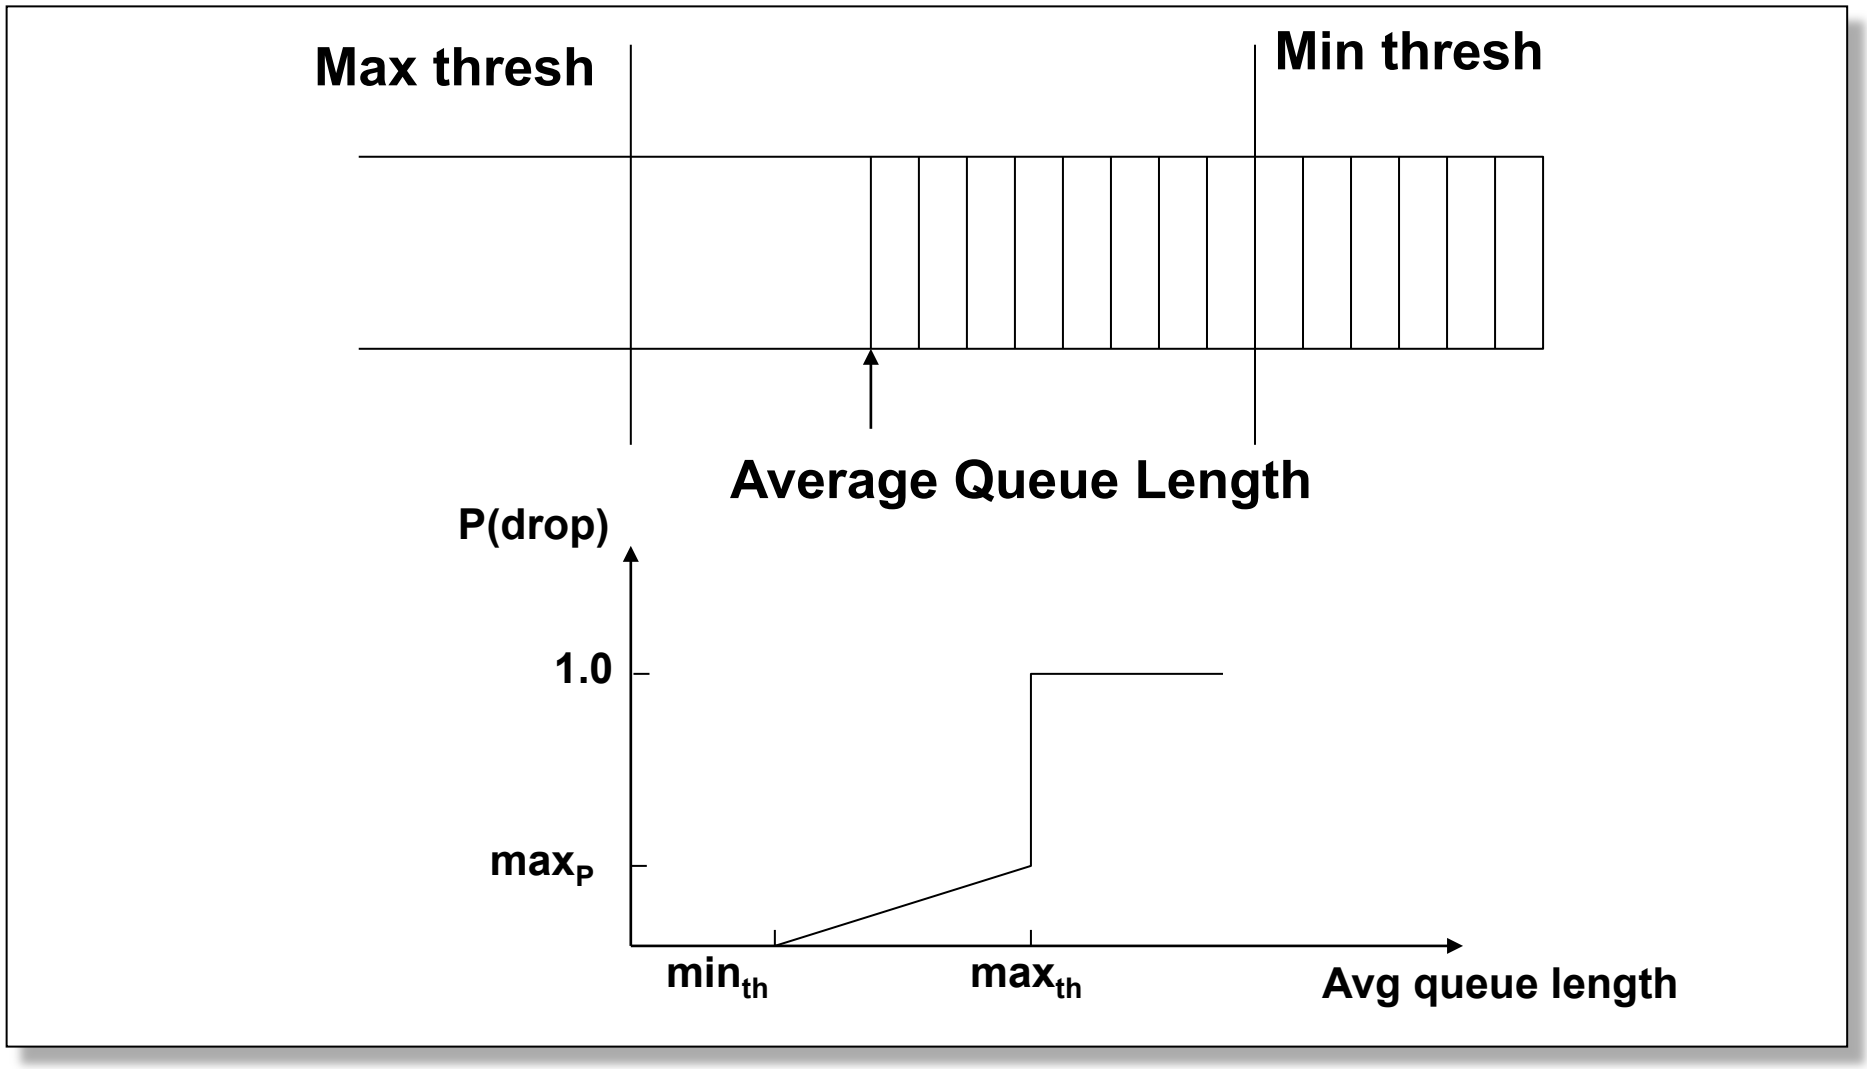

RED with Multiple Thresholds

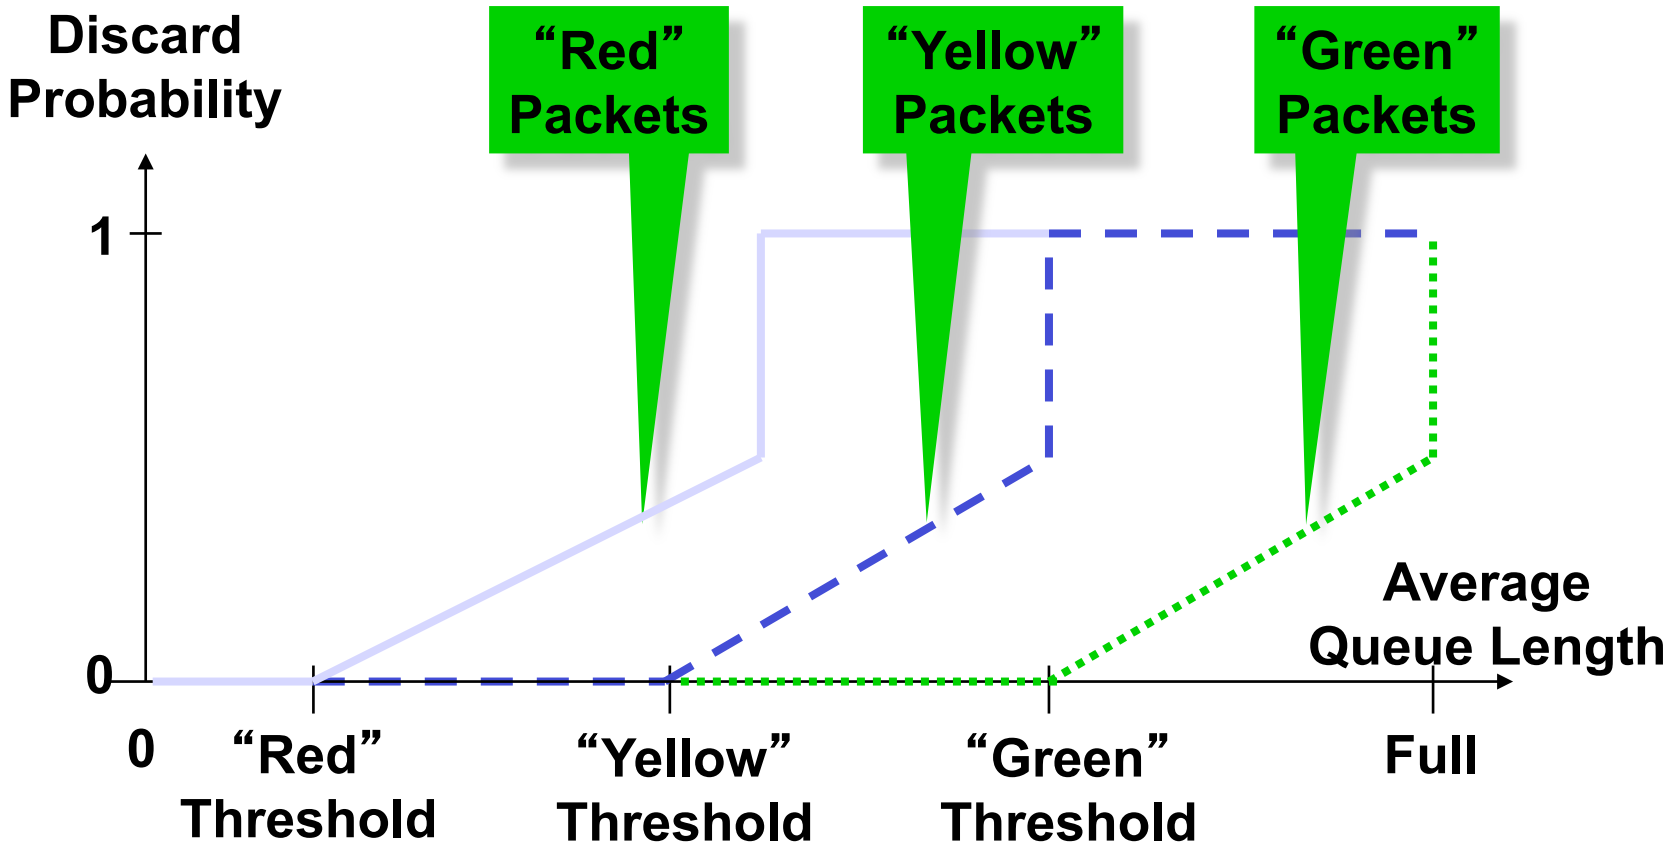

source Juha Heinänen

TRAFFIC, SERVICE CHARACTERIZATION,

