Course on Computer Communication and Networks

Lecture 15
Summary (flashback) and Projection (related topics – continuation of study)

EDA344/DIT 420, CTH/GU

Based on the book Computer Networking: A Top Down Approach, Jim Kurose, Keith Ross, Addison-Wesley.

Important for the exam

When/where: wednesday March 16, 14.00-18.00, M-building

You may have with you:

- English-X dictionary
- no calculators, PDAs, etc (if/where numbers matter, do rounding)

Grading

- 30-40, 41-50, 51-60 (out of 60)= 3, 4, 5 (CTH)
- 30-44, 45-60 (out of 60) = G, VG (GU)

To think during summary-study

Have overview, critical eye; explain; ask yourselves: why is this so? / how does it work (or not work)?

Flashback

Principles, Organisation

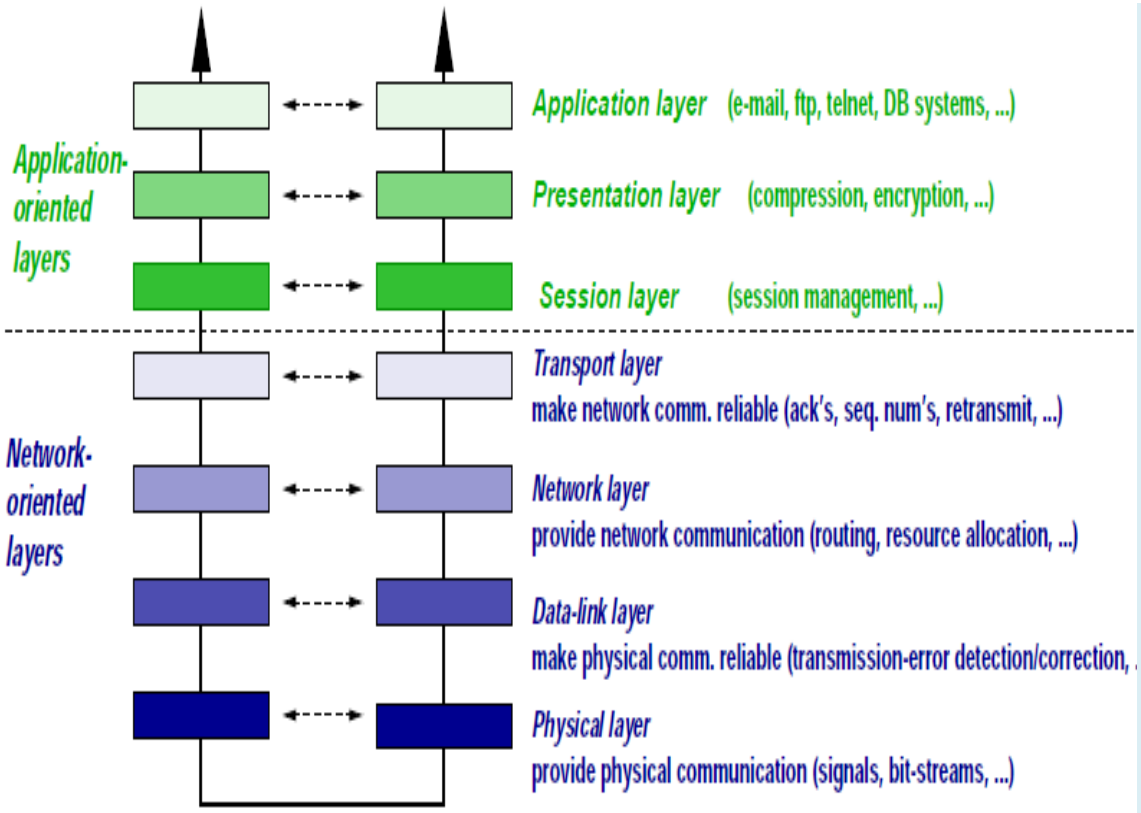

Network Problems (in the order faced in the 1st intro):

- Mobility, performance, security, …, …
- serving different types of traffic,
- connecting transparently different networks,
- routing, congestion control,
- access to shared (broadcast) transmission medium
- producer-consumer problems, flow and error control

Layering : principle, why

Highlights

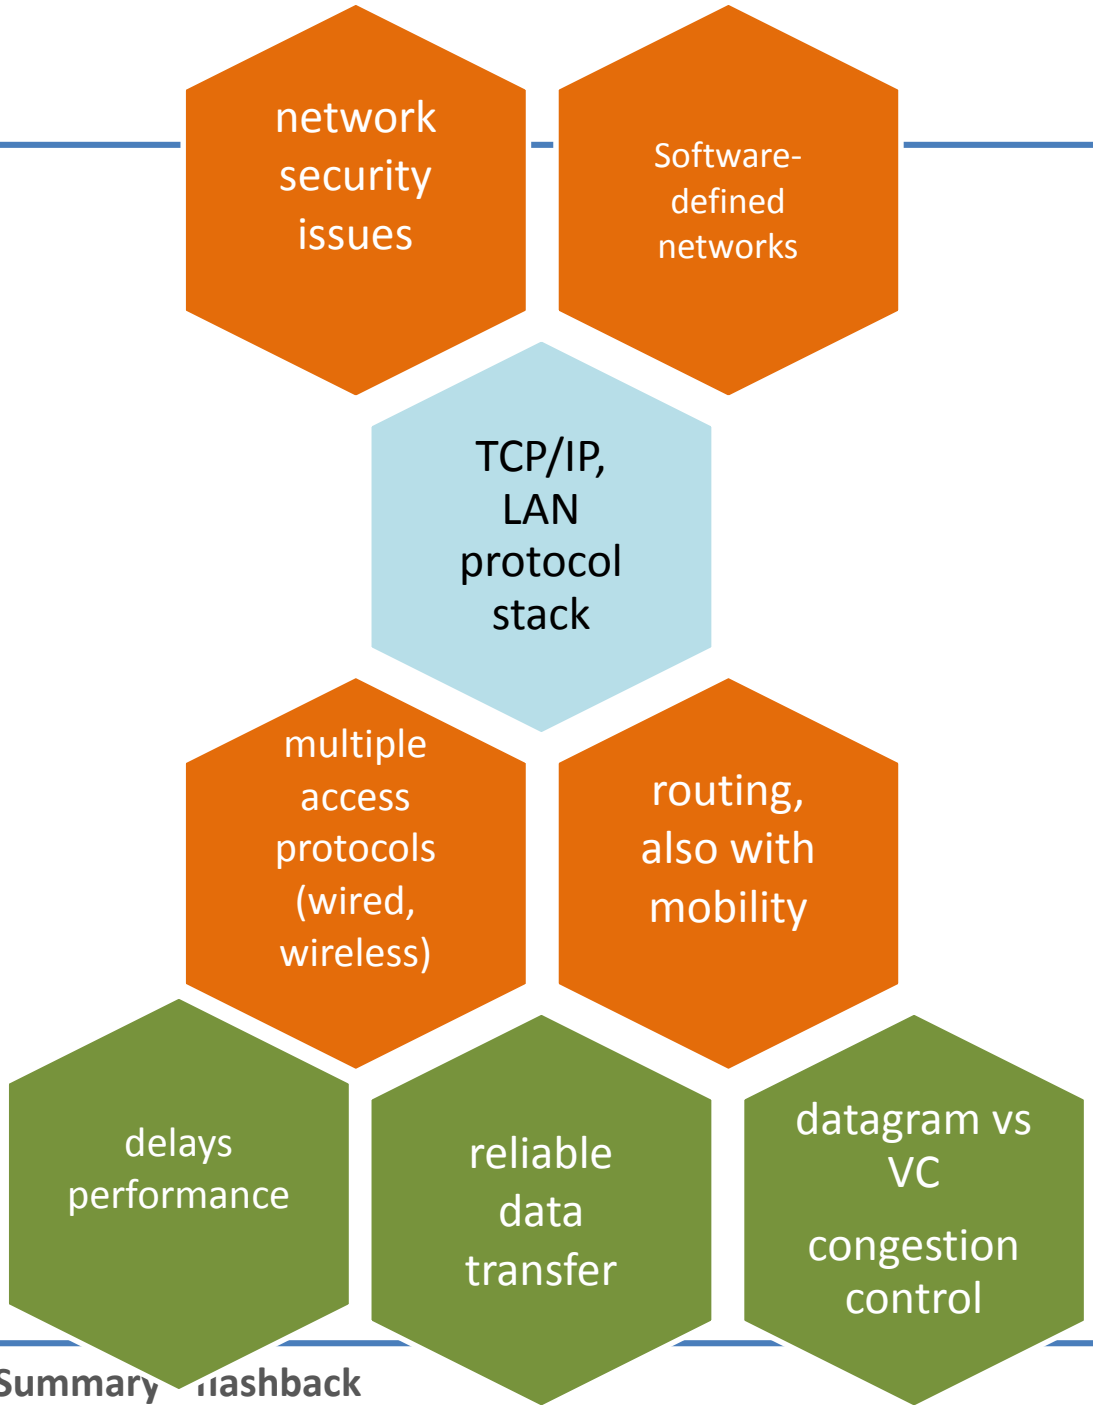
network security issues
Software-defined networks
TCP/IP, LAN protocol stack
multiple access protocols (wired, wireless)
routing, also with mobility
delays performance
reliable data transfer
datagram vs VC
congestion control

Types of delay; performance

delays
performance

- Propagation, transmission, queueing, processing
- Throughput -- effective bandwidth
- Utilization -- efficiency
- Packet-switching: impact of store&forward, pipelines, space-time diagrams
- Sliding windows performance
- Relation between delays-losses

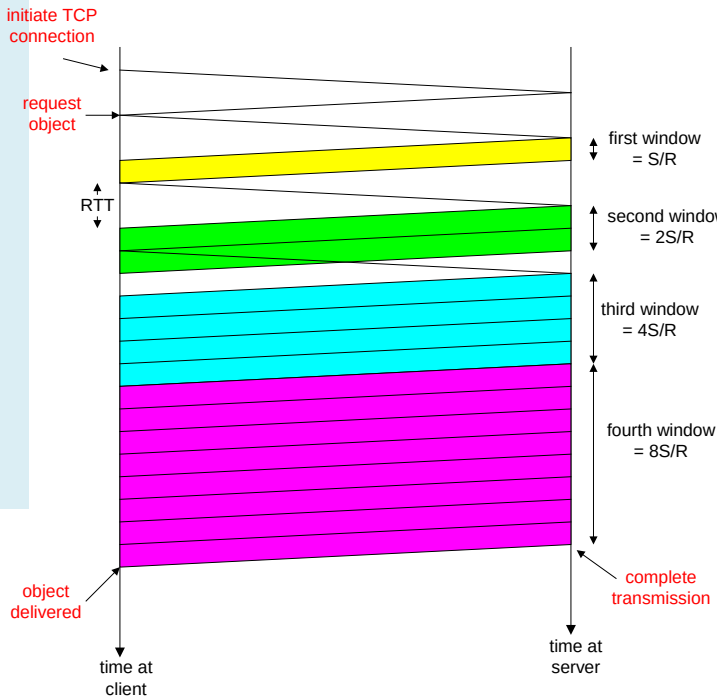

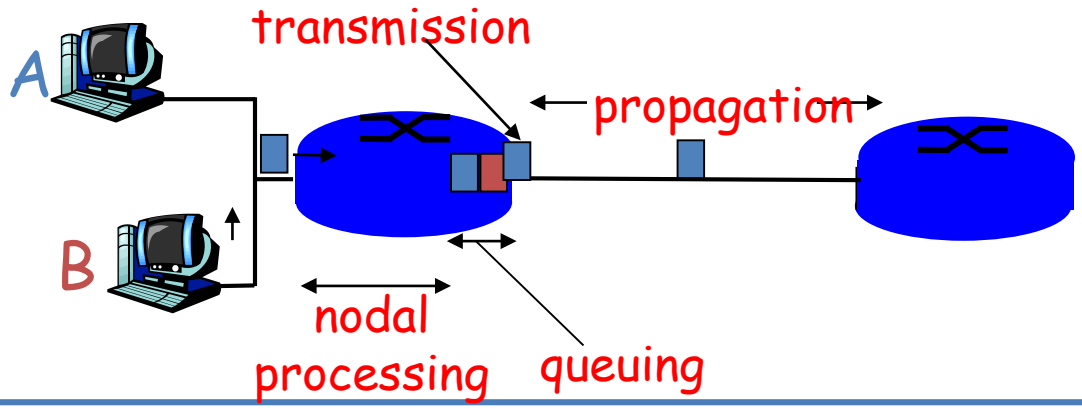

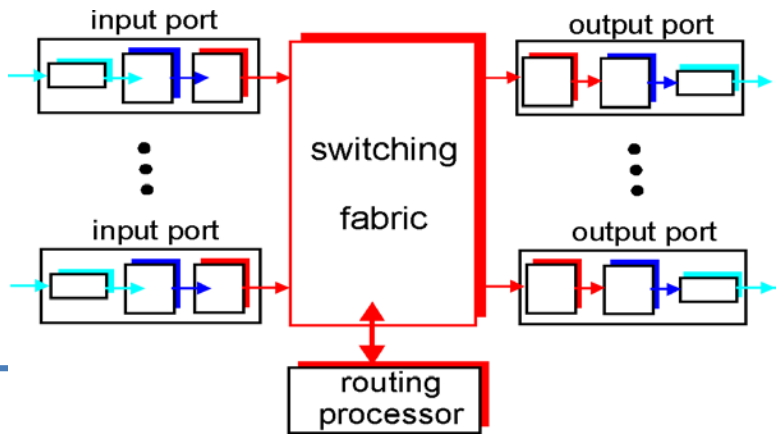

Reliable data transfer

reliable data transfer

Guaranteed, in-order, correct delivery:

- stop&wait
- sliding windows
- sequence numbers
- window sizes
- dynamic windows (TCP)
- performance
- **Flow control**

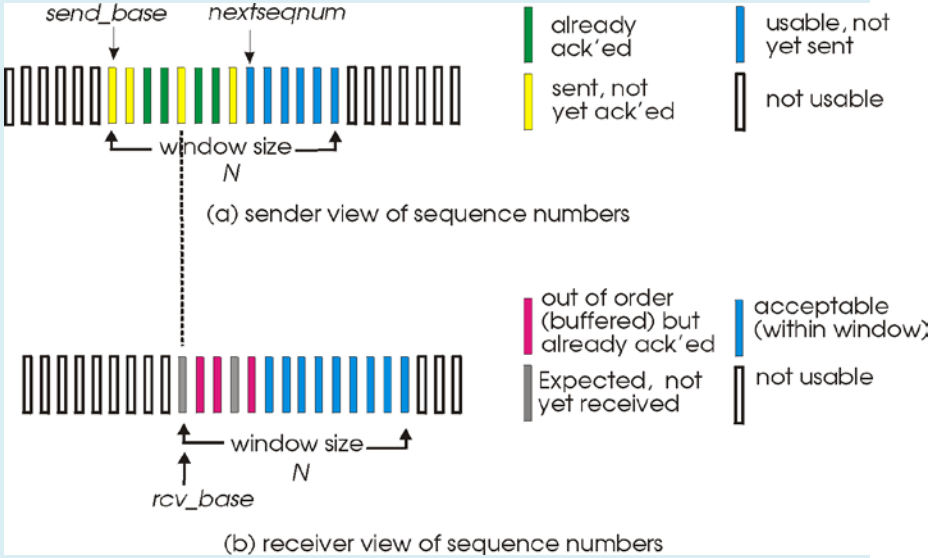

- **Error detection**: checksums
- **Error control**: go-back-n, selective repeat, FEC methods

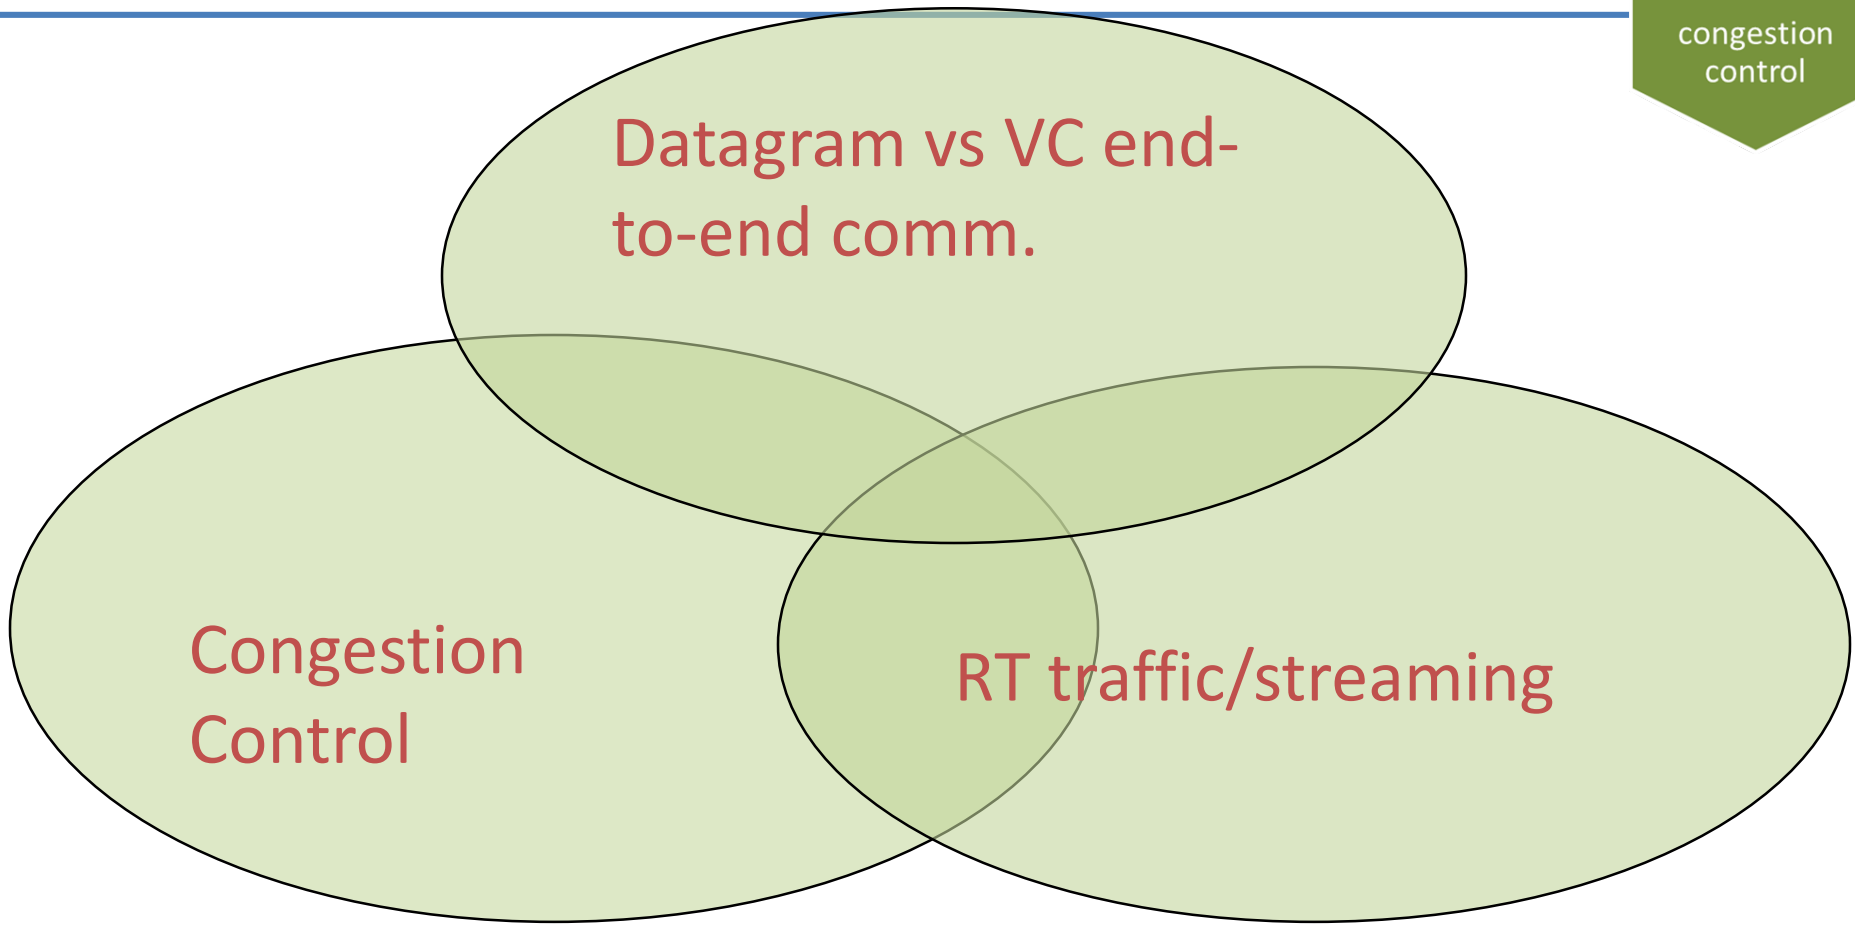
Datagram vs VC end-to-end comm.
Congestion Control
RT traffic/streaming

