High-level summary …

TDTS21 Advanced Networking

Niklas Carlsson, Associate Professor
http://www.ida.liu.se/~nikca/

Kick starting science ...

… well, cable into wall …

What happens there?

Hosts, the Internet architecture, and the E2E arguments …

End hosts …


cingular
9:41 AM
Text
Calendar
Photos
Camera
Calculator
Stocks
Maps
Weather
Notes
Clock
Settings

How to find who to talk to?

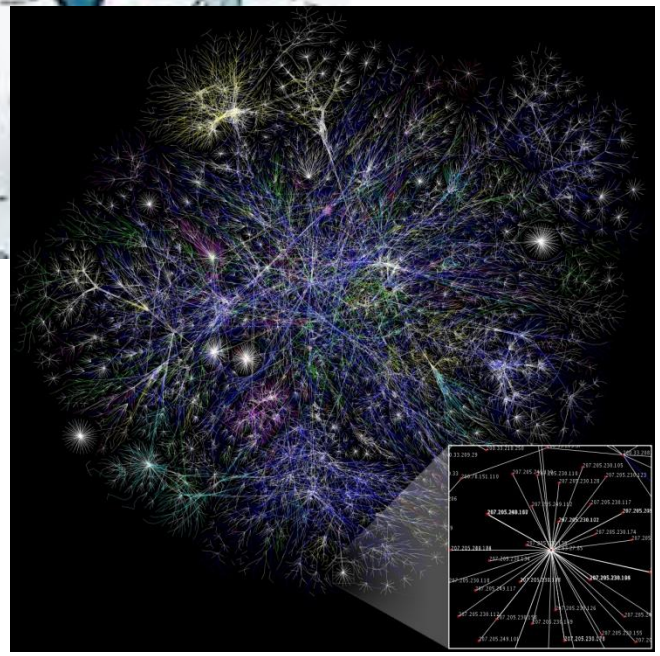

Learning a Host's Address

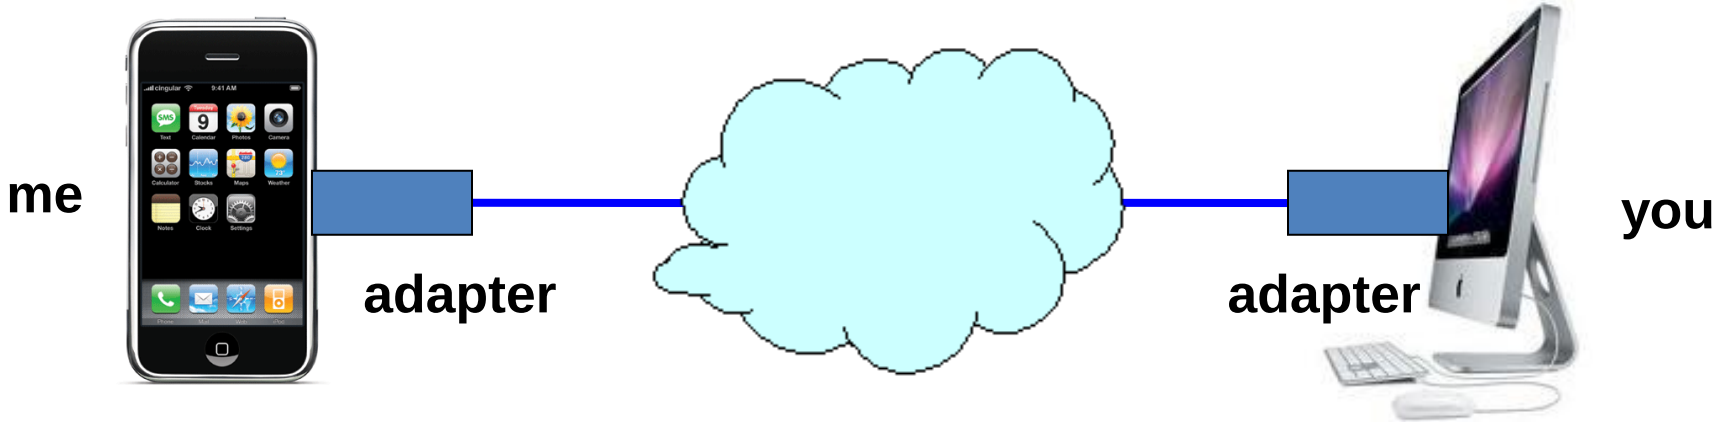

- Who am I?
  - Hard-wired: MAC address
  - Static configuration: IP interface configuration
  - Dynamically learned: IP address configured by DHCP
- Who are you?
  - Hard-wired: IP address in a URL, or in the code
  - Dynamically looked up: ARP or DNS

Goals of the Internet Architecture (Clark '88)

1. **Connect existing networks**
2. Robust in face of failures (not nuclear war…)
3. Support multiple types of services
4. Accommodate a variety of networks
5. Allow distributed management
6. Easy host attachment
7. Cost effective
8. Allow resource accountability

Real Goals

1. **Something that works…..**
2. Connect existing networks
3. Survivability (not nuclear war…)
4. Support multiple types of services
5. Accommodate a variety of networks
6. Allow distributed management
7. Easy host attachment
8. Cost effective
9. Allow resource accountability

Host-Network Division of Labor

- Network
  - Best-effort packet delivery
  - Between two (or more) end-point addresses
- Hosts
  - Everything else