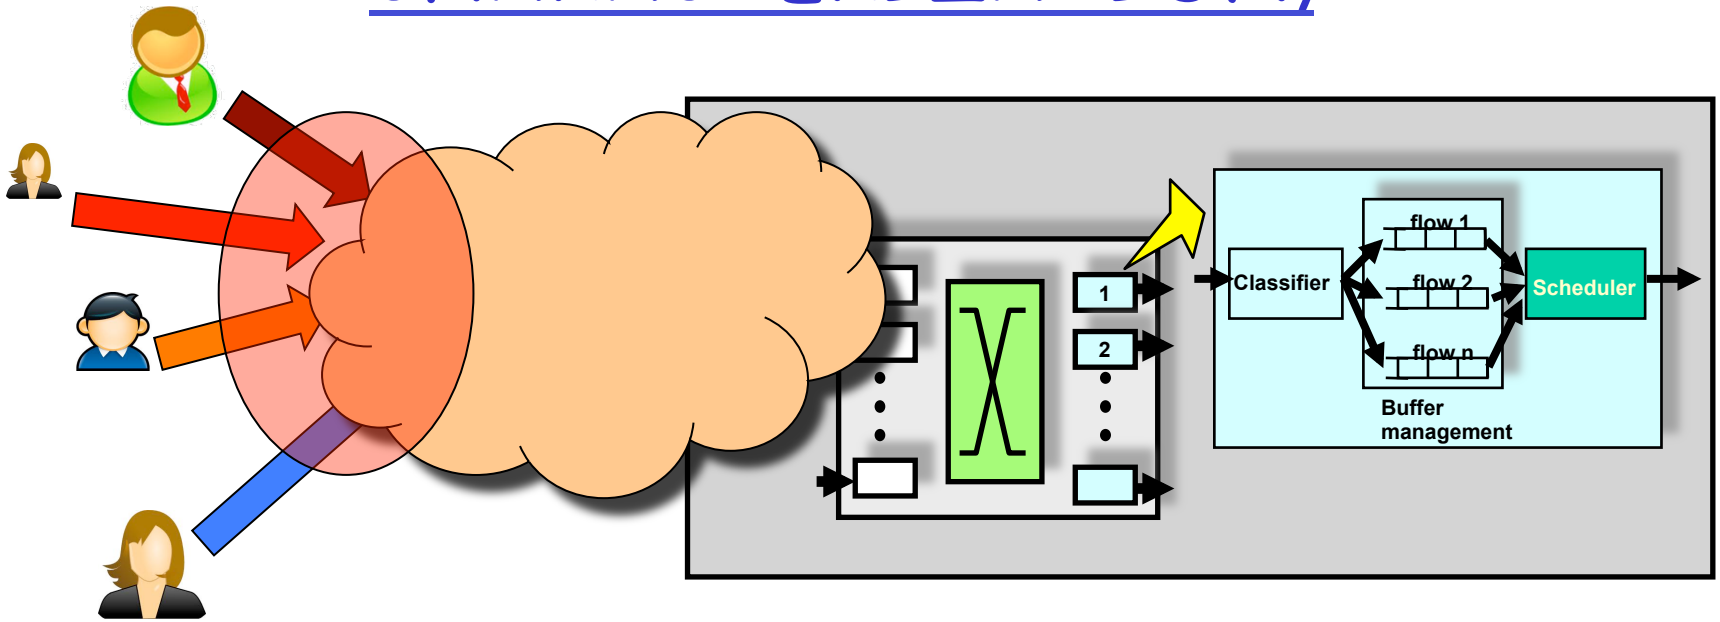

Traffic Conditioning

- User declares traffic profile (eg, rate and burst size); traffic is metered and shaped if non-conforming