Datagram vs VC end-to-end communication

- Conceptual differences
- Decisions, comparison

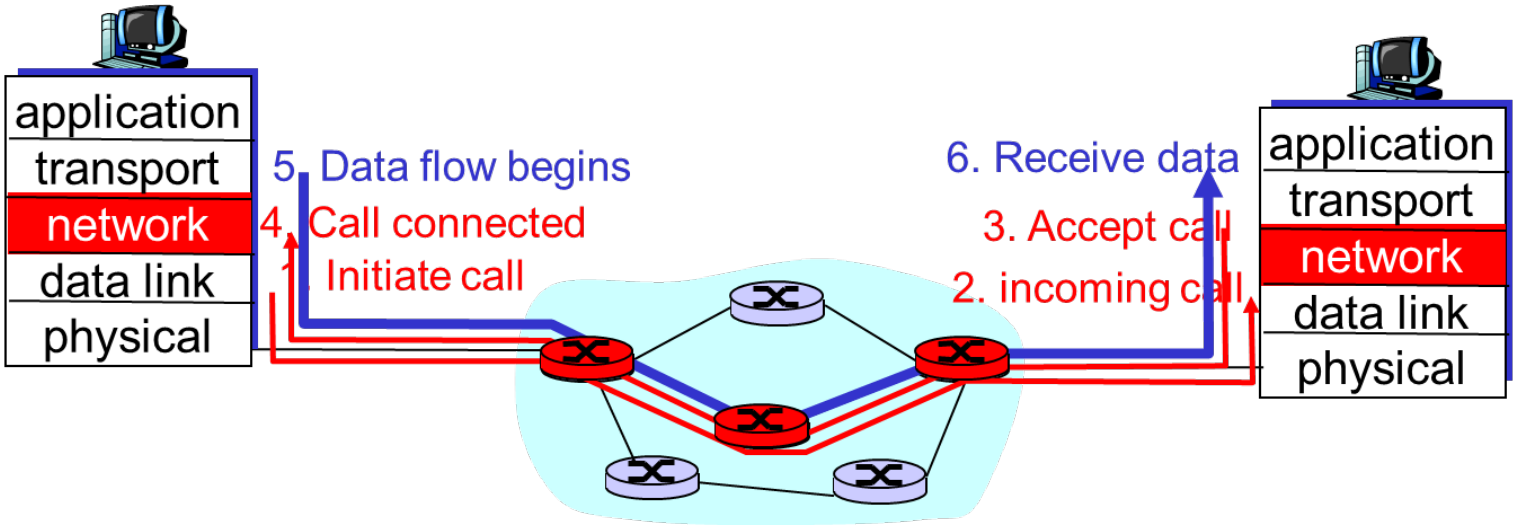

Congestion control (CC)

datagram vs VC

congestion control

- why, how congestion occurs
- CC in TCP and performance; implied weaknesses
- CC in other ways, e.g. VC-based networks
  - Real-time (RT)-traffic resource reservation: traffic shaping and policing
  - rate-based

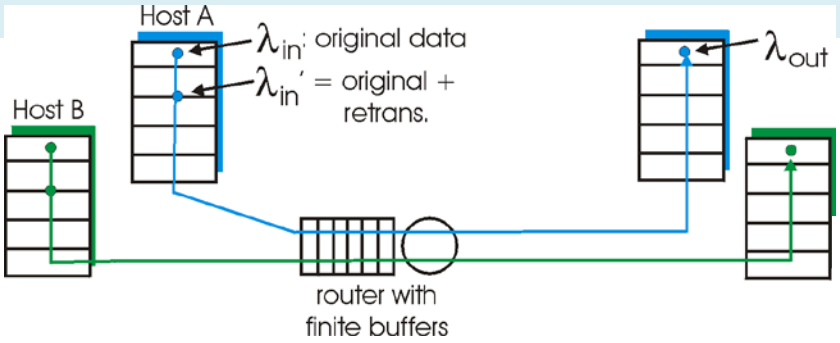

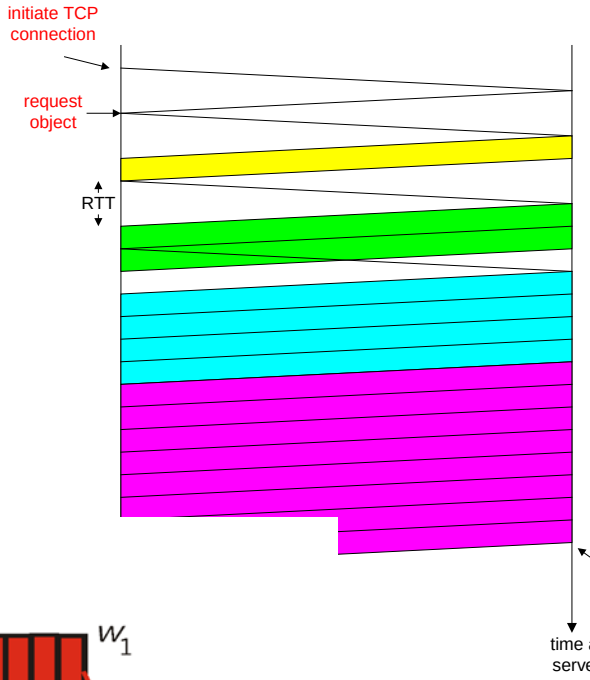

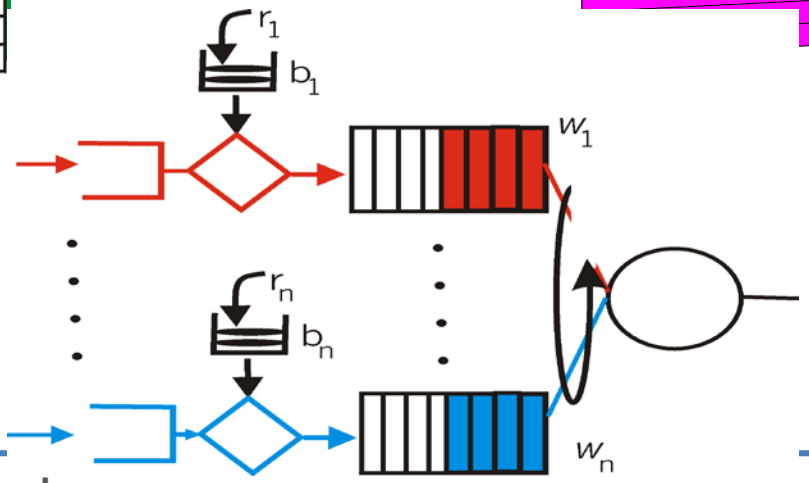

RT/streaming traffic

datagram vs VC

congestion control

Internet context

- Application-level solutions (playout delay, forward-error-control, caching-CDN)
- Intserv, Diffserv, traffic engineering

Conceptual needs:

- packet/flow marking
- Admission control
- Traffic shaping & policing
- Packet scheduling

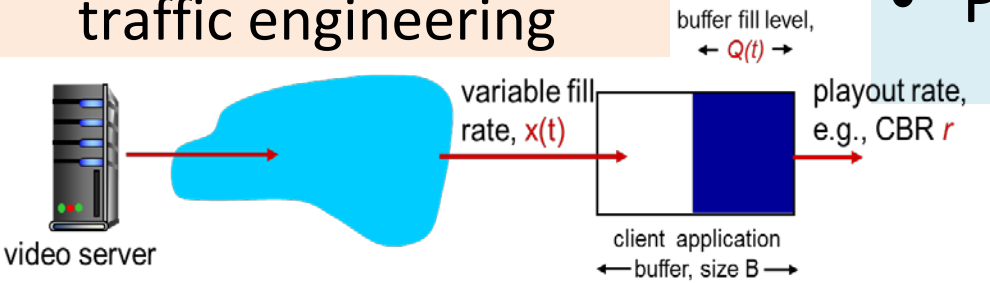

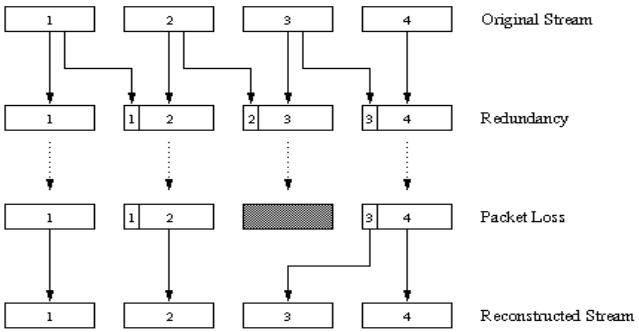

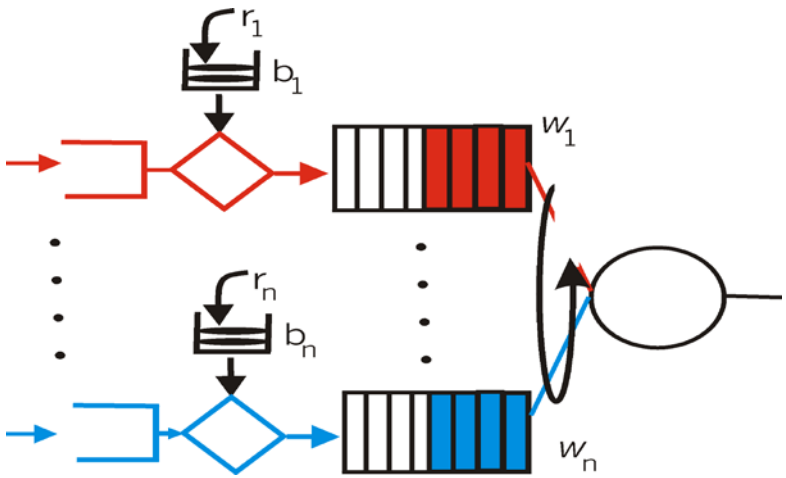

Highlights

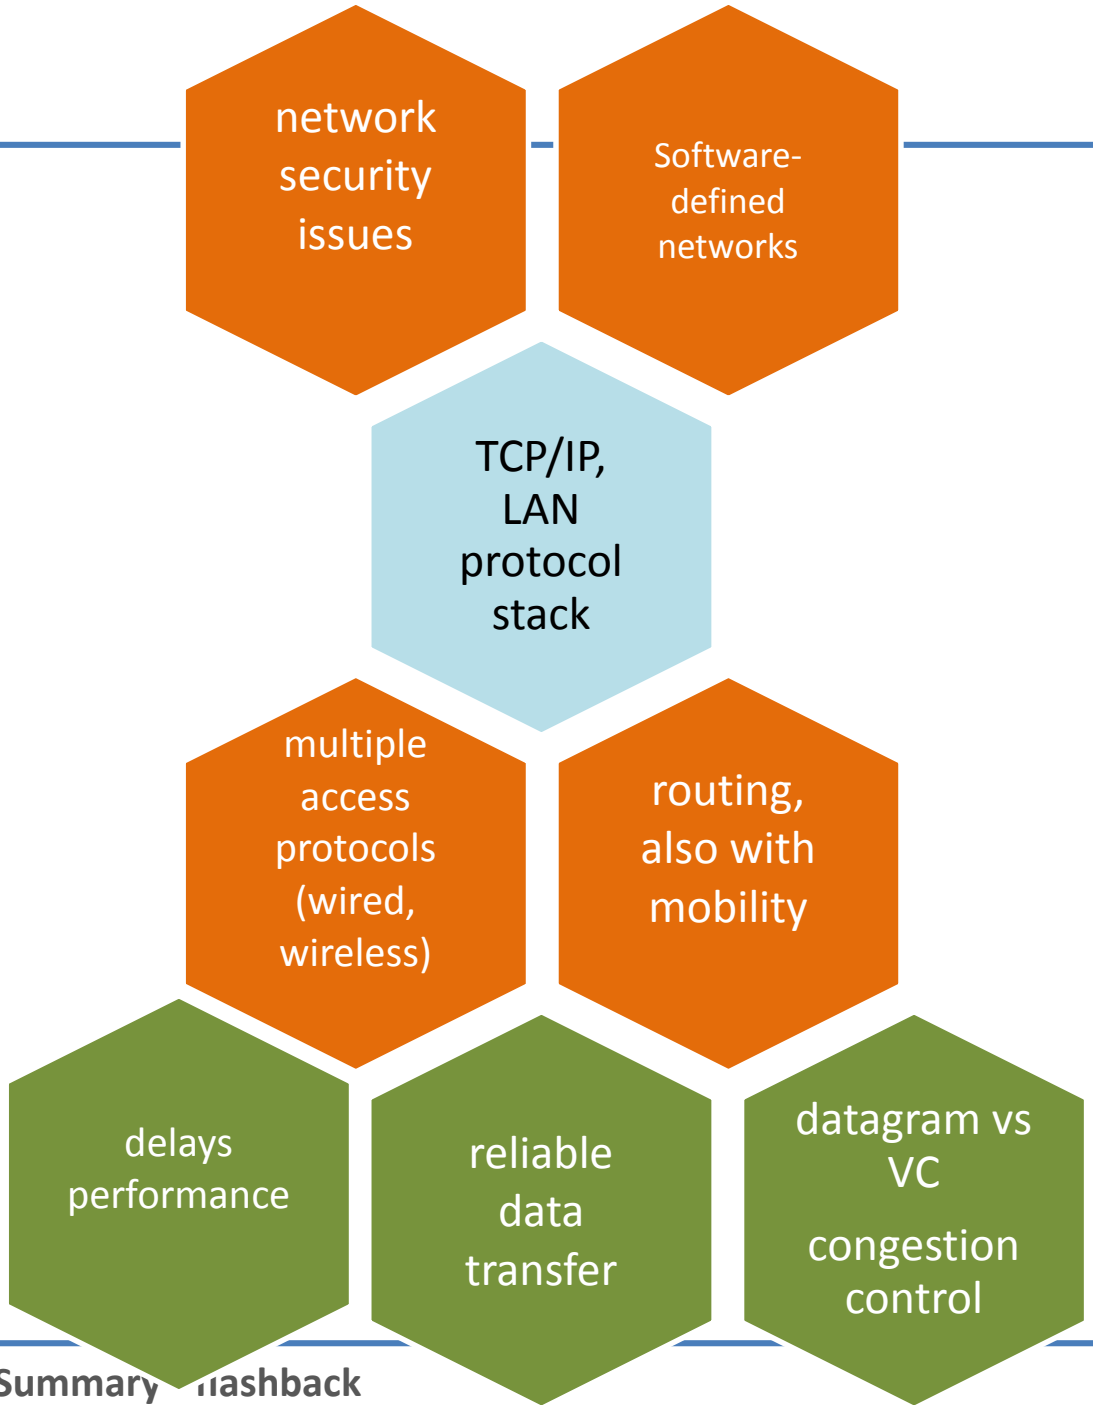
network security issues
Software-defined networks
TCP/IP, LAN protocol stack
multiple access protocols (wired, wireless)
routing, also with mobility
delays performance
reliable data transfer
datagram vs VC
congestion control

Routing, also with mobility

routing, also with mobility

- Routing algorithms, protocols
- Forwarding in routers
- Resource, policy issues
- Addressing mobility, tunneling

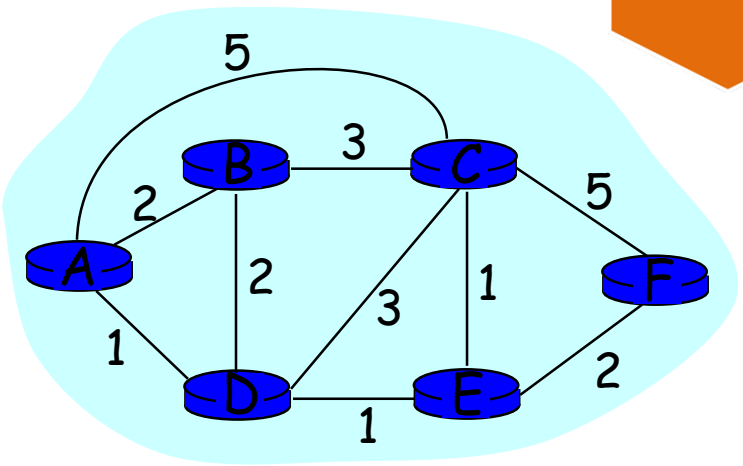

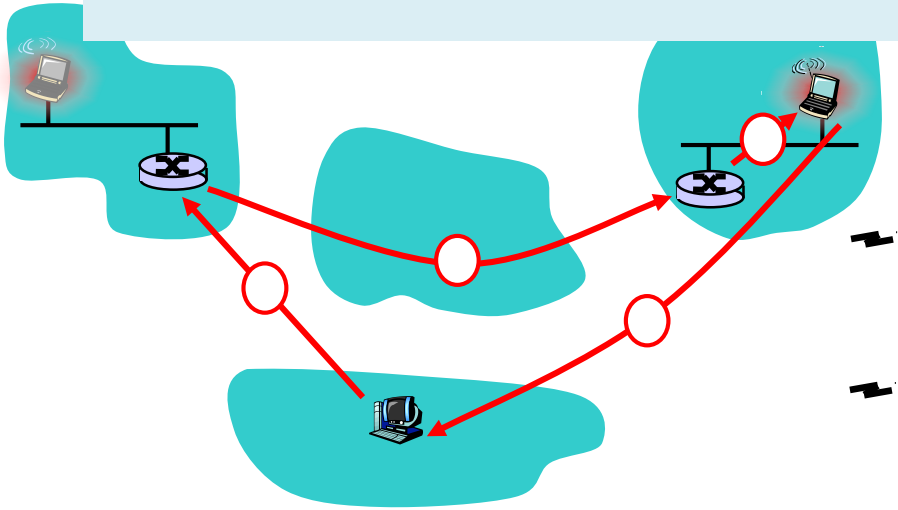

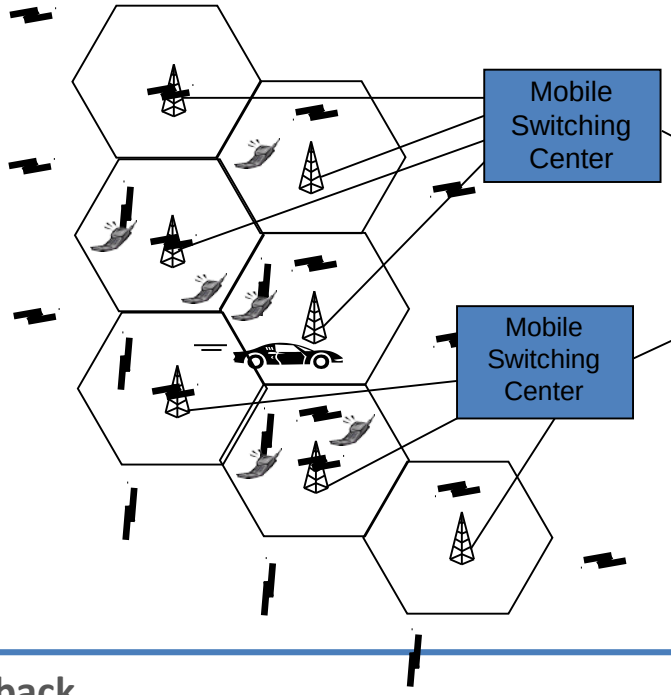

Complementary video links
- IP addresses and subnets
http://www.youtube.com/watch?v=ZTJIkjgyuZE&list=PLE9F3F05C381ED8E8&feature=plcp
- How does BGP choose its routes
http://www.youtube.com/watch?v=RGe0qt9Wz4U&feature=plcp