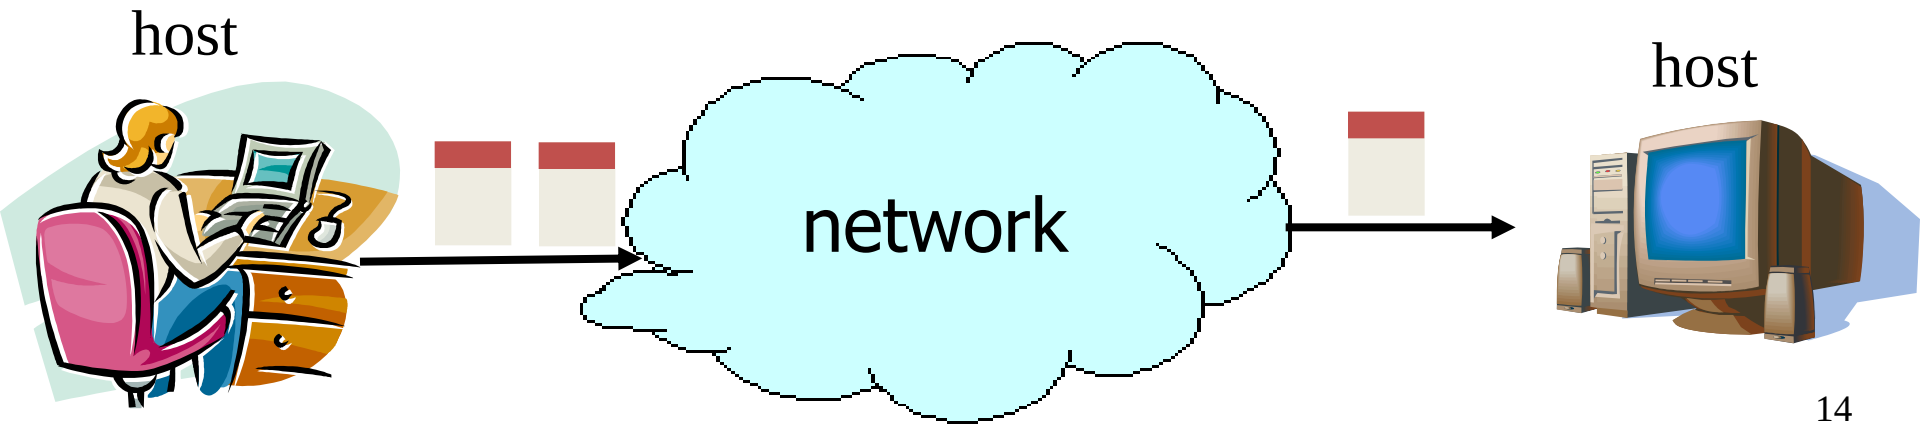