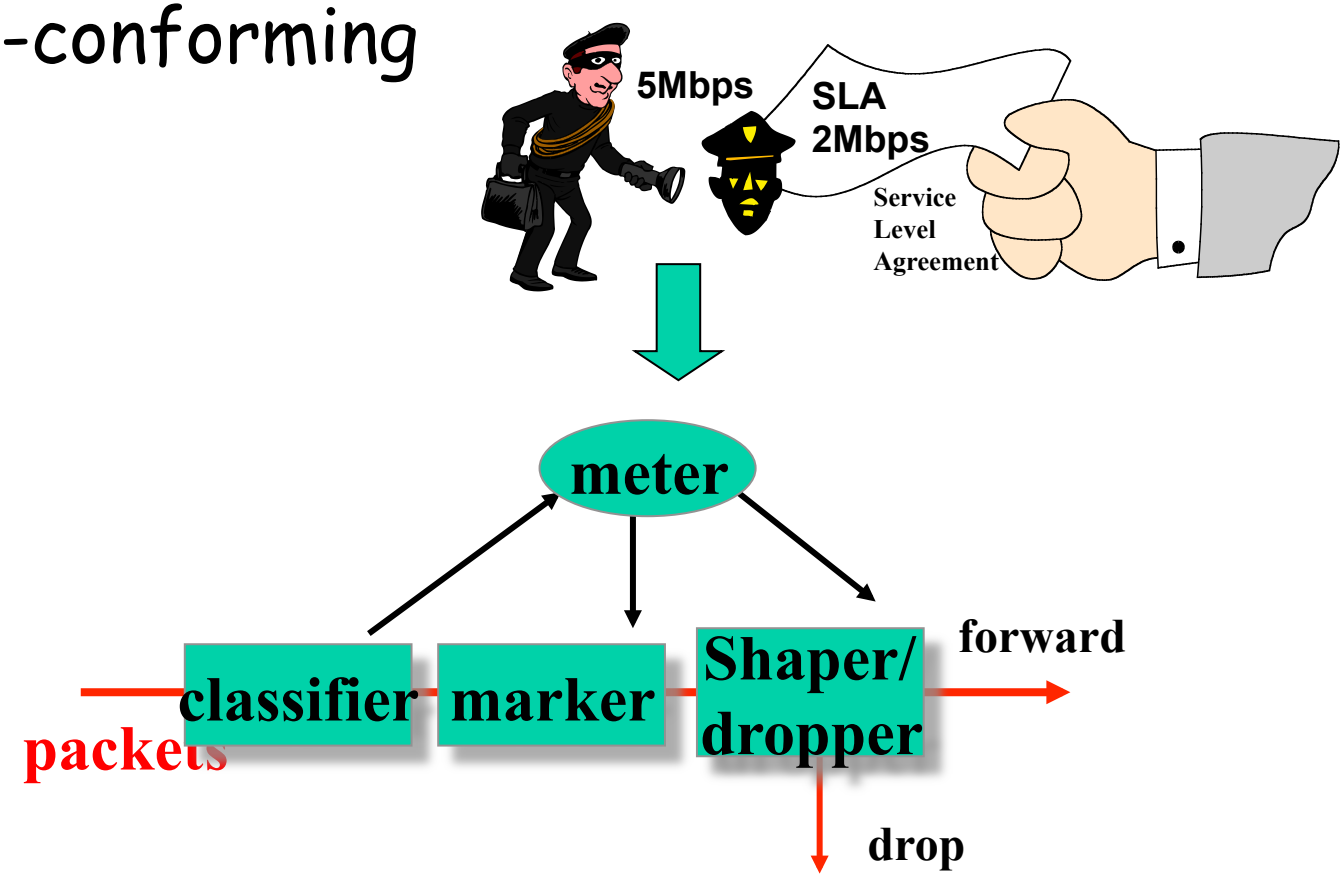

Ex: Token Bucket

- Characterized by three parameters (b, R, C)
  - b – token depth
  - R – average arrival rate
  - C – maximum arrival rate (e.g., link capacity)
- A bit is transmitted only when there is an available token
  - When a bit is transmitted exactly one token is consumed

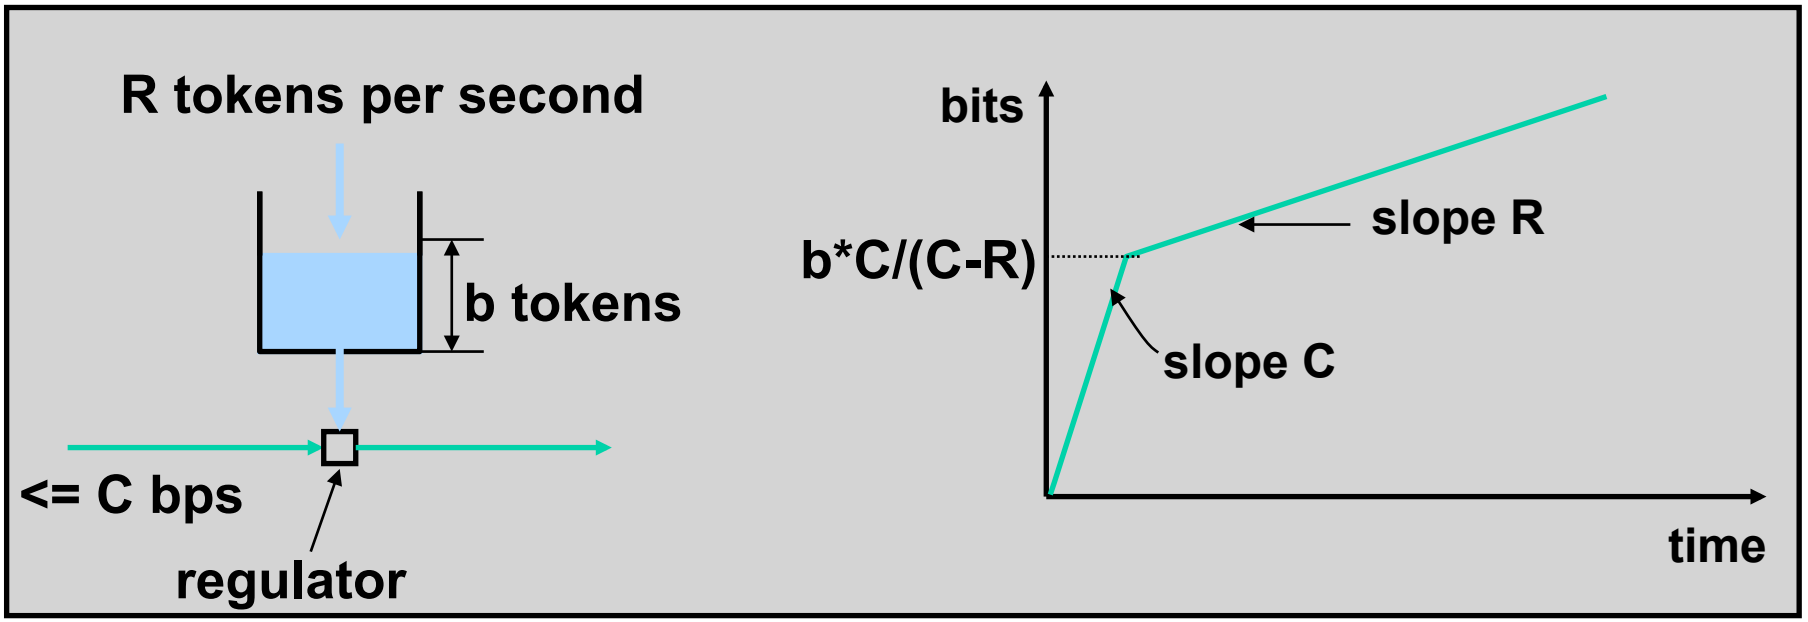

Traffic Envelope (Arrival Curve)