Medium access: multiple access methods

multiple access protocols (wired, wireless)

Strategies: (functionality, appropriateness)

- **Contention-based (random access), wired/wireless:**
  - Aloha, CSMA(CD/CA)
- **Collision-free:**
  - **Channel partitioning:** TDMA, FDMA, CDMA
  - **Taking turns:** e.g. tokens, reservation-based

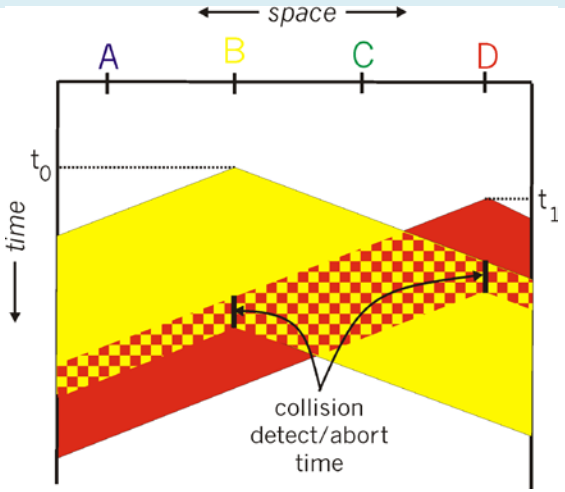

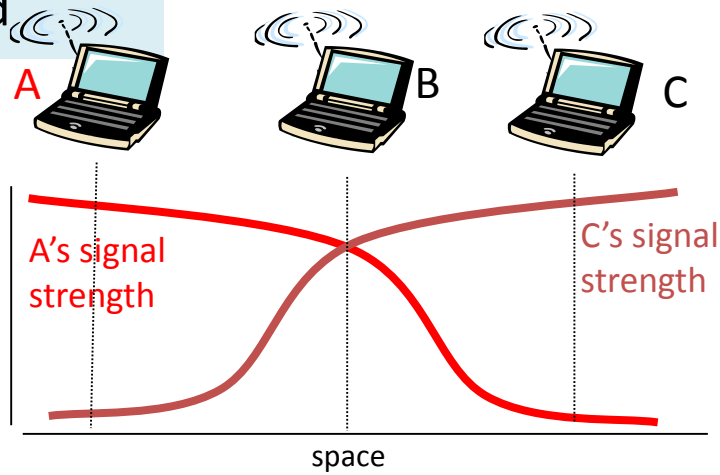

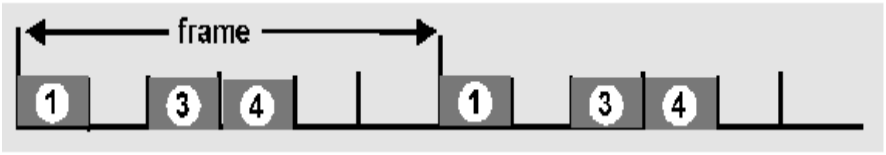

Highlights

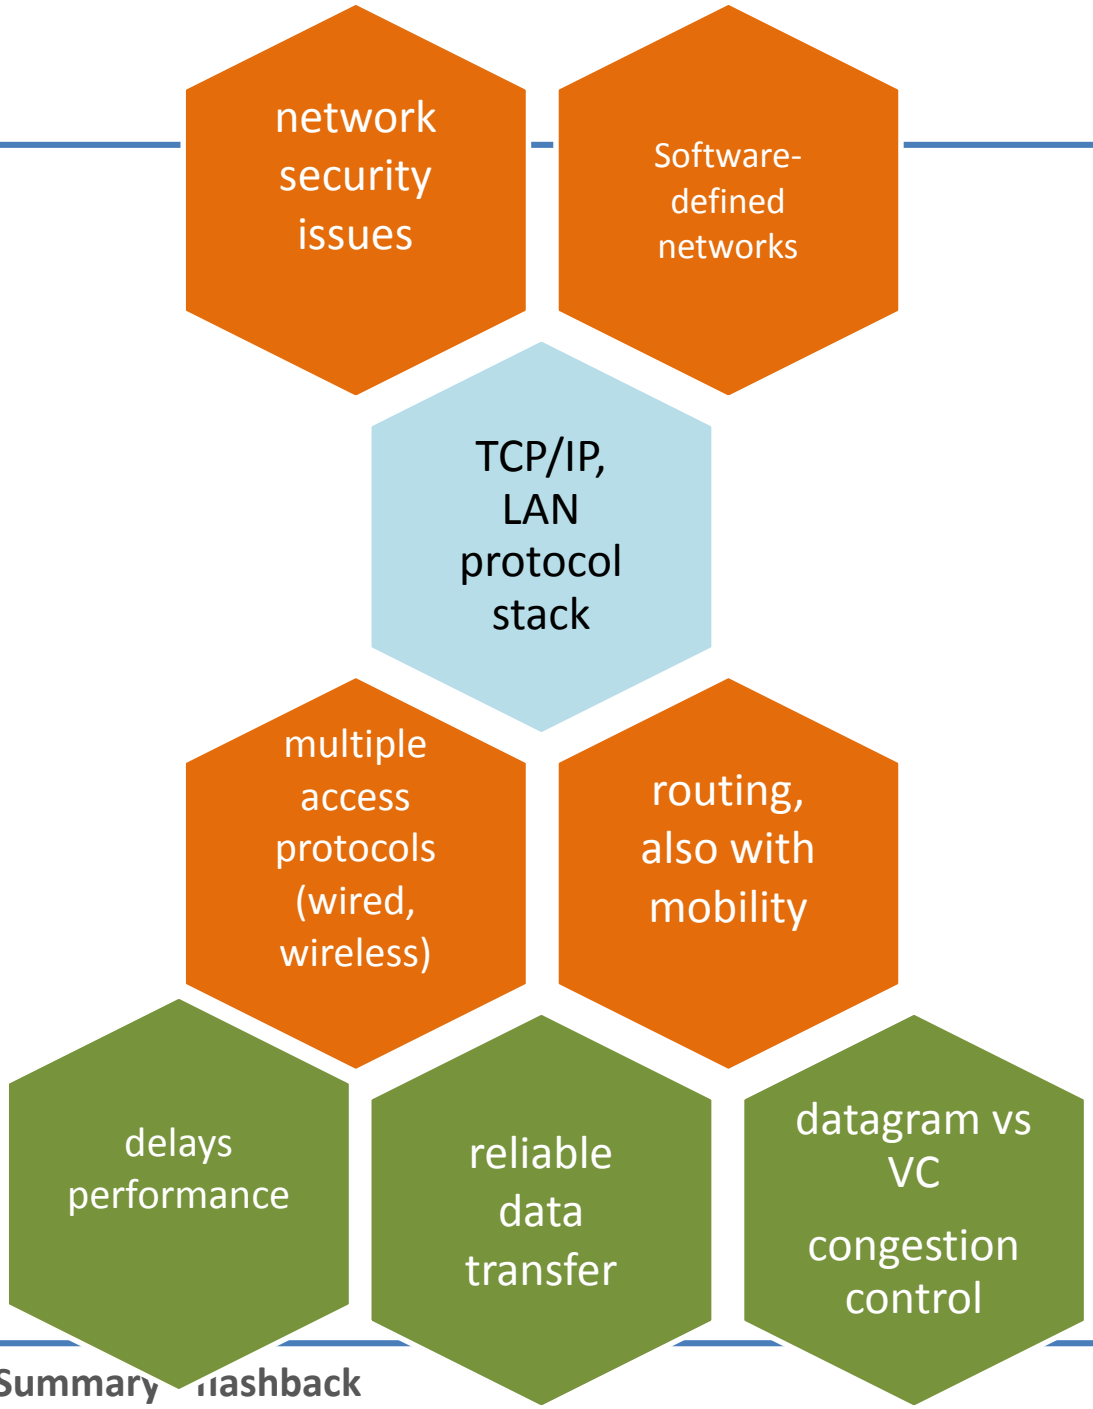
network security issues
Software-defined networks
TCP/IP, LAN protocol stack
multiple access protocols (wired, wireless)
routing, also with mobility
delays performance
reliable data transfer
datagram vs VC
congestion control

TCP/IP protocol stack, applications, evolution

TCP/IP, LAN protocol stack

- Instantiation of network- solutions

(Routing, Congestion Control, Flow & error control, applications, link layer technologies)

- Advantages, limitations, updates
- New types of applications and how they function given the existing state of Internet: multimedia/streaming applications, CDNs, P2P applications, overlays

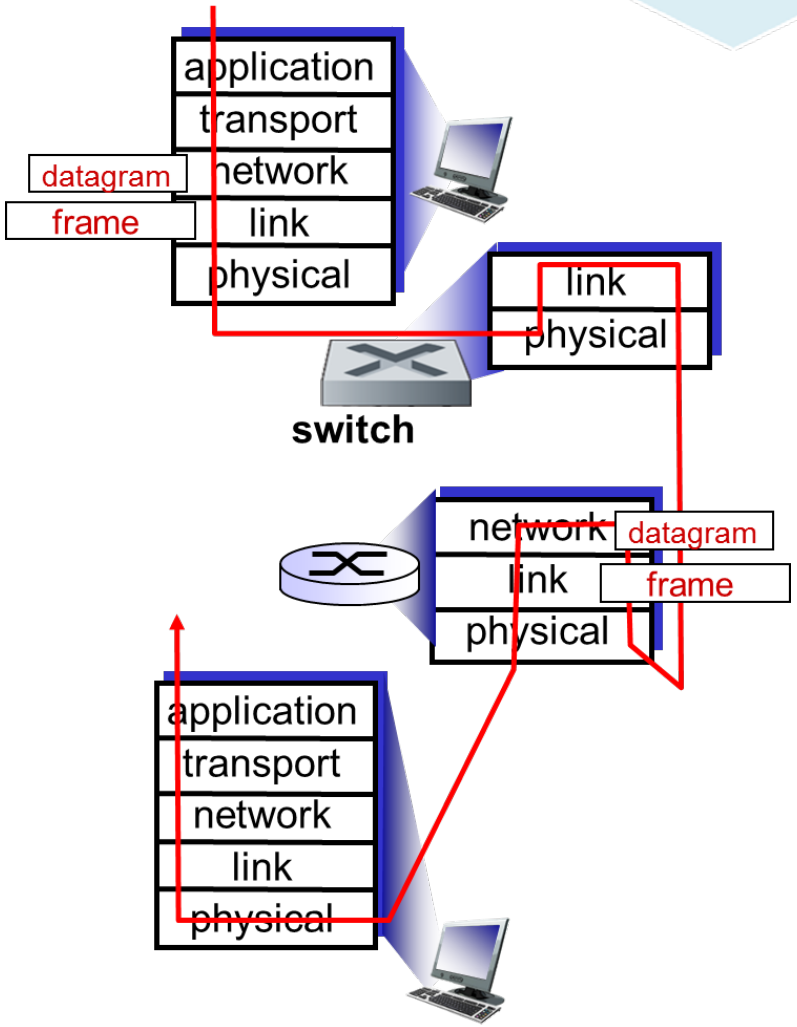

LANs & related link technologies

TCP/IP, LAN protocol stack

- **Protocol Examples: wired, wireless**

  Ethernet, 802.xy, GSM:

  Functionality, performance under low/high load

- **Connecting devices;**
  - functionalities and differences (Hubs, switches)
  - Algorithms for switch-"routing": learning& forwarding of packets
- **ARP**

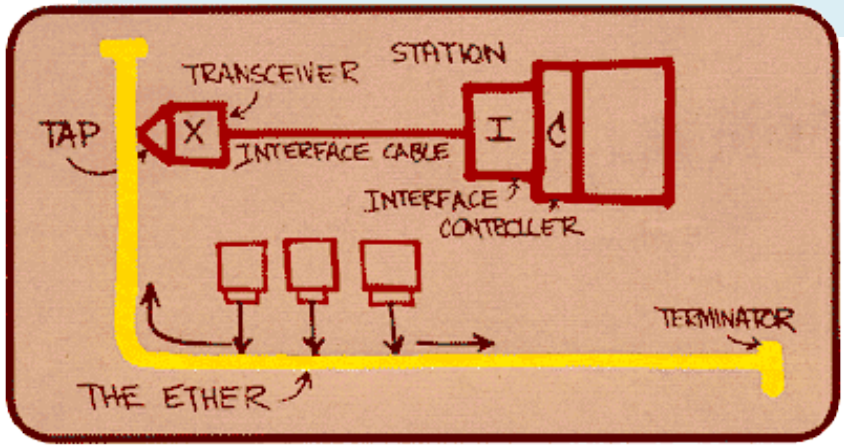

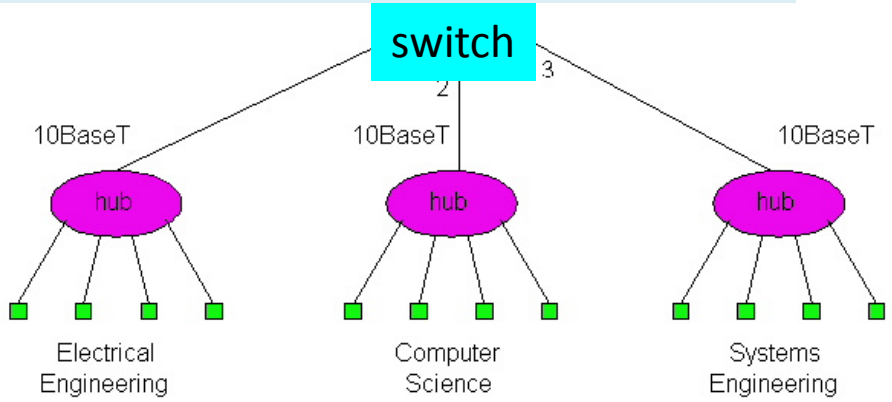

Highlights

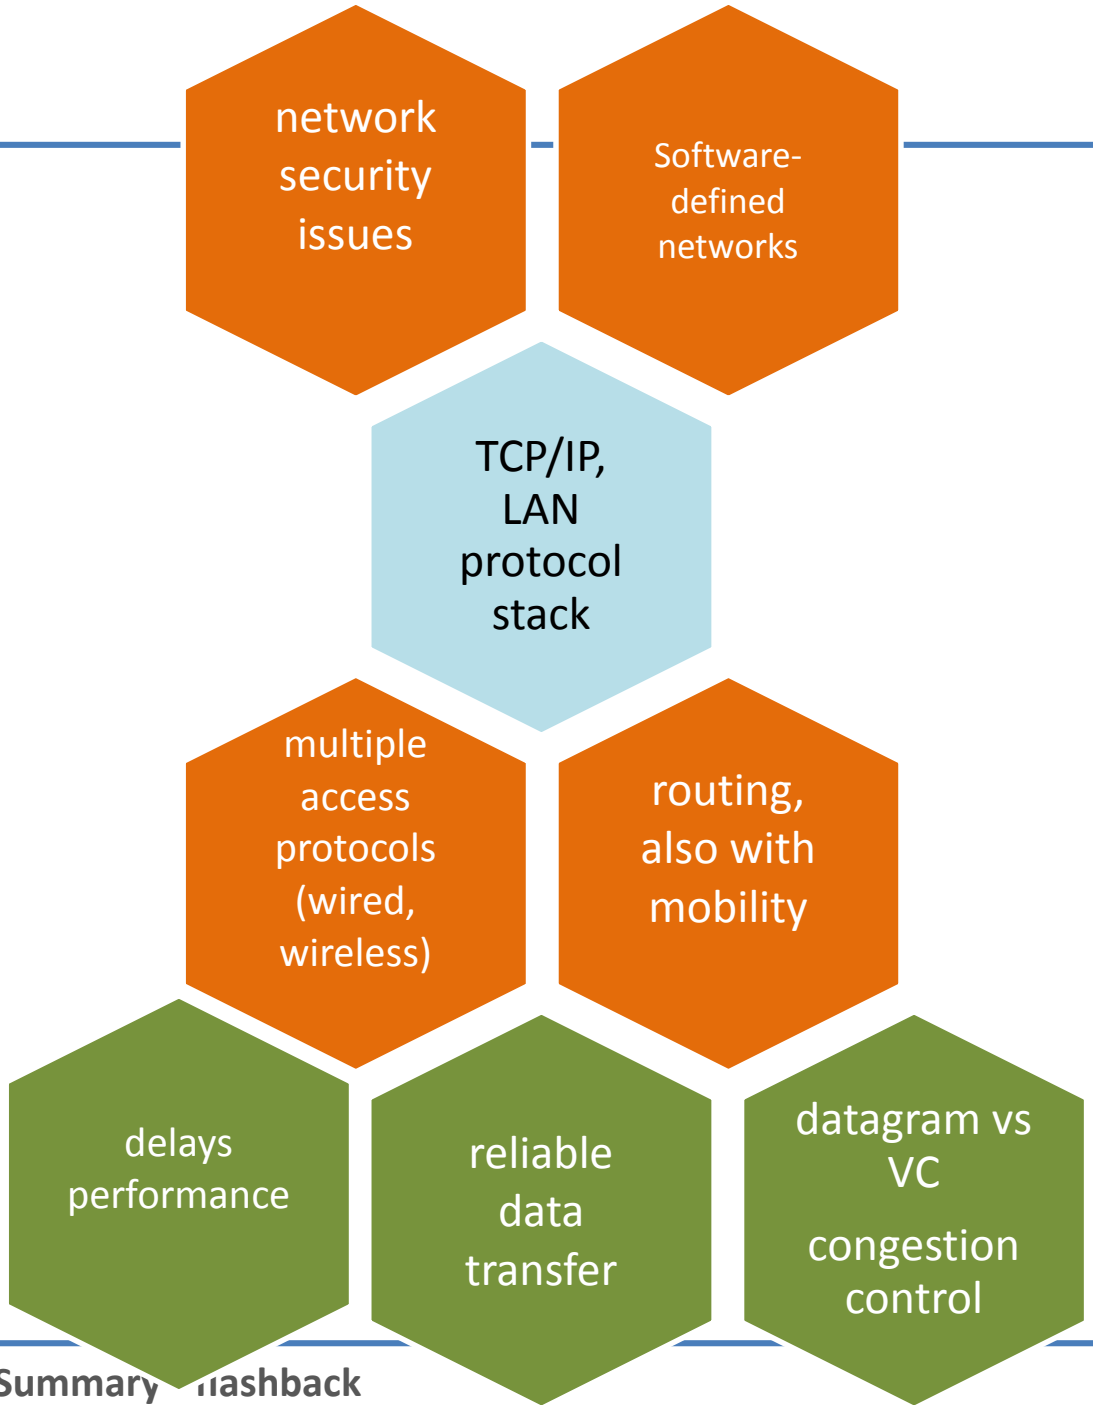
network security issues
Software-defined networks
TCP/IP, LAN protocol stack
multiple access protocols (wired, wireless)
routing, also with mobility
delays performance
reliable data transfer
datagram vs VC
congestion control