Network Stack in Practice

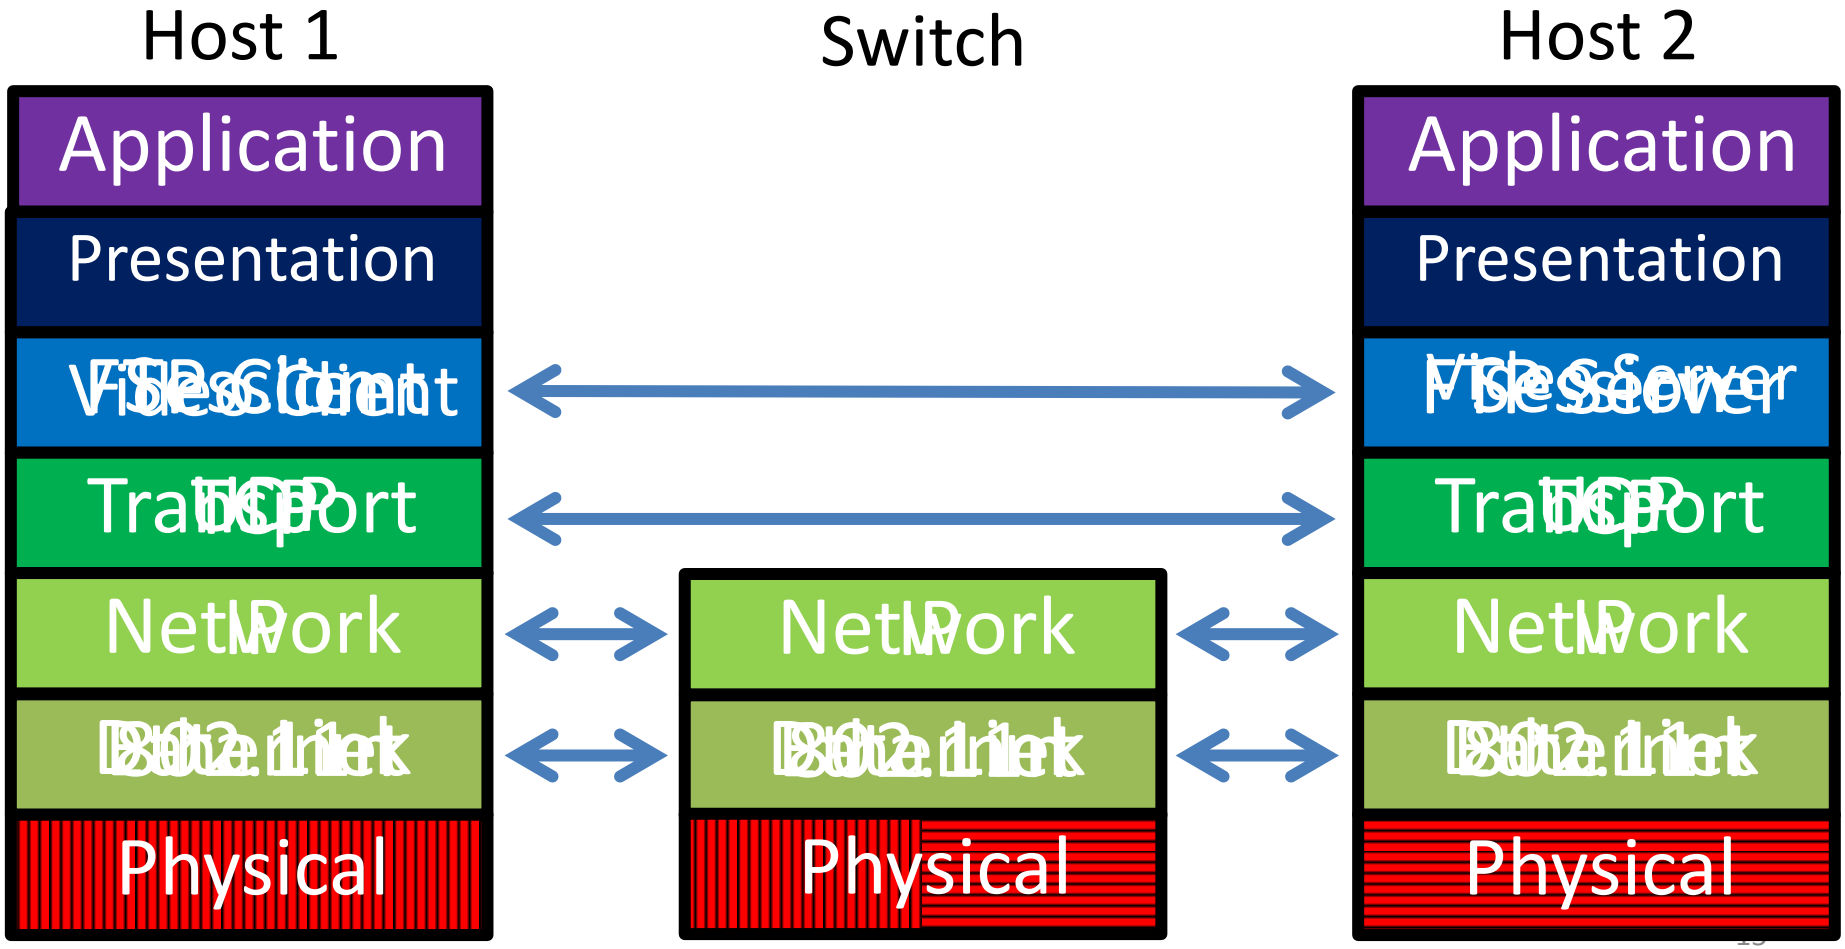