- Maximum amount of service that a flow can send during an interval of time t

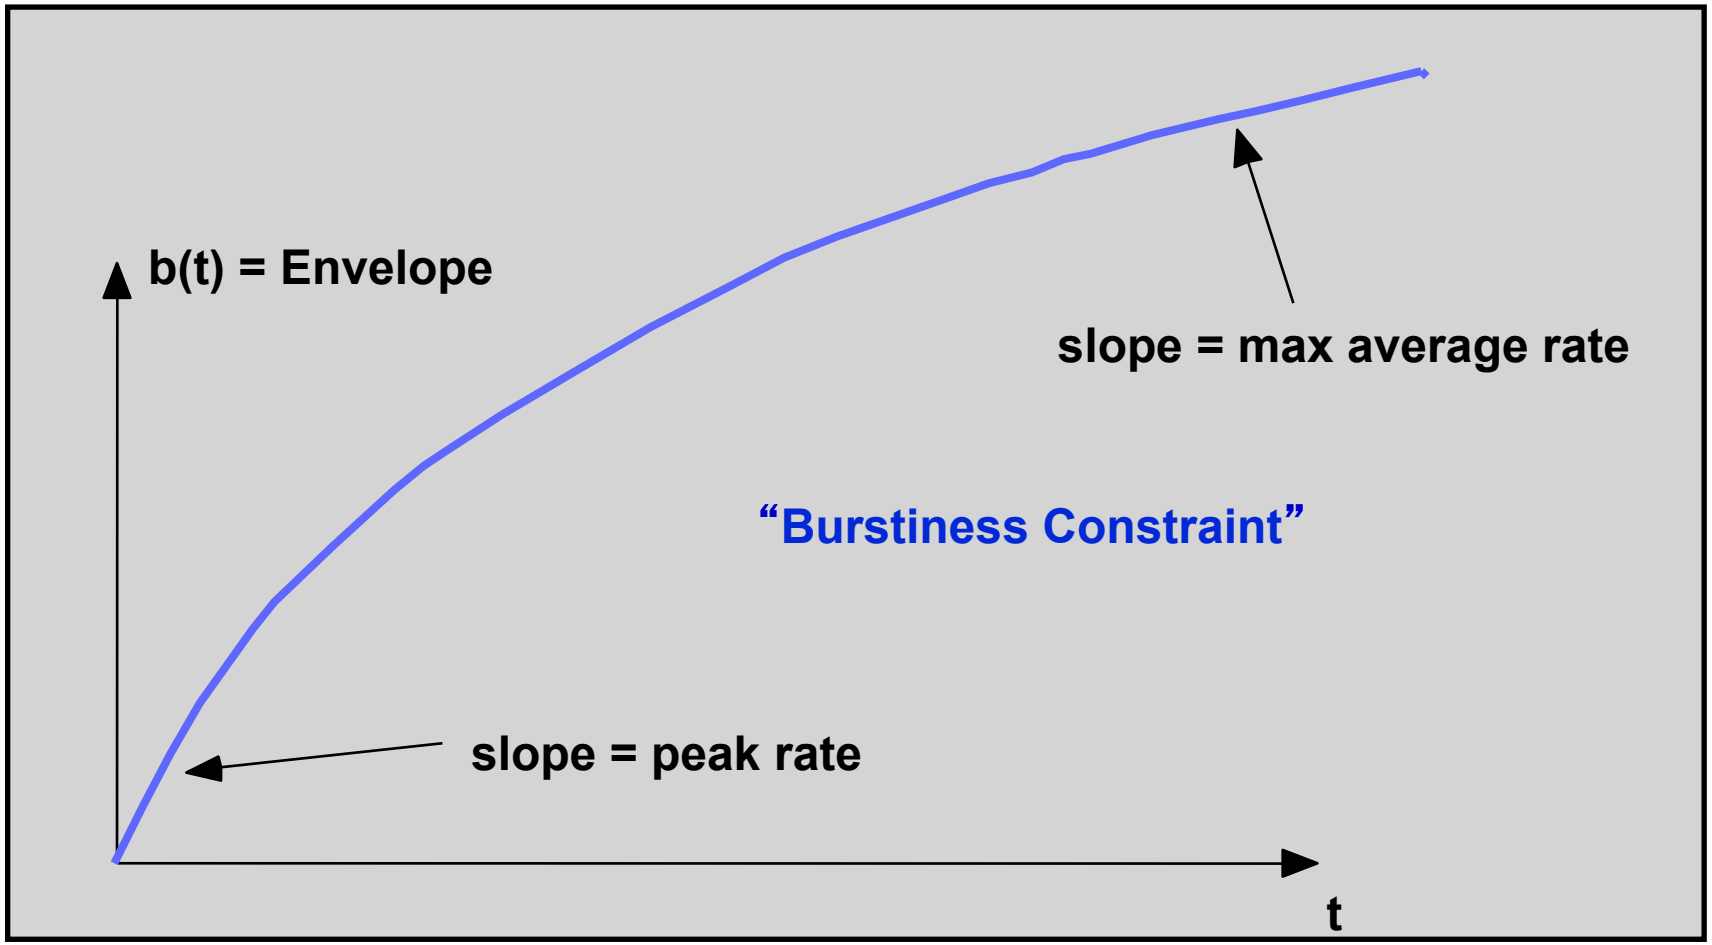

Token Bucket

Example

- B = 4000 bits, R = 1 Mbps, C = 10 Mbps
- Packet length = 1000 bits
- Assume the bucket is initially full and a "large" burst of packets arrives