Security issues

network security issues

- **C, I, A** and methods to achieve them
  - The language of cryptography
  - Message integrity, signatures
- Instantiation in Internet: SSL, IPsec
- Firewalls

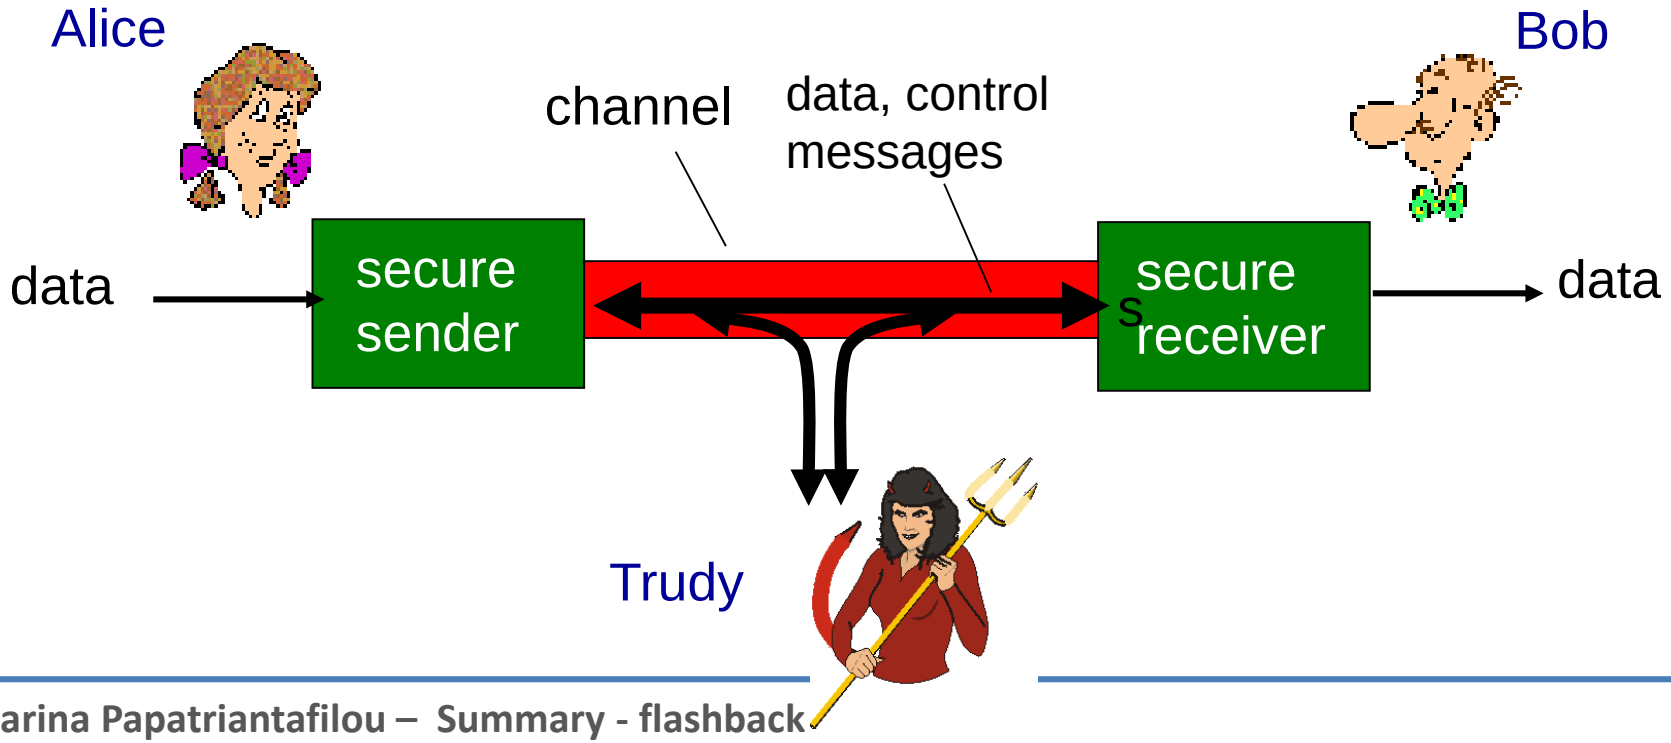

Overlays, software-defined networks

Software-defined networks

- P2P/streaming applications-infrastructure (application-layer networking)
- traffic engineering, tunneling
- **Software-defined networks**: separation of control and execution planes; virtualization of ”layers”: eg. routing table updates implemented elsewhere (not in particular routers)
- 5G; Internet of things concepts

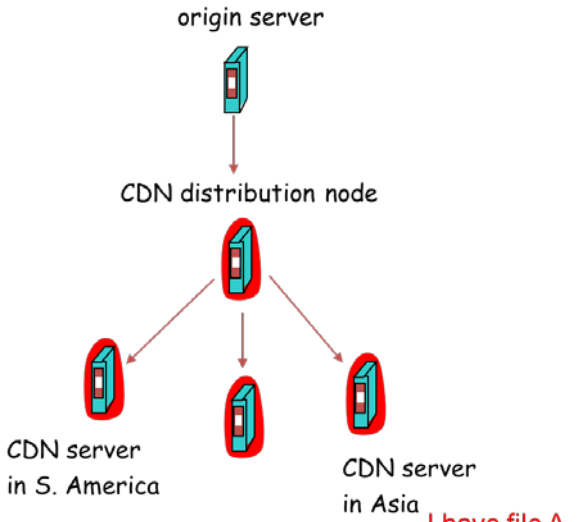

… complement the networking infrastructure …

…taking advantage of the network resources at the edge of the network…

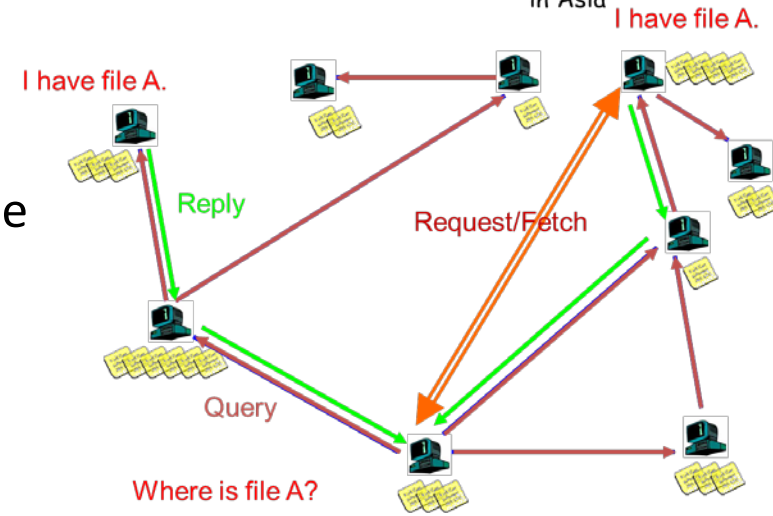

Synthesis: a day in the life of a web request

- putting-it-all-together: synthesis!
  - *goal:* identify, review protocols (at all layers) involved in seemingly simple scenario: requesting www page
  - *scenario:* student attaches laptop to campus network, requests/receives www.google.com

A day in the life …. : scenario

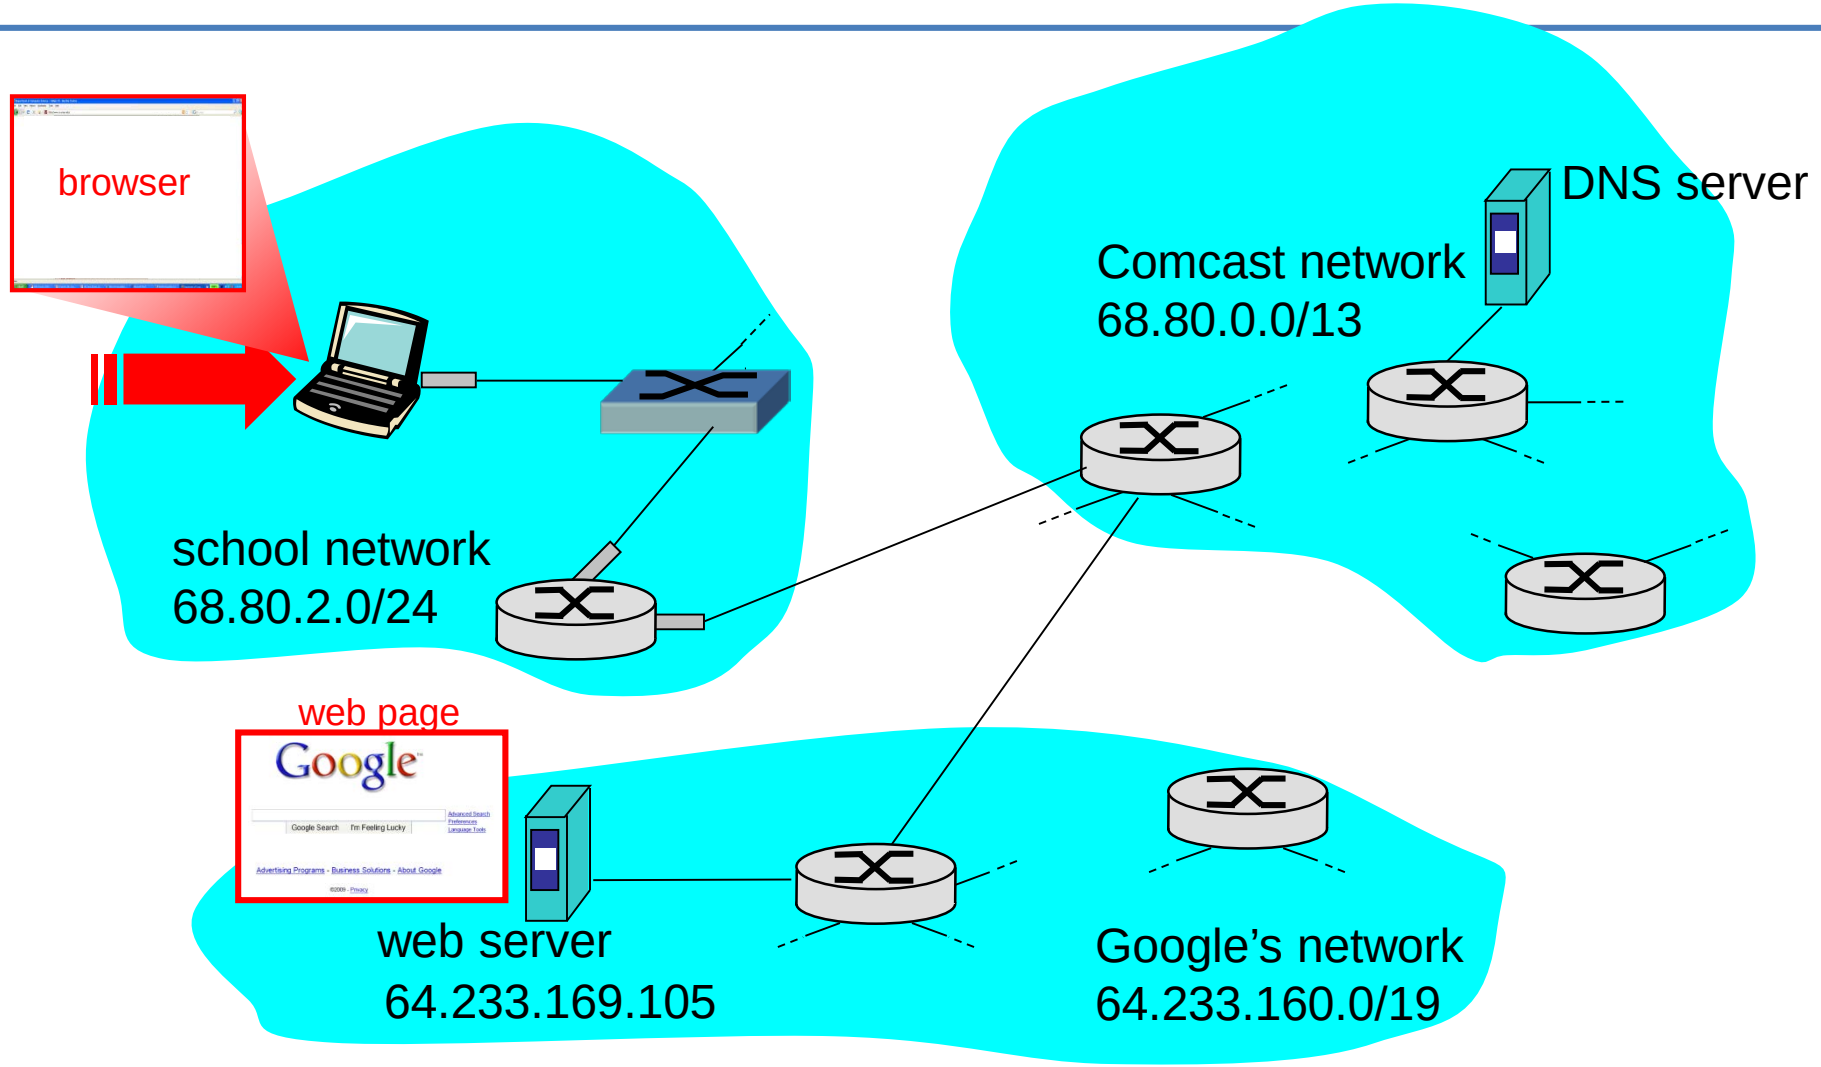

A day in the life… connecting to the Internet

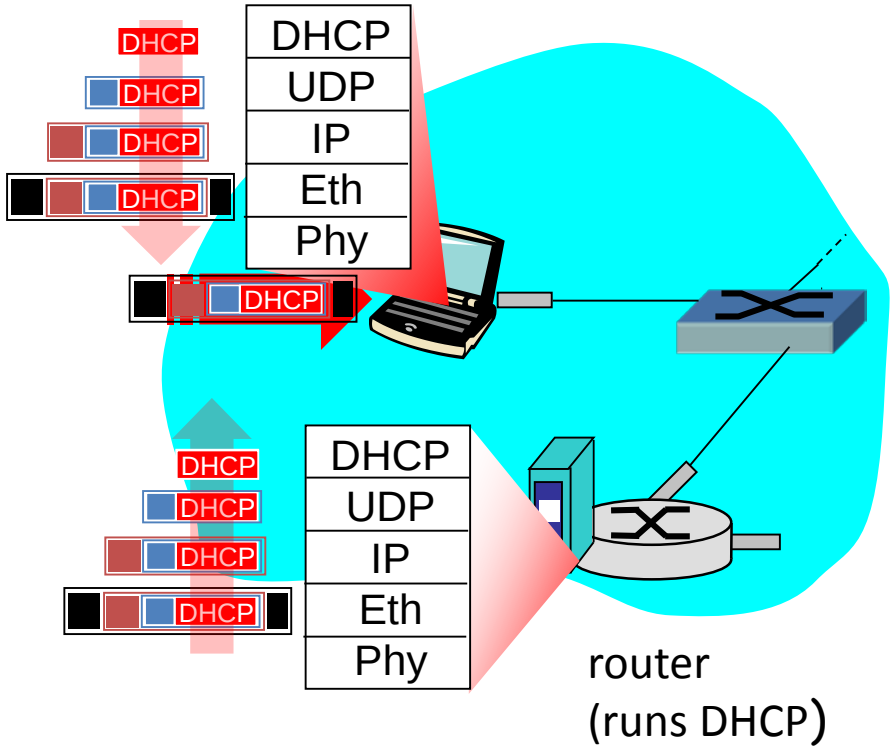

connecting laptop needs to get its own IP address: use ***DHCP***

- DHCP request encapsulated in UDP, encapsulated in IP, encapsulated in Ethernet
- Ethernet frame broadcast (dest: FFFFFFFFFFFF) on LAN, received at router running DHCP server
- Ethernet demux'ed to IP demux'ed to UDP demux'ed to DHCP

A day in the life… connecting to the Internet

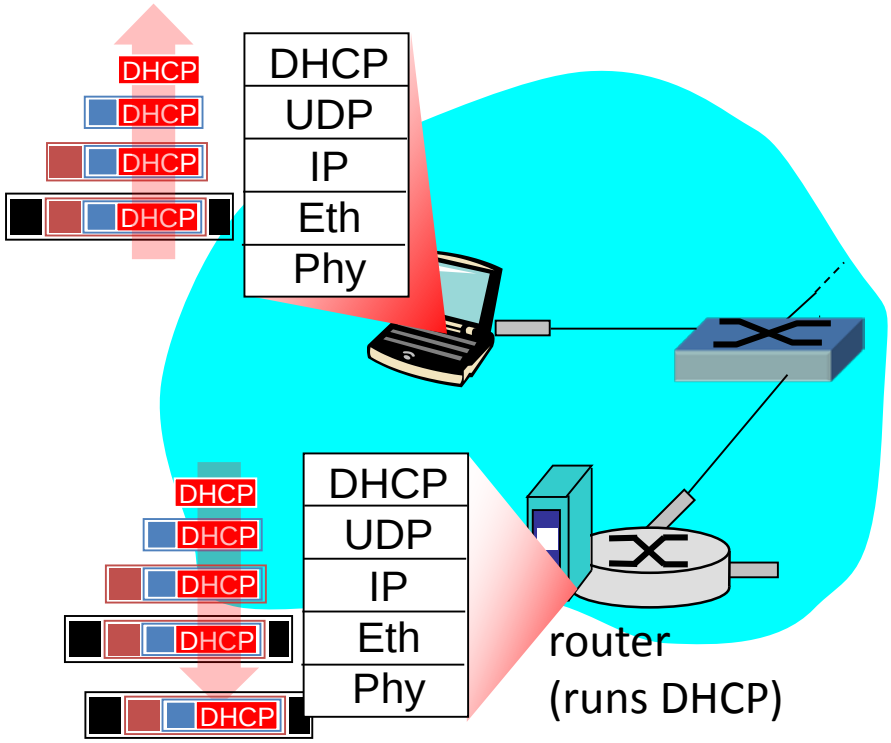

DHCP server formulates ***DHCP ACK*** containing client's IP address (**and also IP address of first-hop router for client, name & IP address of DNS server**)

- frame forwarded (switch learning) through LAN, demultiplexing at client
- DHCP client receives DHCP ACK reply

Client now has IP address, knows name & addr of DNS server, IP address of its first-hop router

A day in the life… ARP (before DNS, before HTTP)