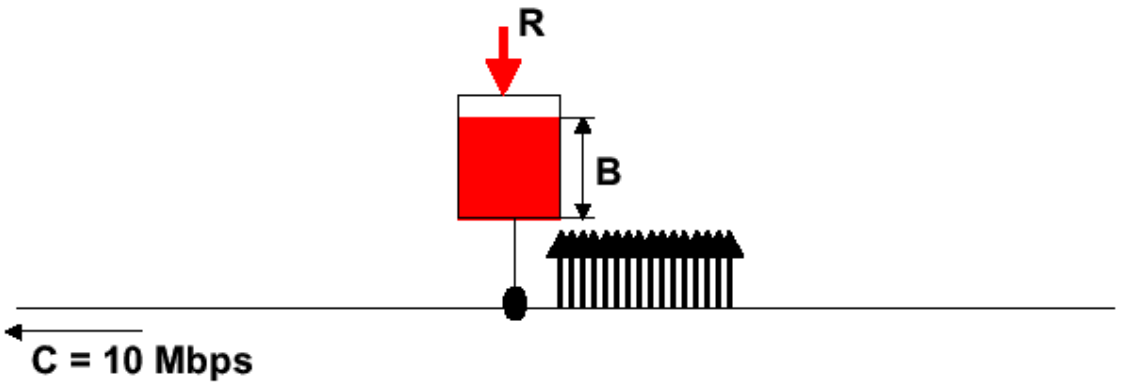

istoica@cs.cmu.edu

Token Bucket

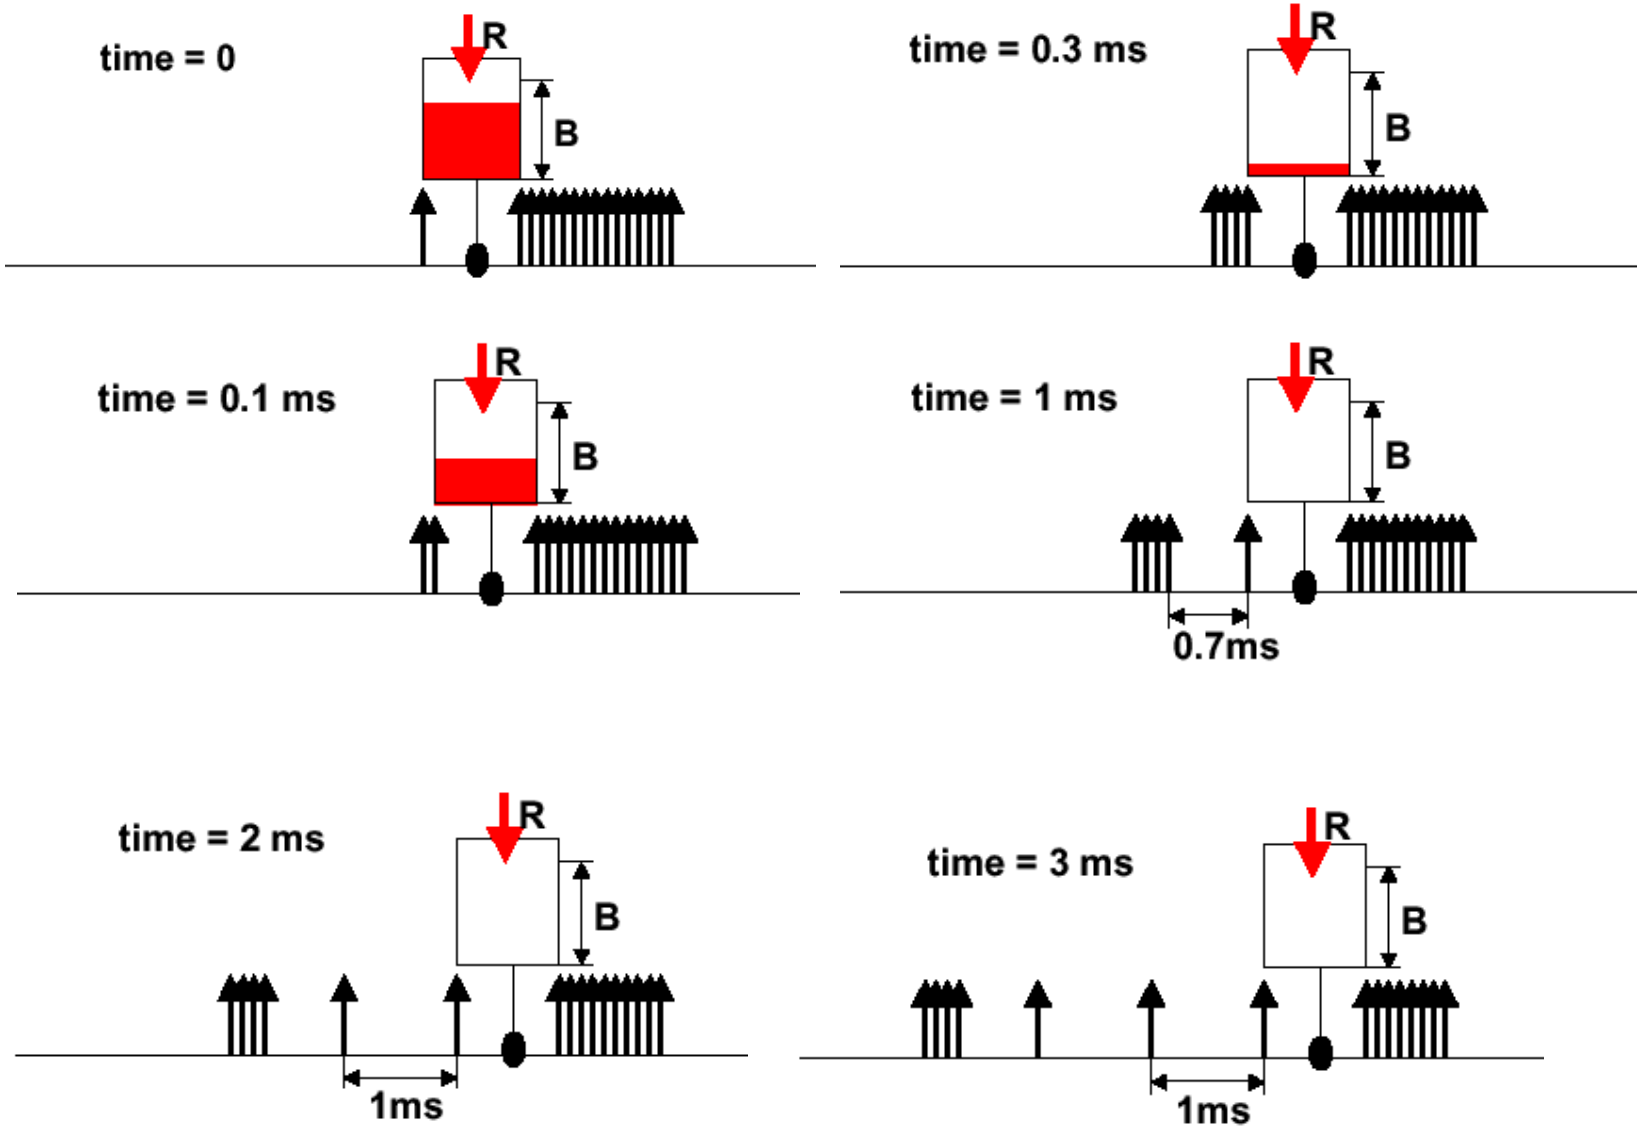

Arrival curve

A(t) – number of bits received up to time t

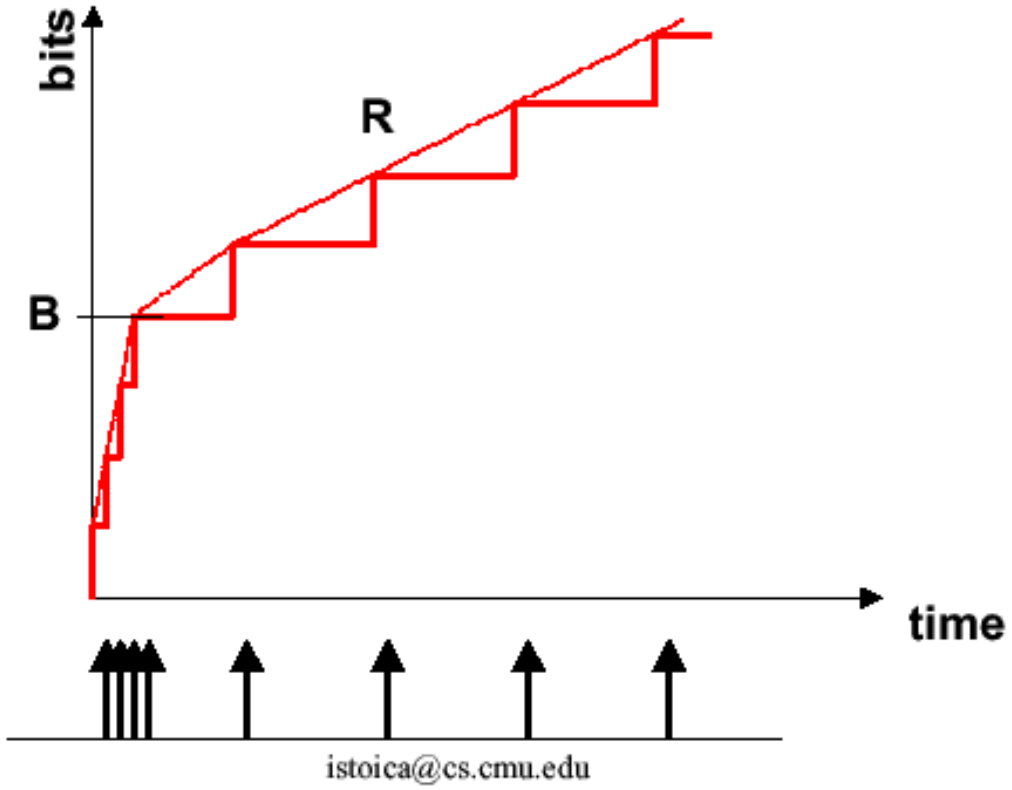

istoica@cs.cmu.edu

Characterizing a Source by Token Bucket

- Arrival curve – maximum amount of bits transmitted by time t
- Use token bucket to bound the arrival curve

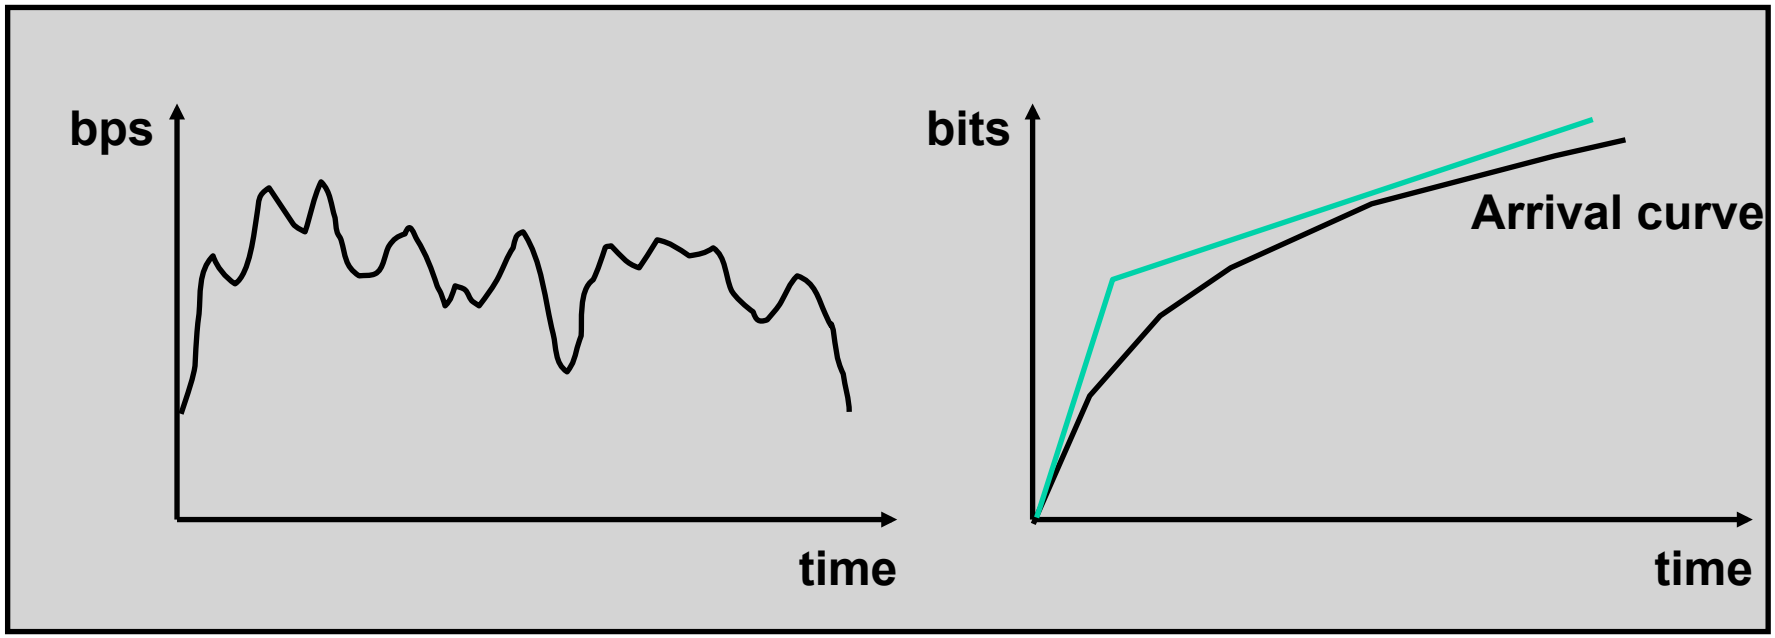

Arrival and Departure Process

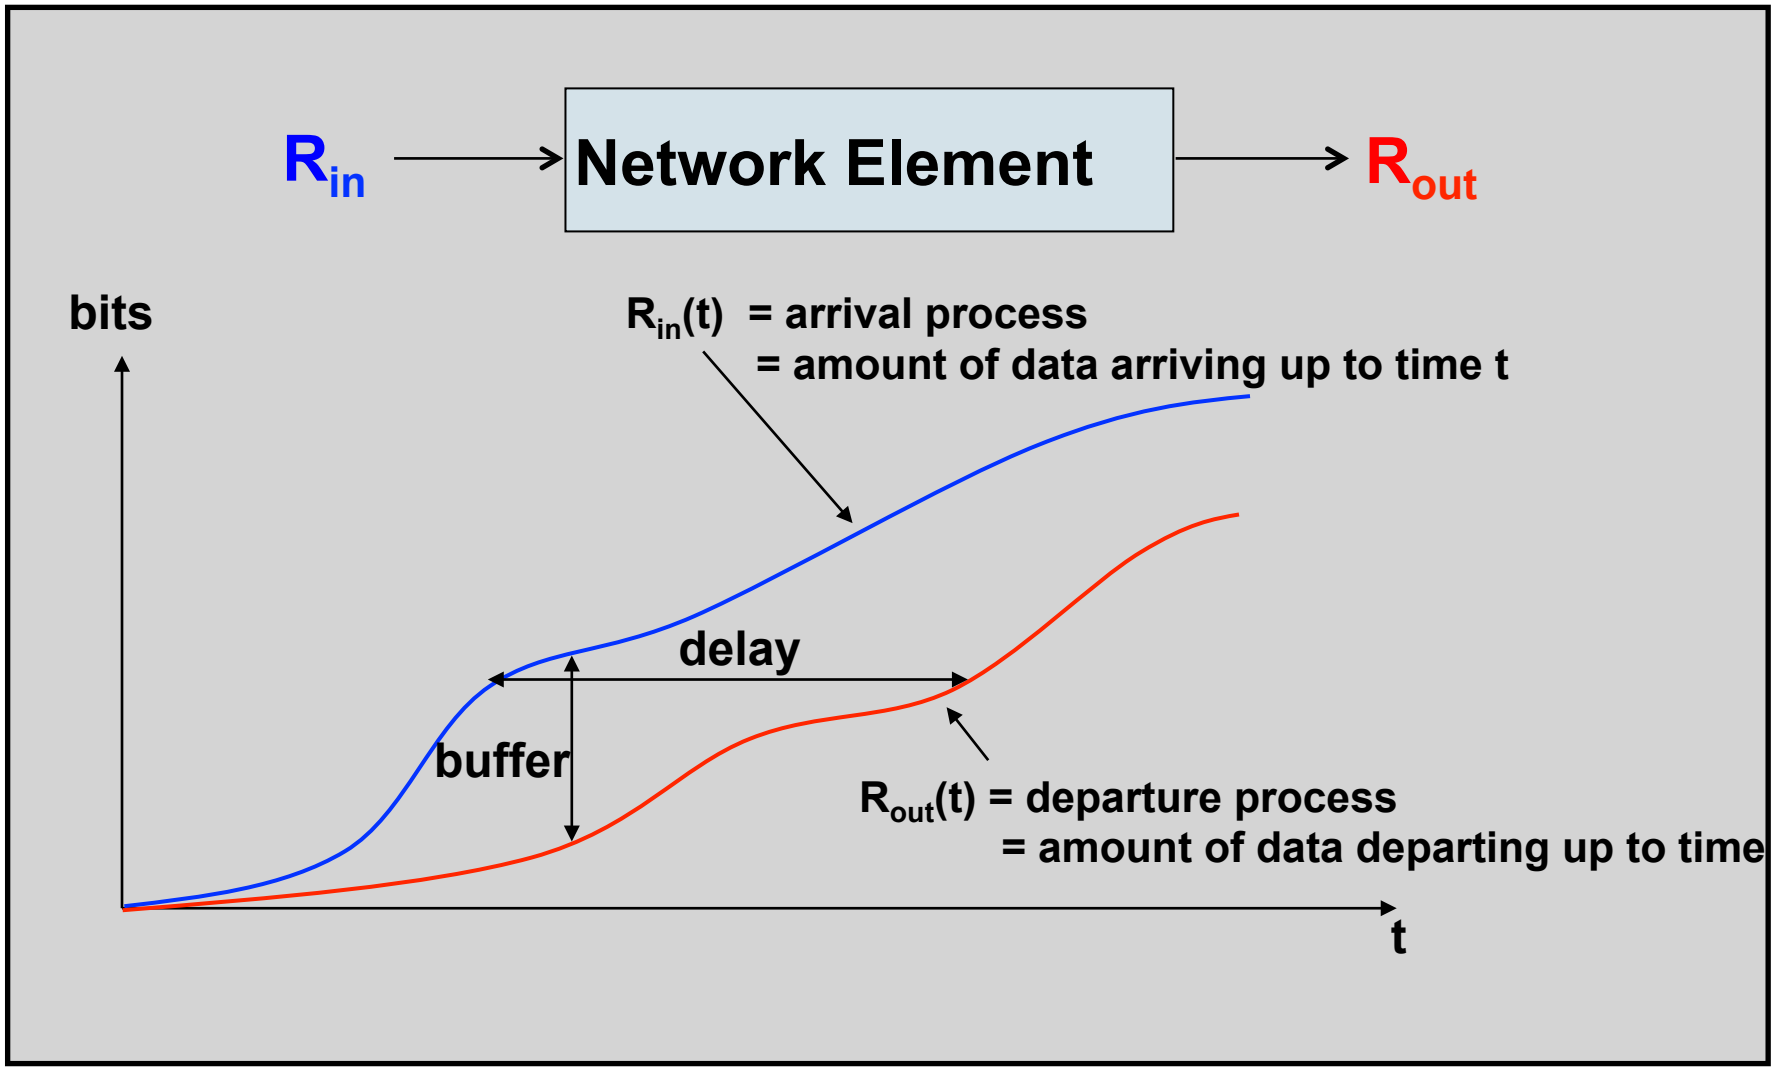

Service Differentiation

The real question is to choose which packets shall be dropped. The first definition of differential service is something like "not mine."
-- Christian Huitema

- Differentiated services provide a way to specify the relative priority of packets
- Some data is more important than other
- People who pay for better service get it!