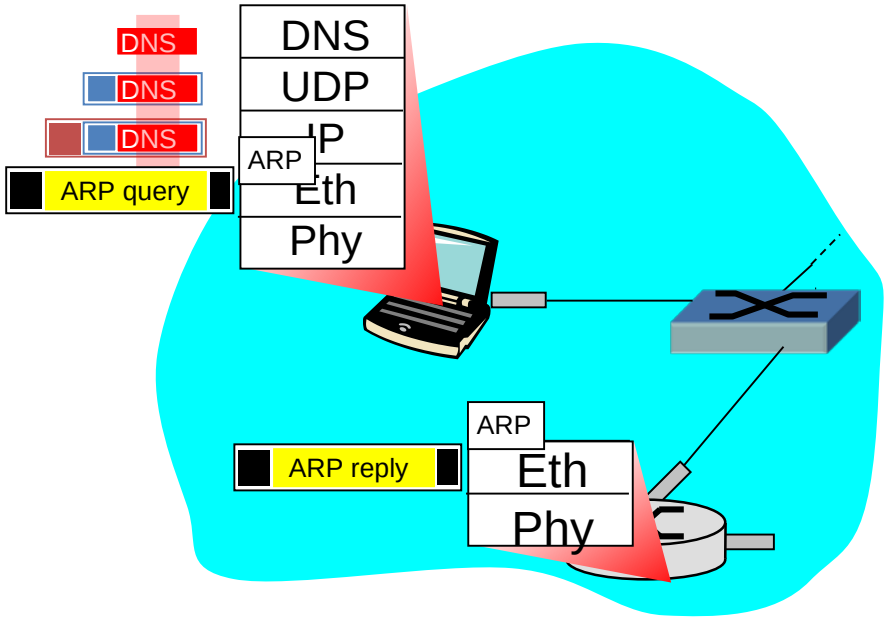

before sending ***HTTP*** request, need IP address of www.google.com: ***DNS***

- DNS query created, encapsulated in UDP, encapsulated in IP, encasulated in Eth. In order to send frame to router, need MAC address of router interface: ARP
- ARP query broadcast, received by router, which replies with ARP reply giving MAC address of router interface
- client now knows MAC address of first hop router, so can now send frame containing DNS query

A day in the life… using DNS

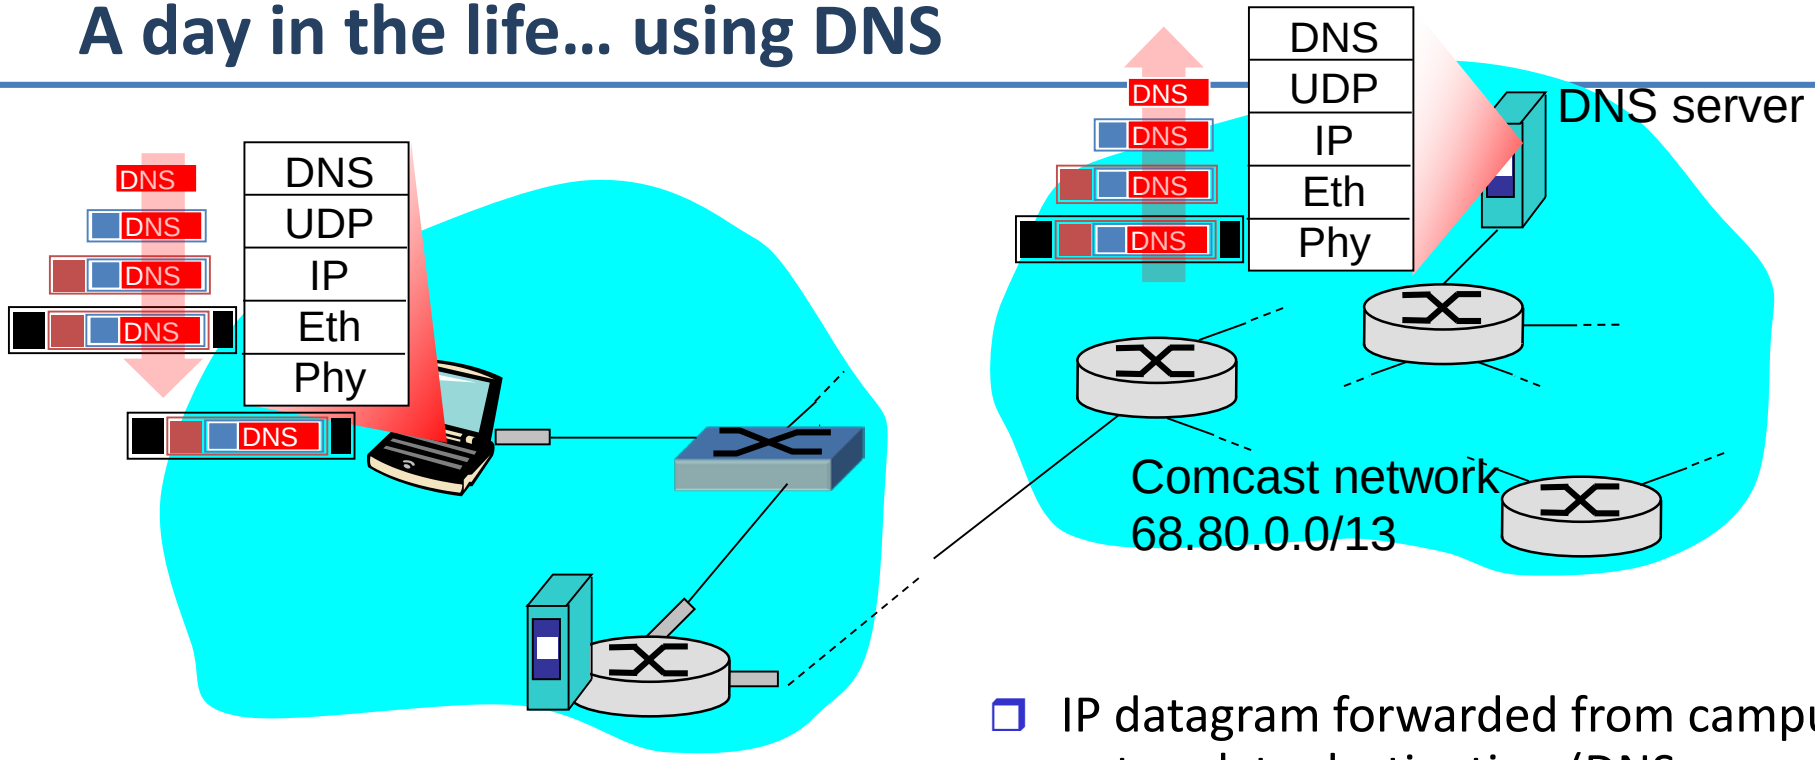

- IP datagram containing DNS query forwarded via LAN switch from client to 1st hop router
- IP datagram forwarded from campus network to destination (DNS-server) network, routed (tables created by RIP, OSPF and BGP routing protocols) to DNS server
- demux'ed to DNS server
- DNS server replies to client with IP address of www.google.com

A day in the life… TCP connection carrying HTTP

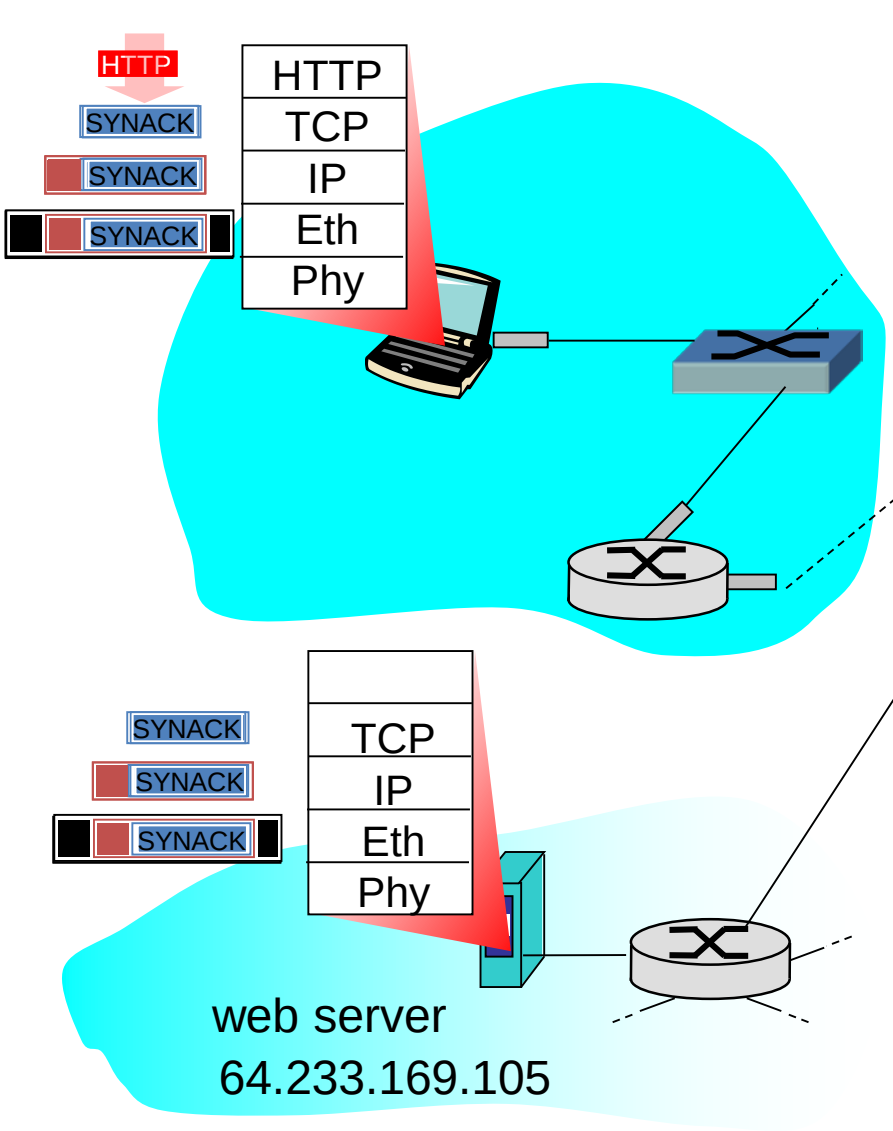

- to send HTTP request, client first opens TCP socket to web server
- TCP SYN segment (step 1 in 3-way handshake) inter-domain routed to web server
- web server responds with TCP SYNACK
- TCP connection established!

A day in the life… HTTP request/reply

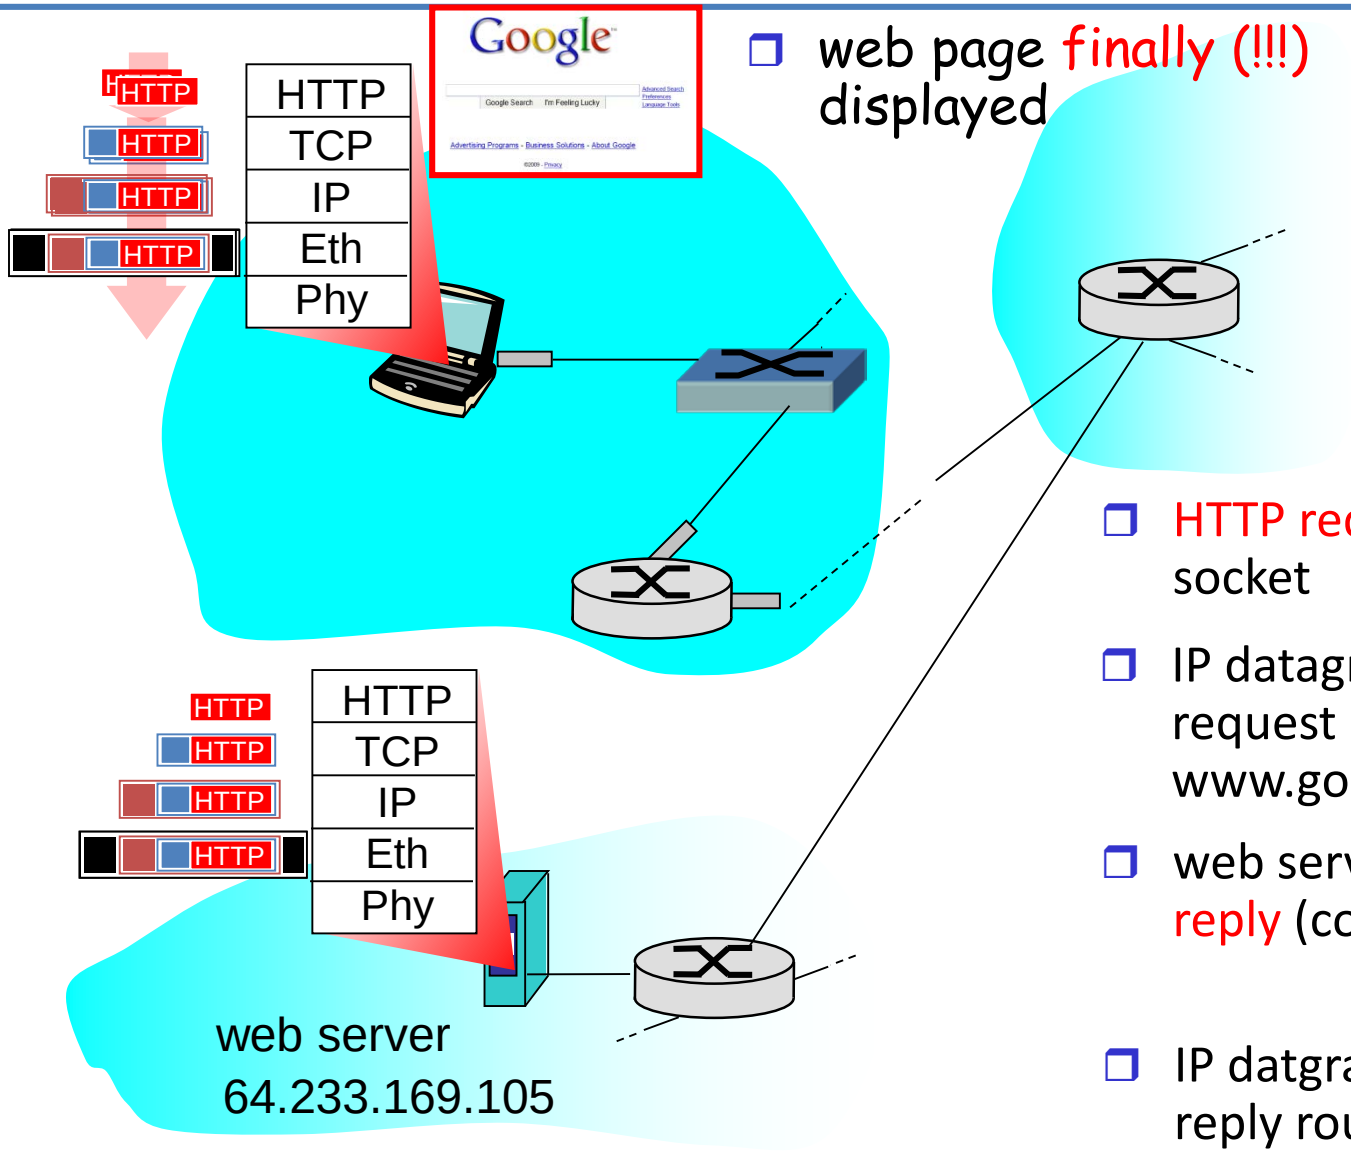

- web page finally (!!!) displayed

- HTTP request sent into TCP socket
- IP datagram containing HTTP request routed to www.google.com
- web server responds with HTTP reply (containing web page)
- IP datgram containing HTTP reply routed back to client

Synthesis cont.

1. Reflections, prespectives
2. Networking constantly evolving

The Internet: virtualizing networks

1974: multiple unconnected nets
- ARPAnet
- data-over-cable networks
- packet satellite network (Aloha)
- packet radio network

… differing in:
- addressing conventions
- packet formats
- error recovery
- routing

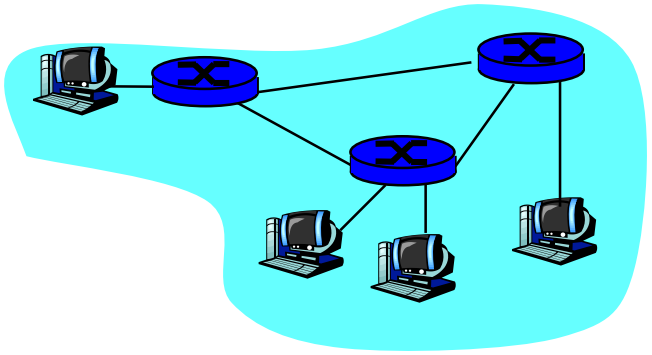

ARPAnet

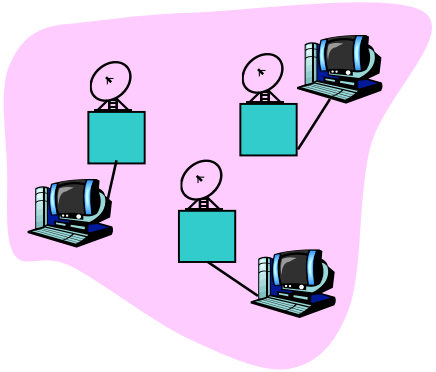

satellite net

"A Protocol for Packet Network Intercommunication", V. Cerf, R. Kahn, IEEE Transactions on Communications, May, 1974, pp. 637-648.

The Internet: virtualizing networks

Internetwork layer (IP):

- addressing: internetwork appears as single, uniform entity, despite underlying local network heterogeneity
- network of networks

Gateway:

- "embed internetwork packets in local packet format"
- route (at internetwork level) to next gateway

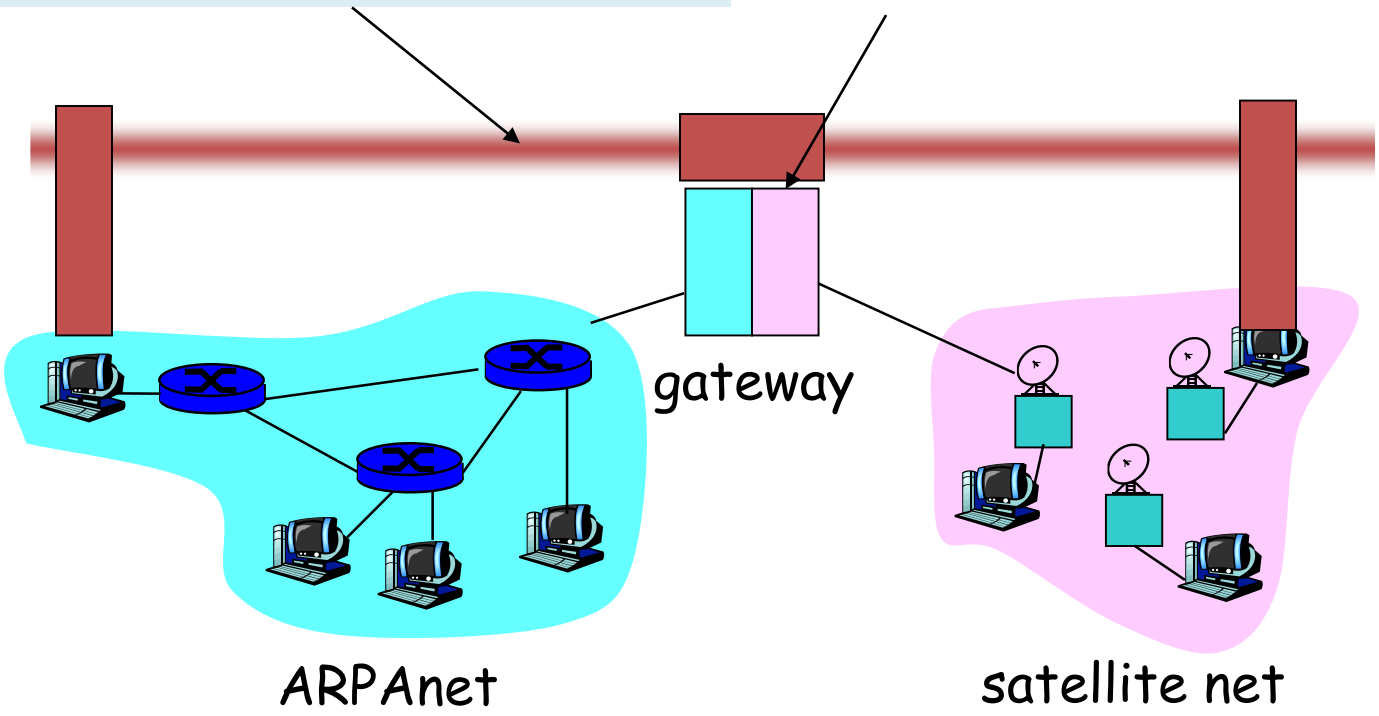

Internet structure: network of networks

Question: given *millions* of access ISPs, how to connect them together?