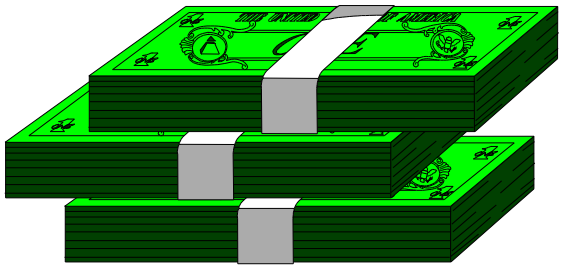

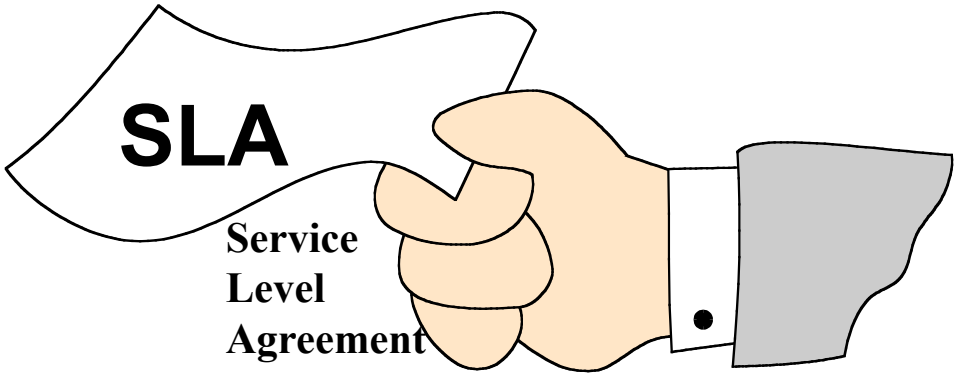

Divide traffic into classes

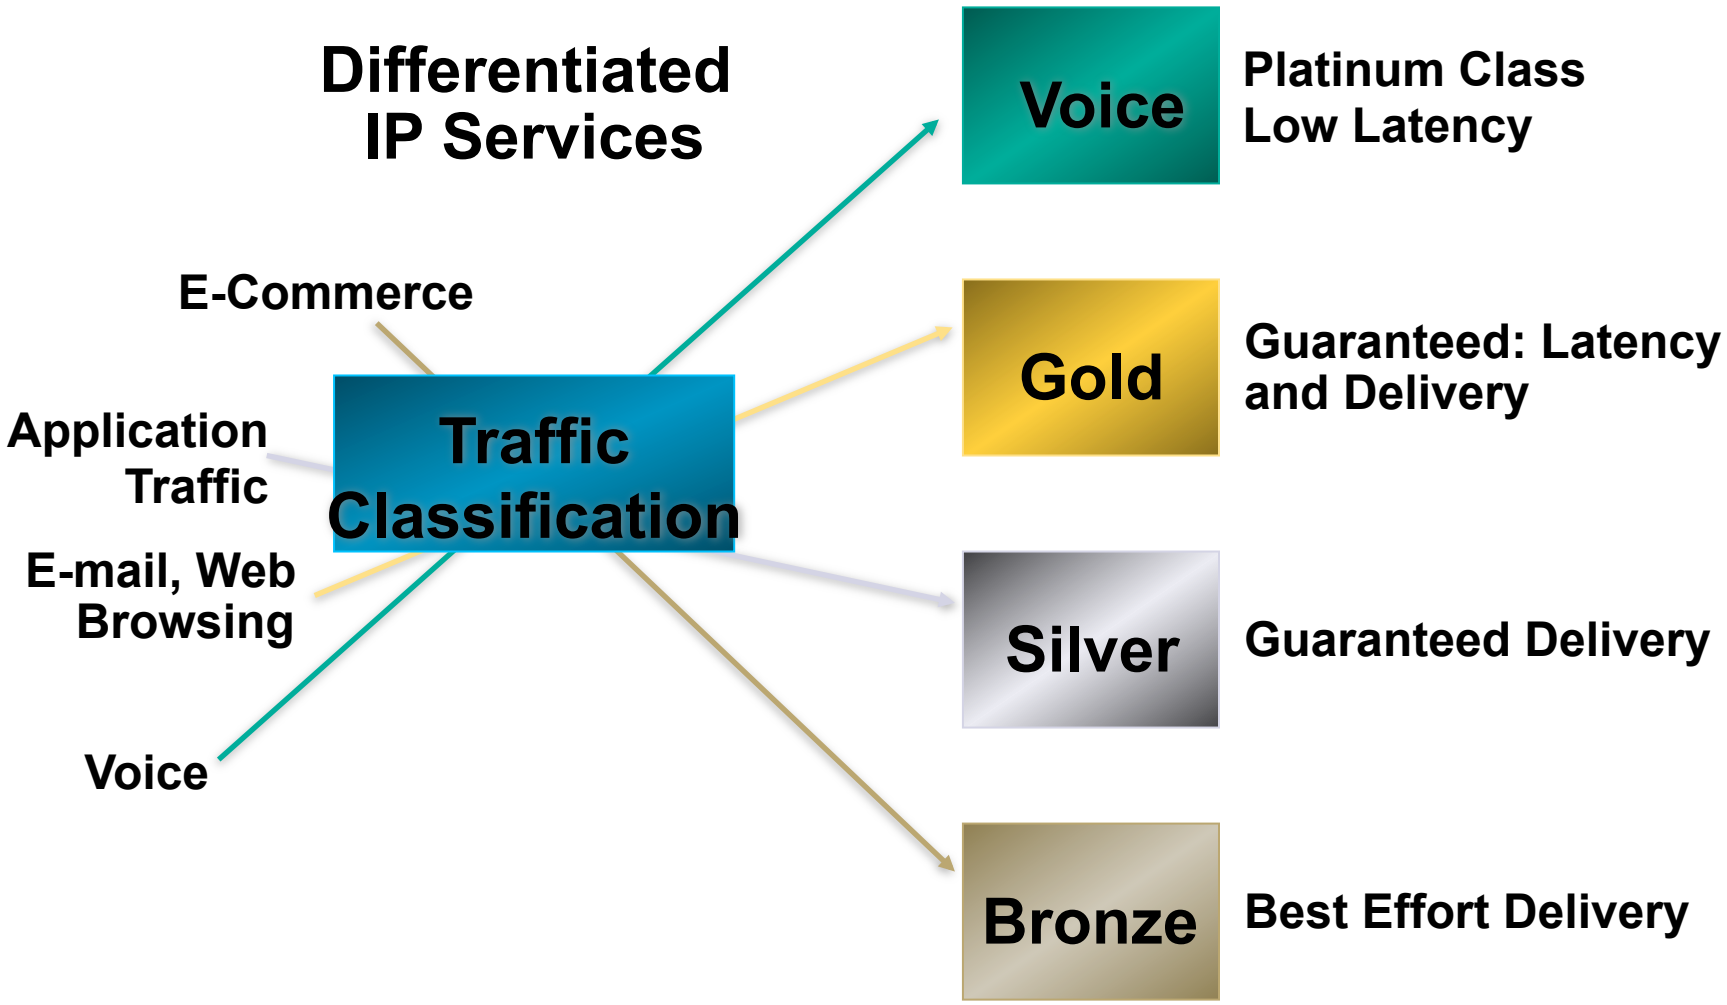