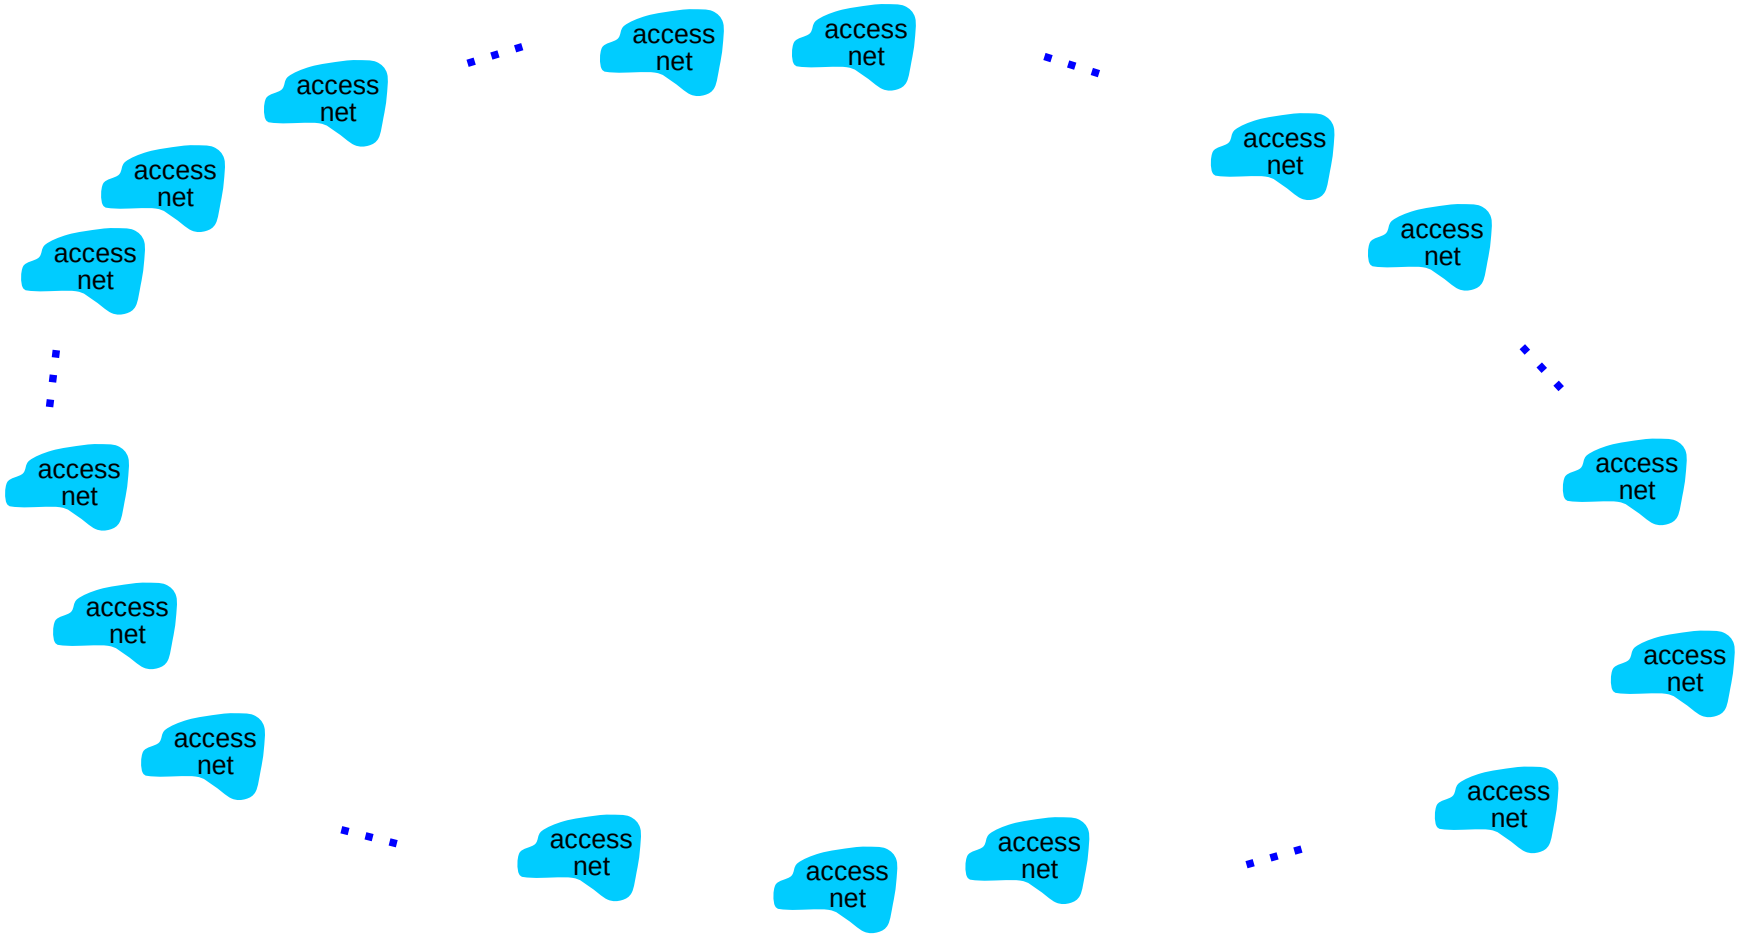

Internet structure: network of networks

Option: connect each access ISP to every other access ISP?

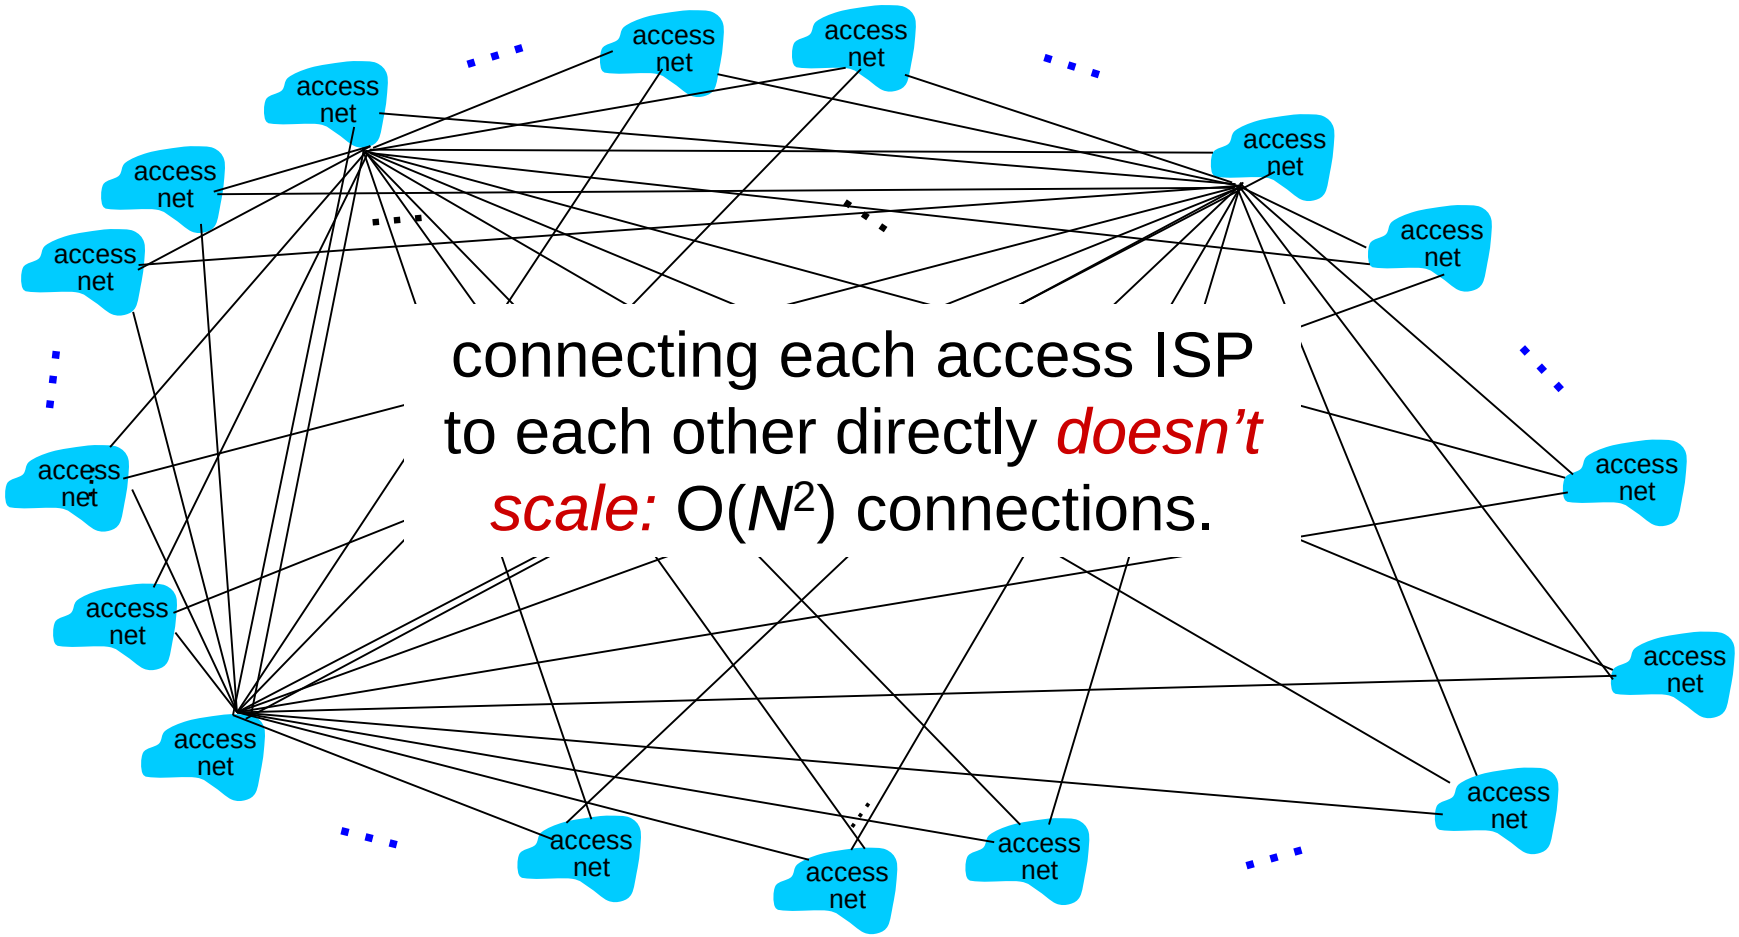

Internet structure: network of networks

Option: connect each access ISP to a global transit (imaginary) ISP?
Customer and provider ISPs have economic agreement.

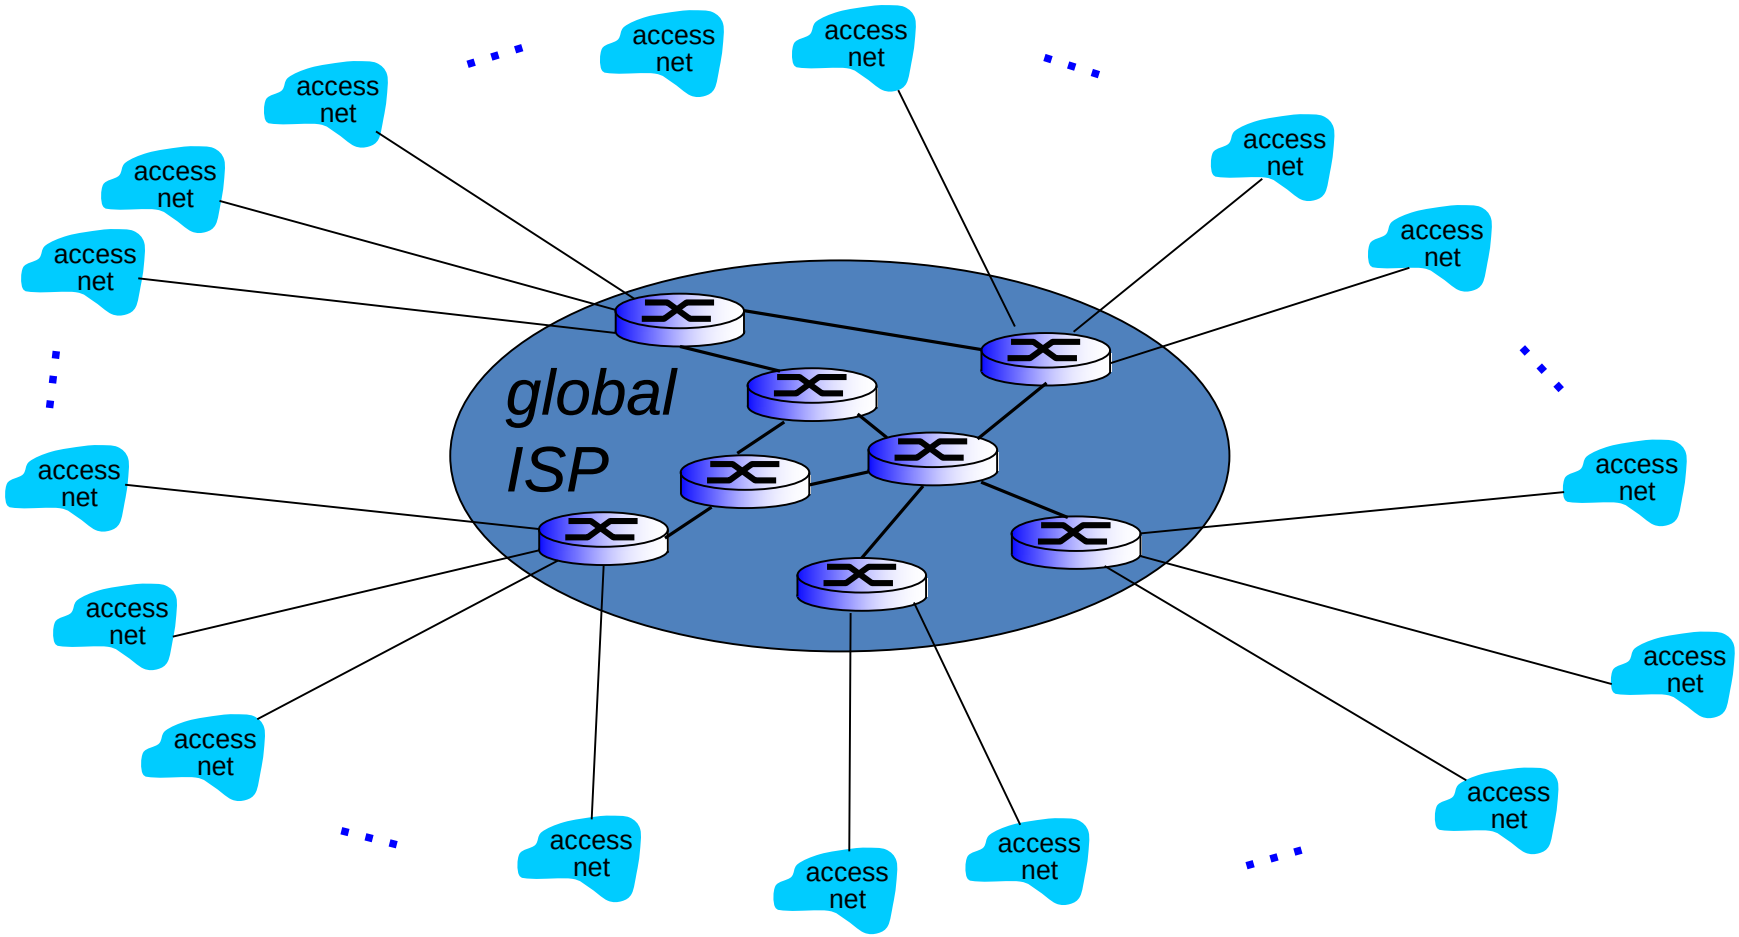

Internet structure: network of networks

But if one global ISP is viable business, there will be competitors ….

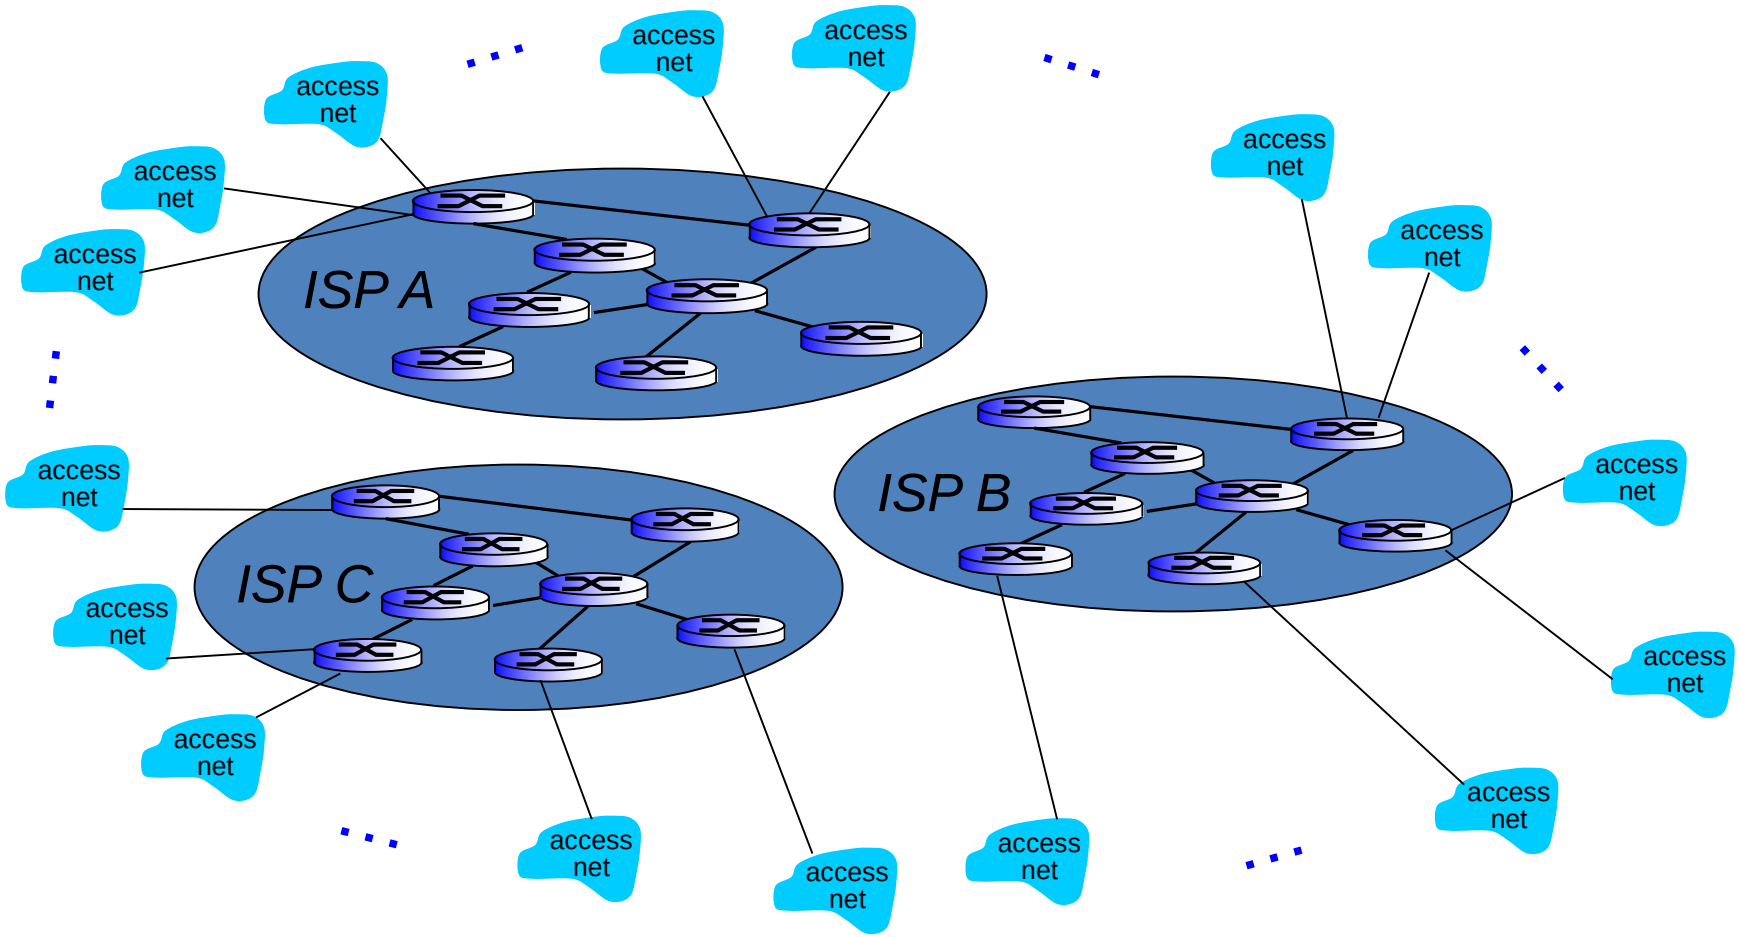

Internet structure: network of networks

But if one global ISP is viable business, there will be competitors …. which must be interconnected

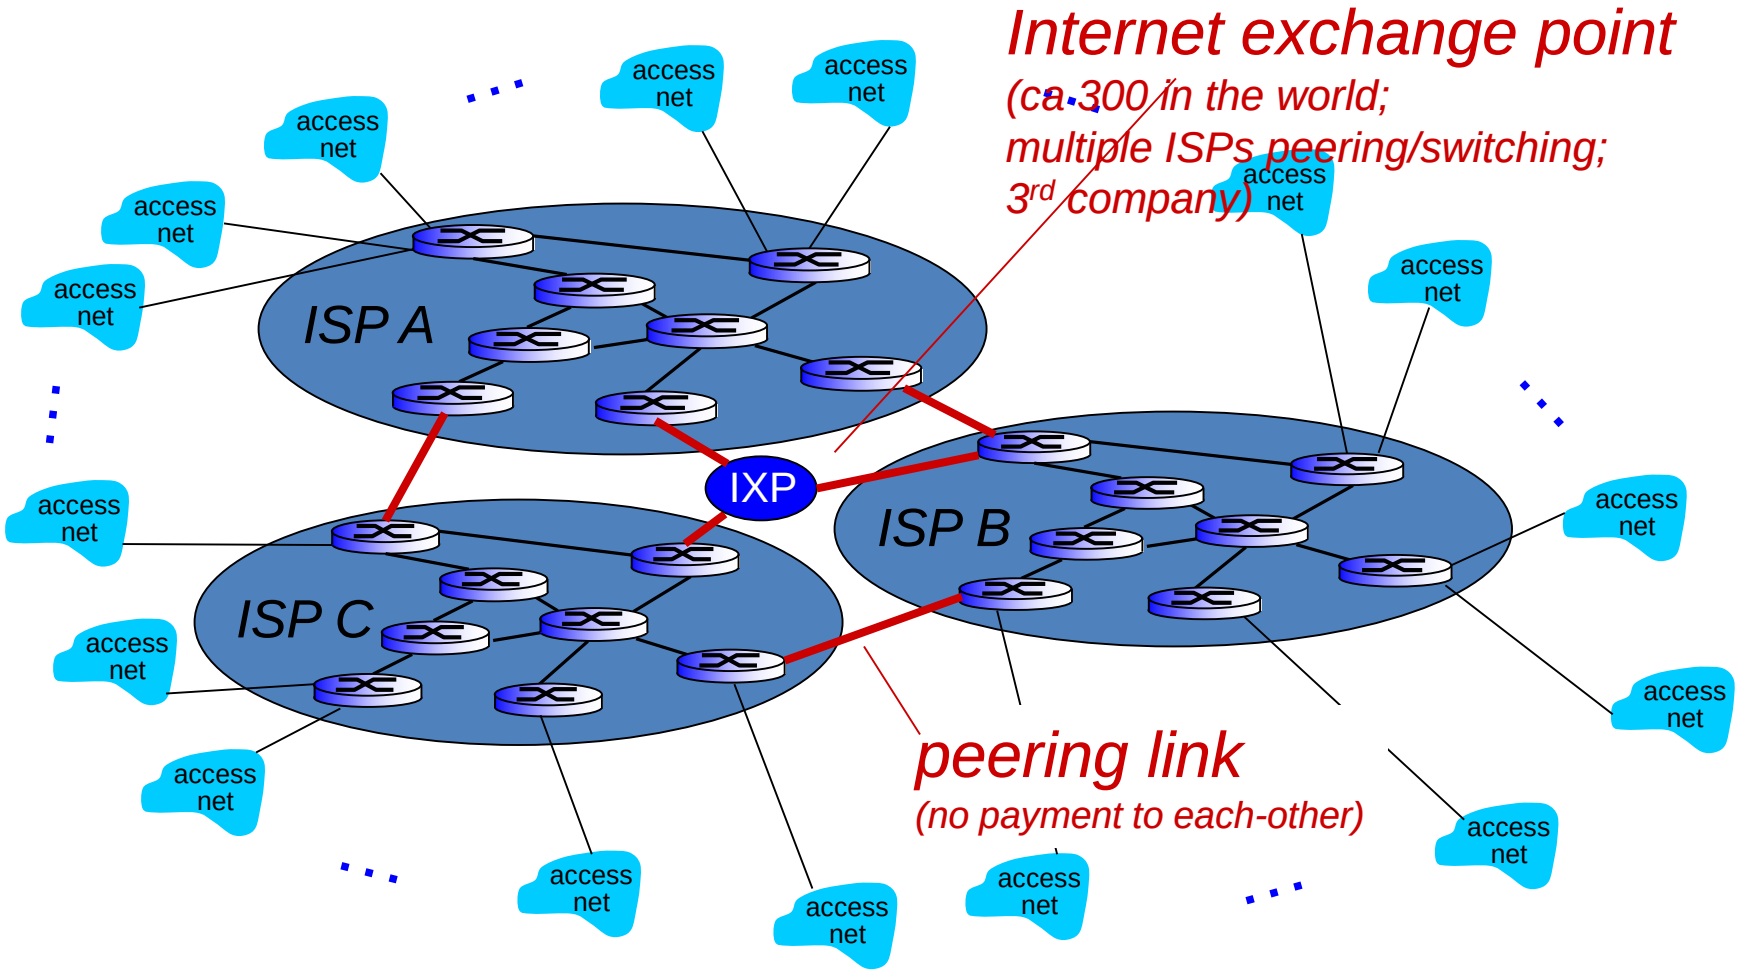

Internet structure: network of networks

… and regional networks may arise to connect access nets to ISPS

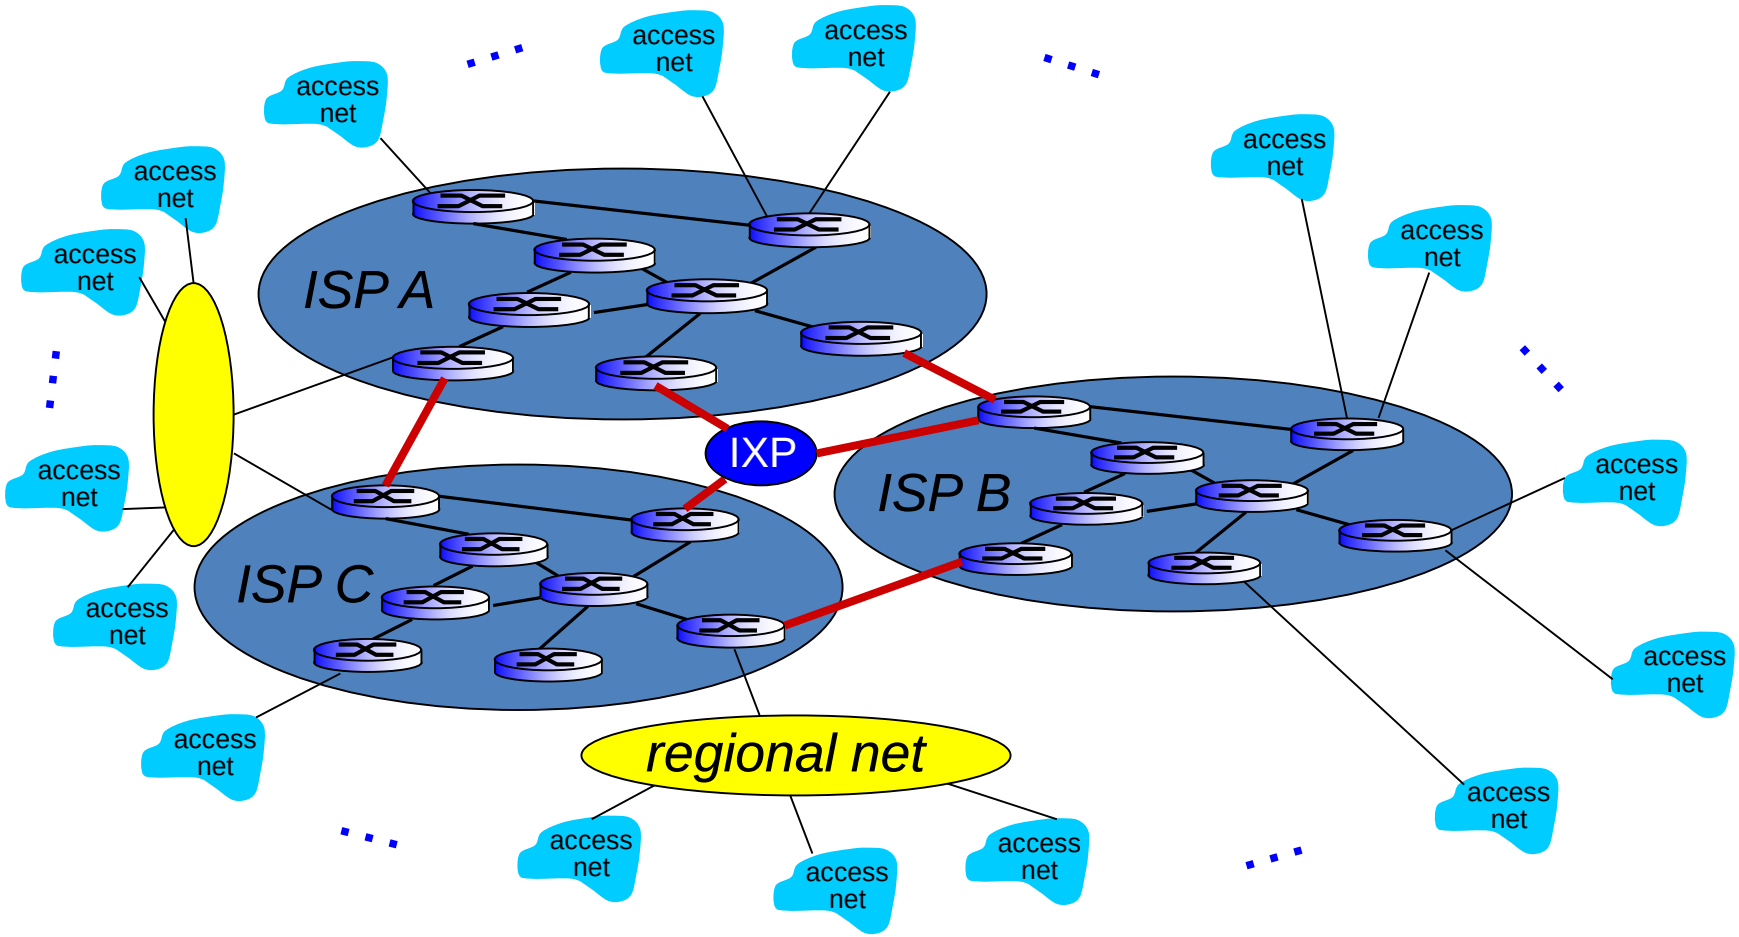

Internet structure: network of networks

… and content provider networks (e.g., Google, Microsoft, Akamai) may run their own network, to bring services, content close to end users