Borrowed from Cisco

DiffServ building blocks

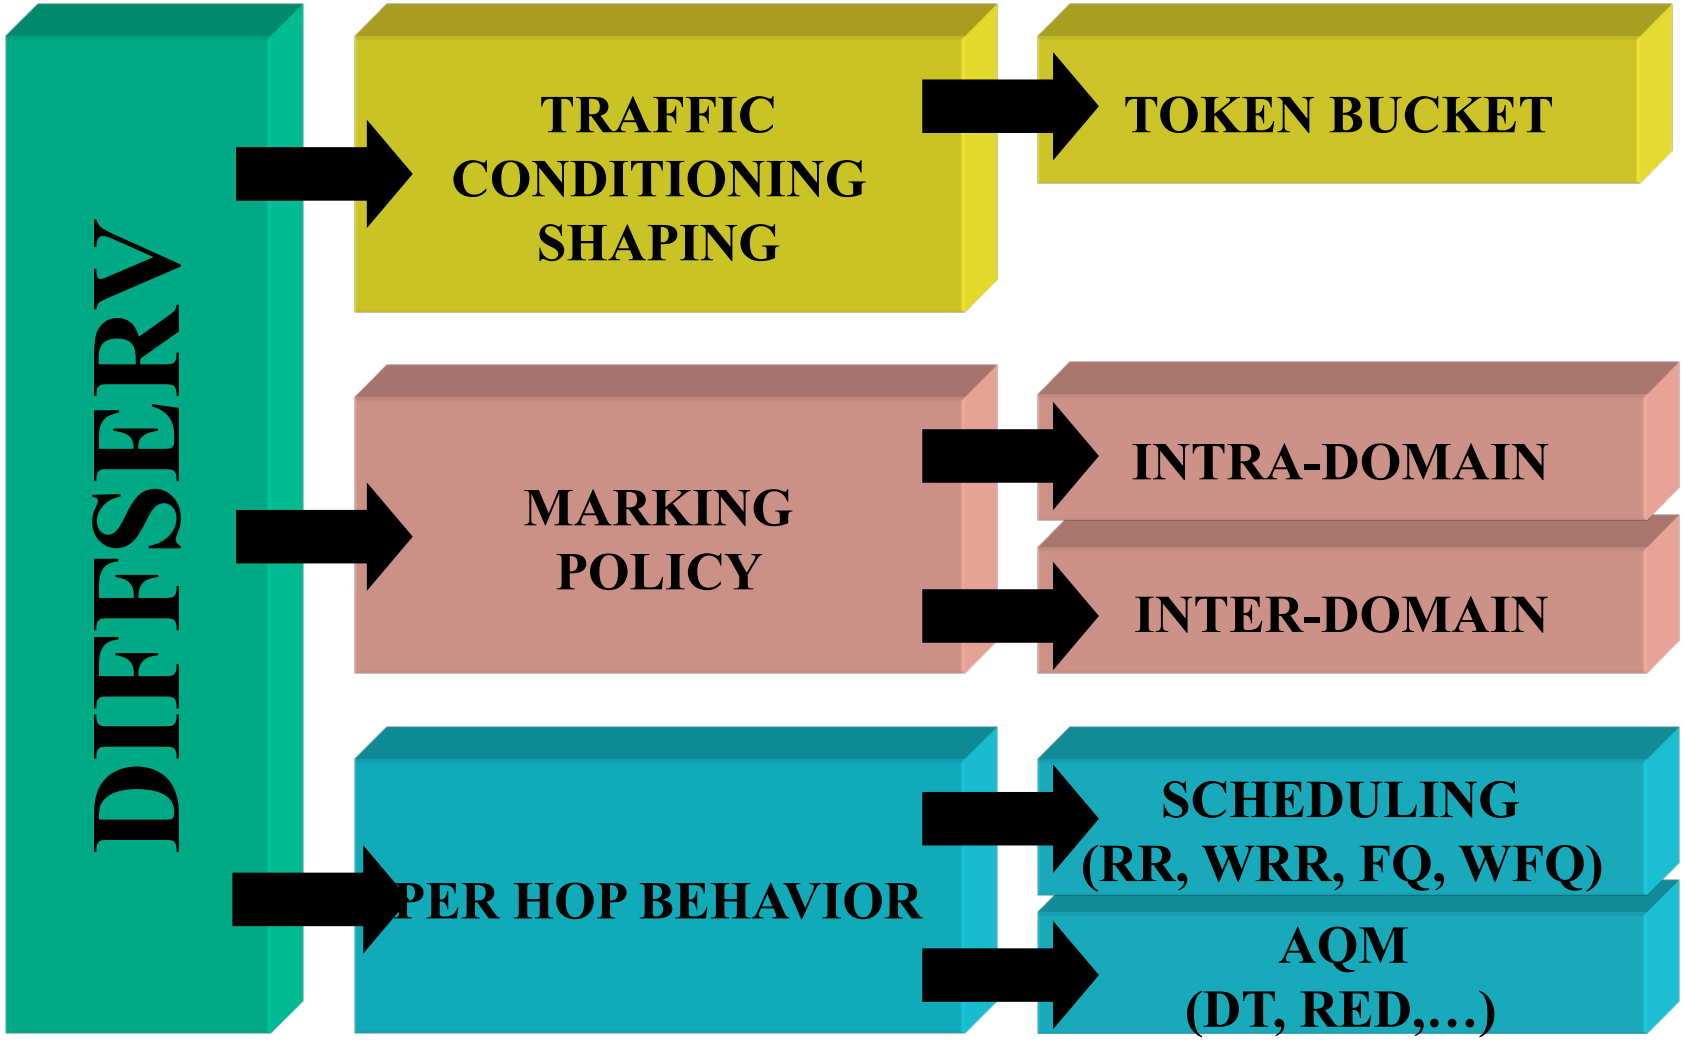

Summary
Towards IP/DiffServ/(G)MPLS/DWDM

From cisco