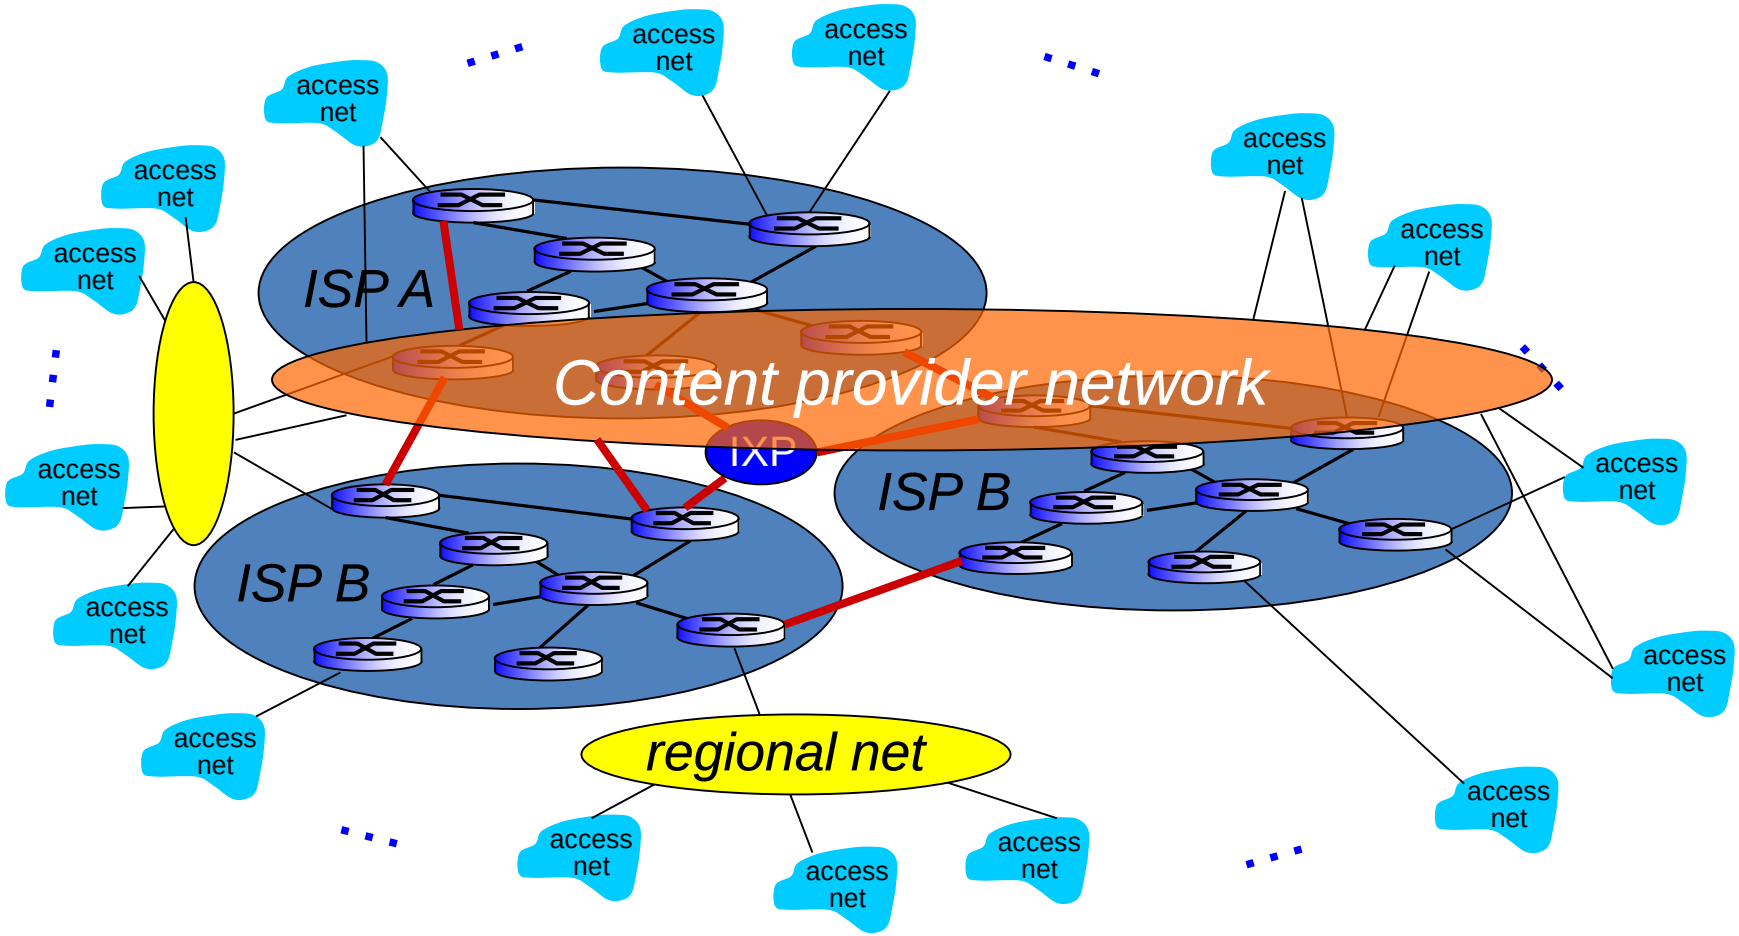

Internet structure: network of networks

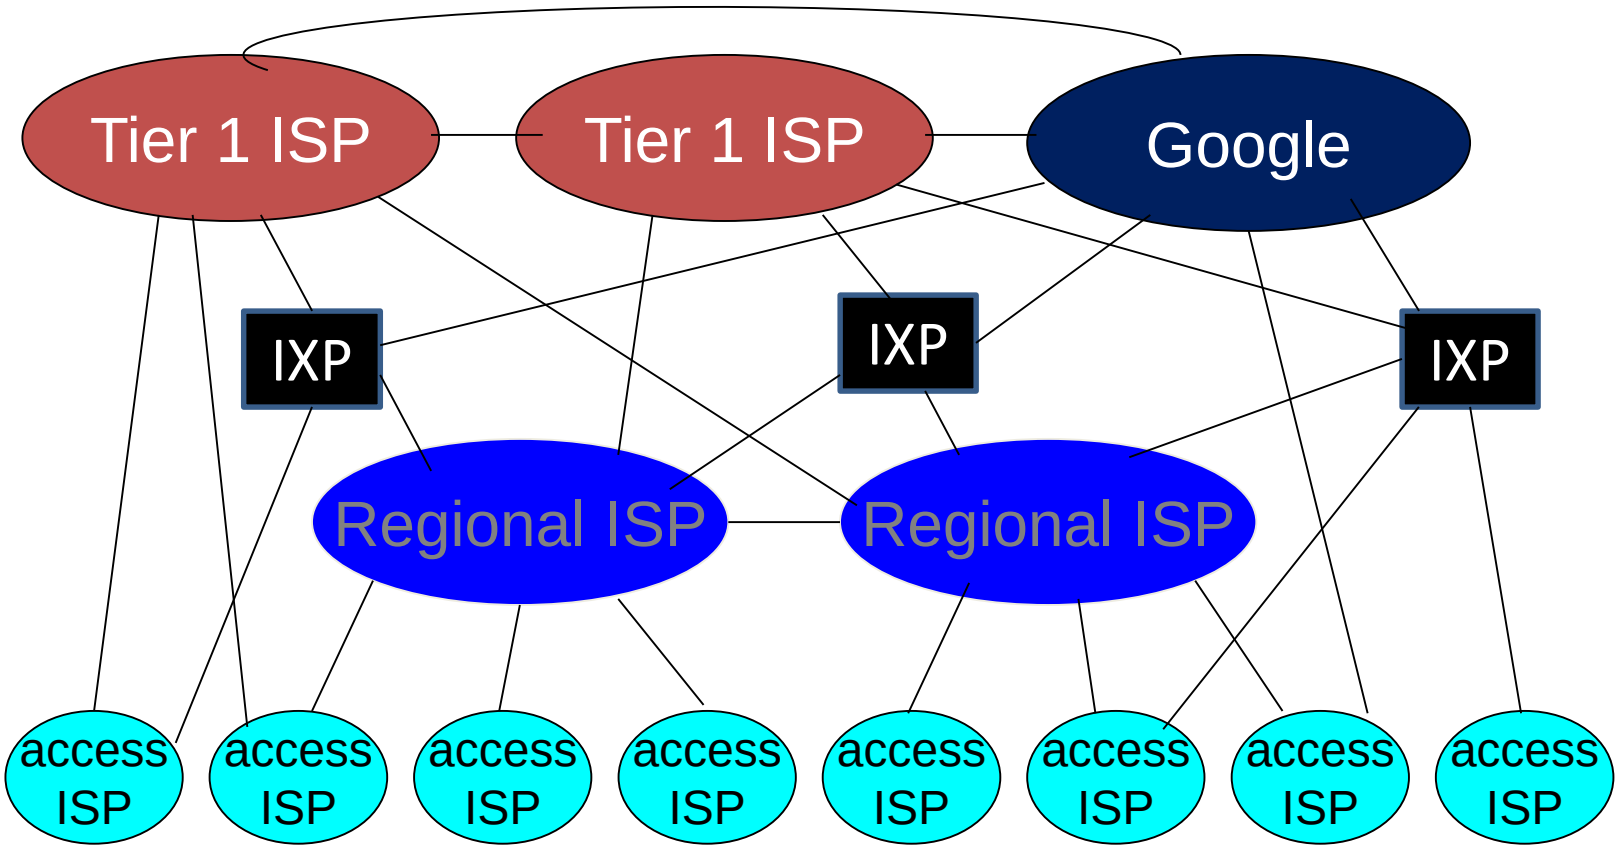

- at center: small # of well-connected large networks
  - "tier-1" commercial ISPs (e.g., AT&T, NTT, TeliaSonera, DeutcheTelecom), national & international coverage
  - A new form of content provider network (e.g, Google): private network that connects it data centers to Internet, often bypassing tier-1, regional ISPs

Synthesis cont.

1. Reflections, prespectives
2. Networking constantly evolving

Data center networks

- 10's to 100's of thousands of hosts, often closely coupled, in close proximity:
  - e-business (e.g. Amazon)
  - content-servers (e.g., YouTube, Akamai, Apple, Microsoft)
  - search engines, data mining (e.g., Google)

- challenges:
  - multiple applications, each serving massive numbers of clients
  - managing/balancing load, networking, data bottlenecks

Inside a 40-ft Microsoft container, Chicago data center

Data center networks

load balancer: application-layer routing

- receives external client requests
- directs workload within data center
- returns results to external client (hiding data cente internals from client)
- Distributed systems & networks working together

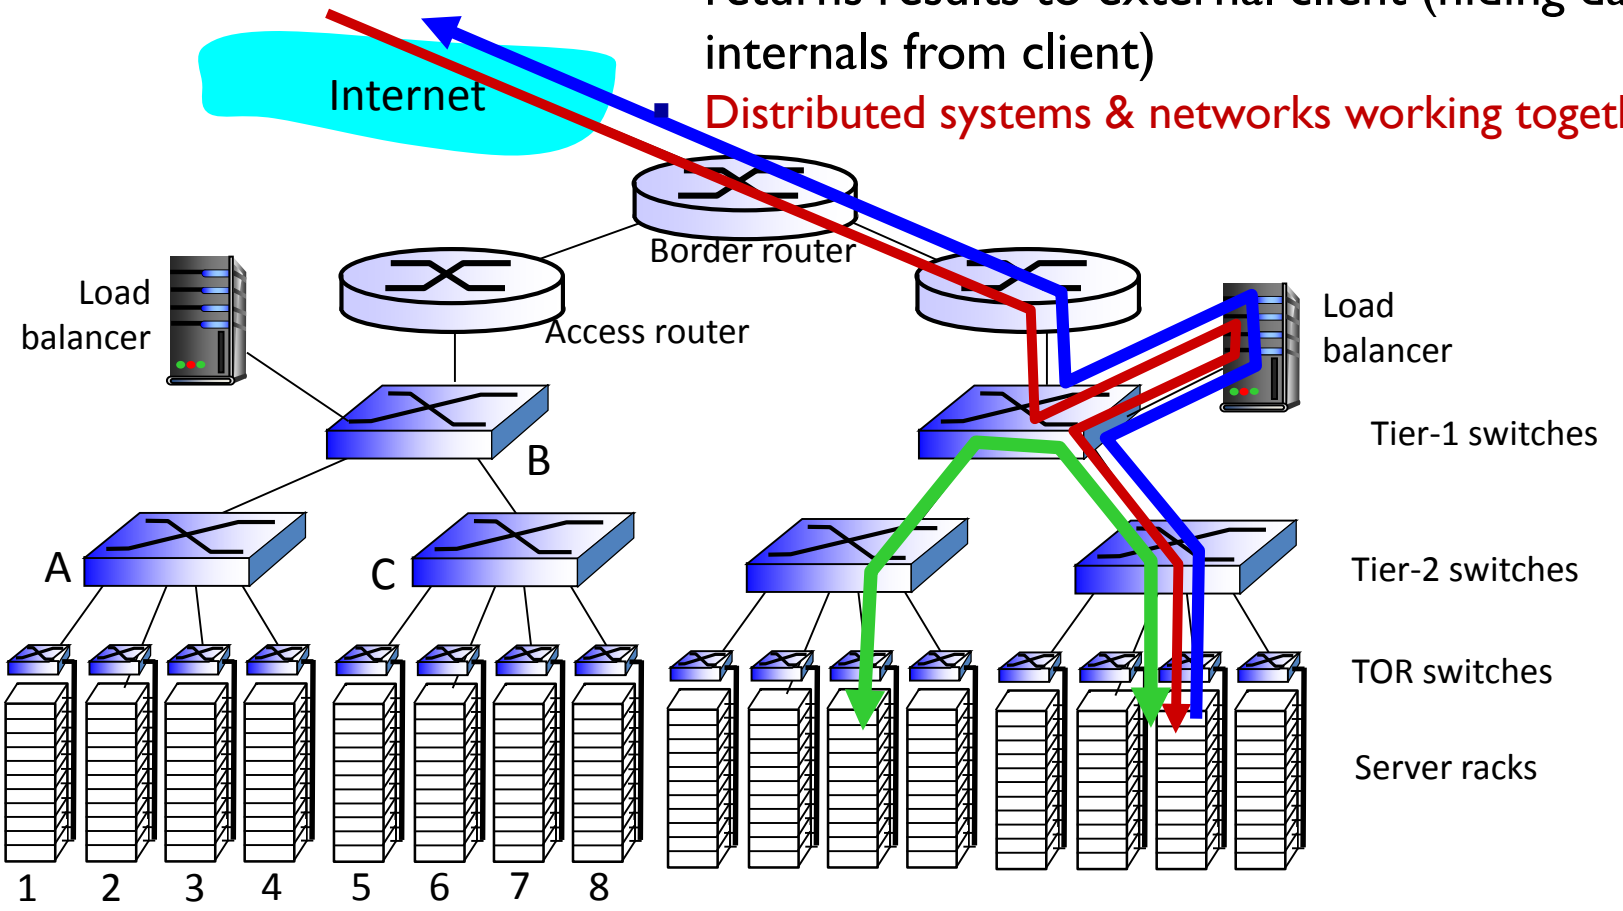

Data center networks

- rich interconnection among switches, racks:
  - increased throughput between racks (multiple routing paths possible)
  - increased reliability via redundancy
  - Distributed systems & networks working together

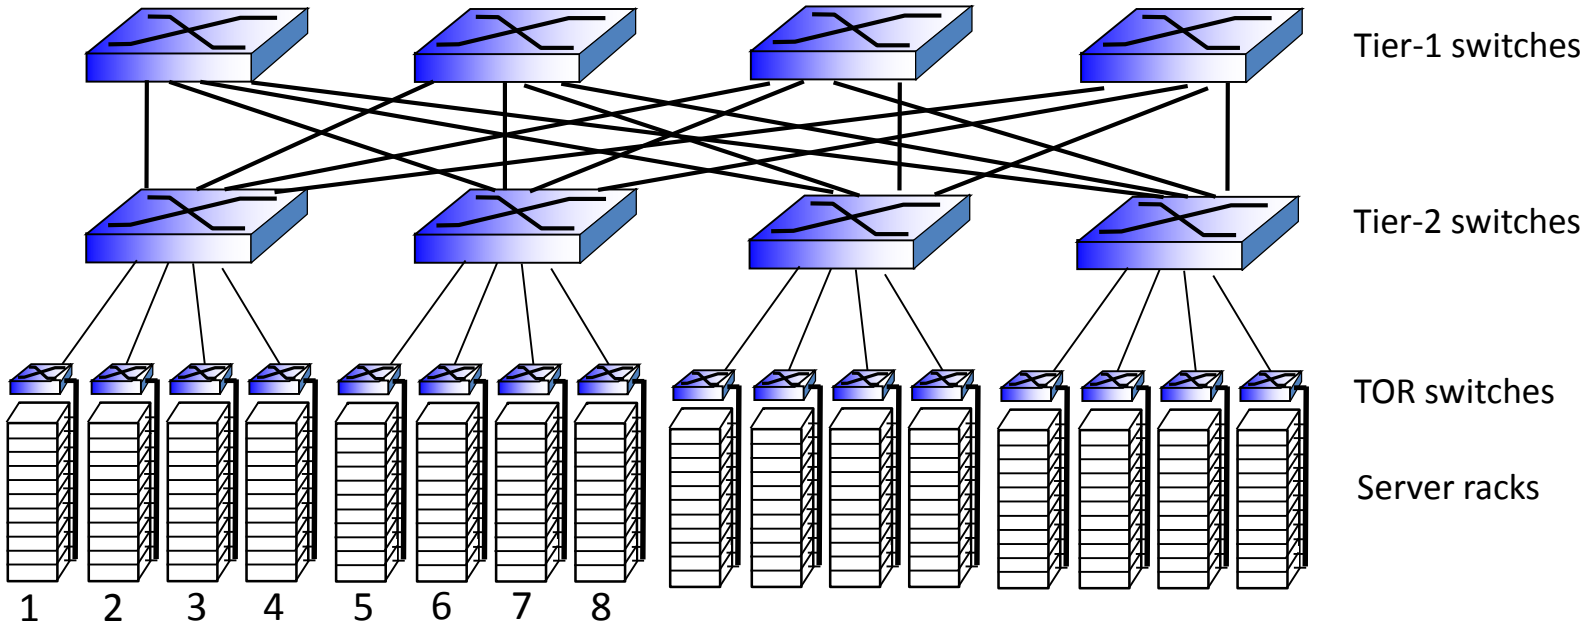

Extra (guest) lecture (joint with Adv. Distributed Systems course), Friday 10-11 room EE

Niklas Gustavsson, backend engineering team leader at Spotify, Gothenburg

Title:

"Gossiping (ie randomized multicast routing) and Conflict-free_replicated_data_types" (eg distributed hash tables and/or related)

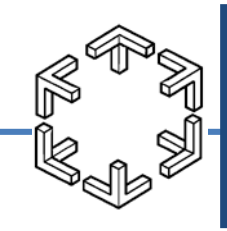

More examples: a story in progress + possible followup course...

Overlays useful here, too:

New power grids: be adaptive!

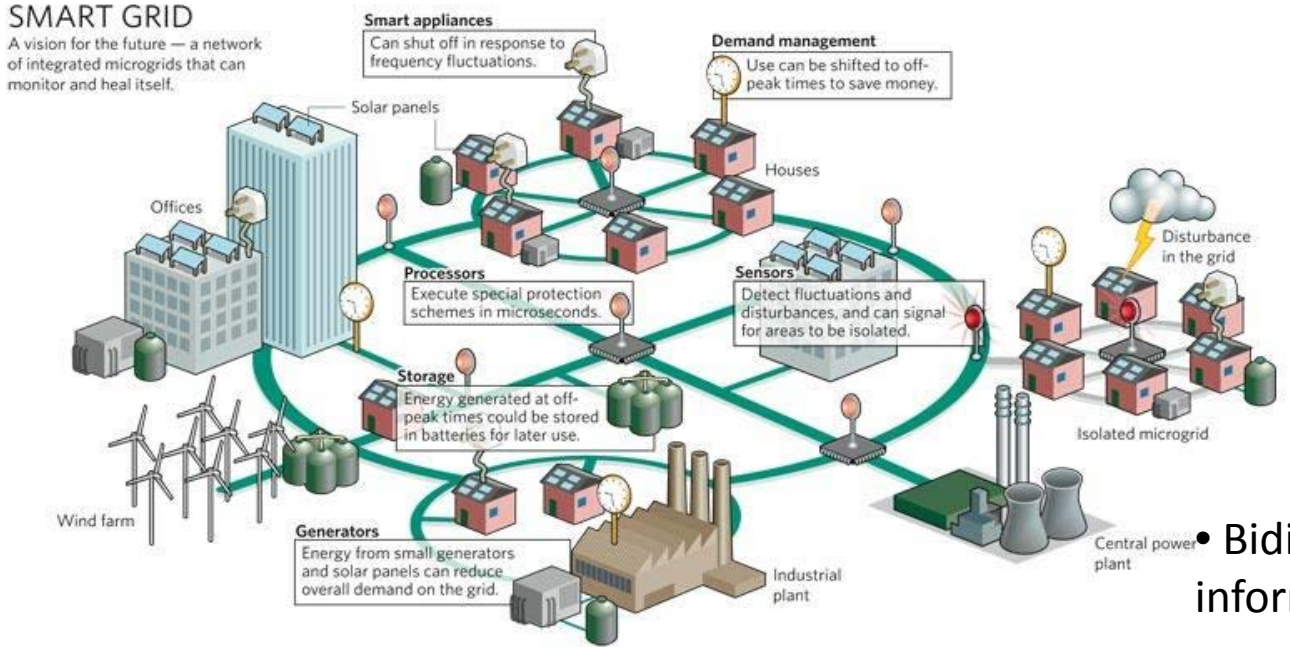

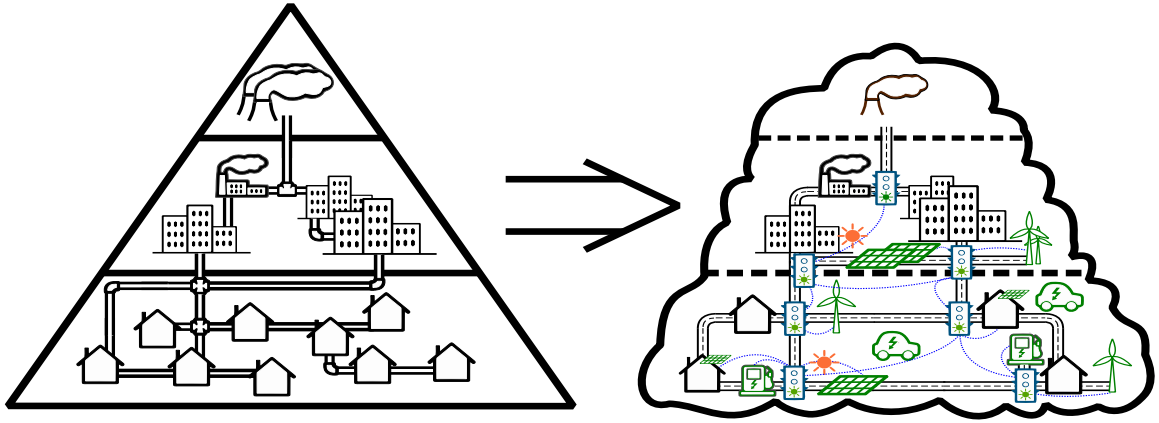

- Bidirectional power and information flow
  - – Micro-producers or "prosumers", can share resources
  - – Distributed energy resources
- Communication + resource-administration (distributed system) layer: IoT

El-networks as distributed cyber-physical systems

Overlay network

El- link and/or communication link

Computing+ communicating device

Cyber system

Why adding "complexity" in the infrastructure?

Motivation: enable renewables, better use of el-power

Course/Masterclass: (DAT300, LP1)
ICT Support for Adaptiveness and Security in the Smart Grid

- Goals
  - Students (CSE and other disciplines) get introduced to advanced interdisciplinary concepts related to the smart grid, thus
  - building an understanding of essential notions in the individual disciplines, and
  - investigating a domain-specific problem relevant to the smart grid that need an understanding beyond the traditional ICT field.

Environment

- Based on both the present and future design of smart systems.
  - How can techniques from **networks/distributed systems be applied to large, heterogeneous systems** where a **massive amount of data** must be collected/processed?
  - How can such a system, containing legacy components with no security primitives, be **made secure** when the communication is added by interconnecting the systems?
- The students will have access to a hands-on lab, where they can run and test their design and code.

Course Setup

- The course is given on an advanced master's level, resulting in 7.5 points.
- Study Period 1
  - Can also define individual, "research internship courses", 7.5, 15p or MS thesis, starting earlier
- The course structure
  - lectures to introduce the two disciplines ("crash course-like"); invited talks by industry and other collaborators
  - second part: seminar-style where research papers from both disciplines are presented & discussed.
  - At the end of the course the students are also expected to present their projects.

Thank you

Recall, important for the exam:

When/where: wednesday March 16, 14.00-18.00, M

You may have with you:

- English-X dictionary
- no calculators, PDAs, etc (if/where numbers matter, do rounding)

To think during last, summary-study

Overview; critical eye; explain; ask yourselves: why is this so? / How does it work?

Good luck with all your efforts!!!

"If you hear a voice within you say 'you cannot paint,' then by all means paint, and that voice will be silenced." – Vincent Van Gogh