Encapsulation, Revisited

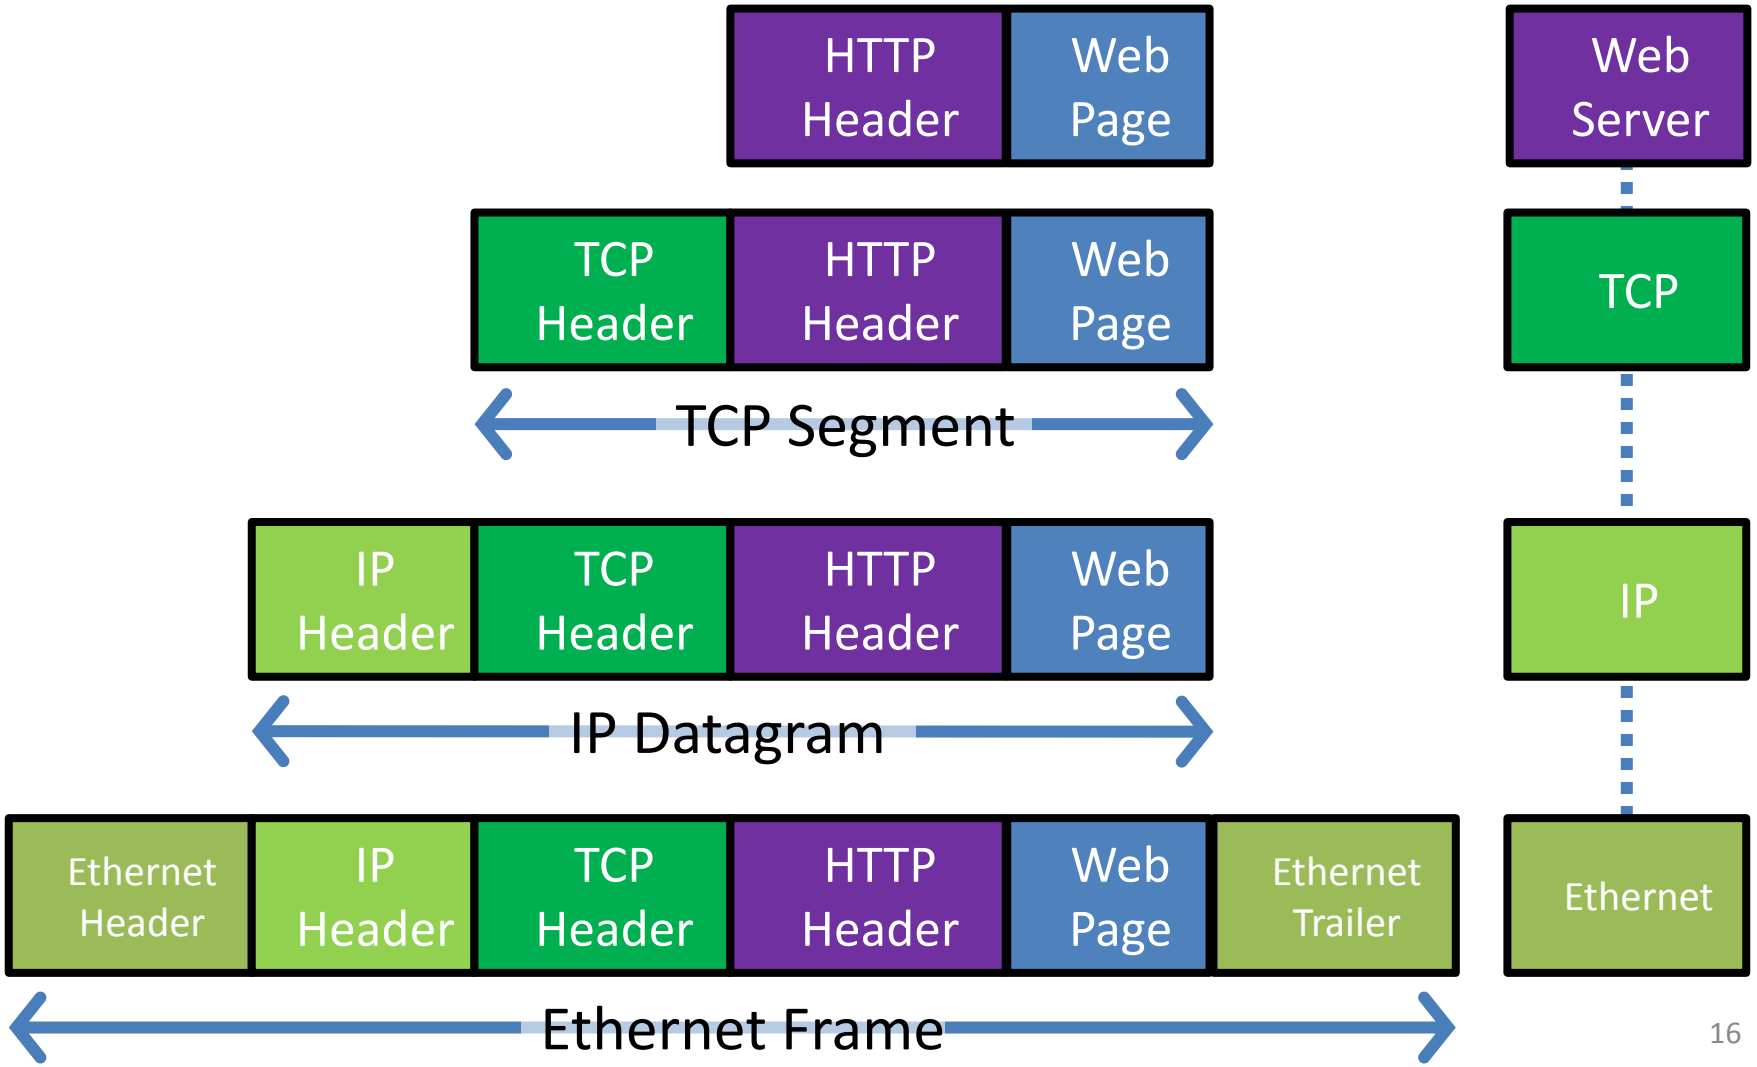

The Hourglass

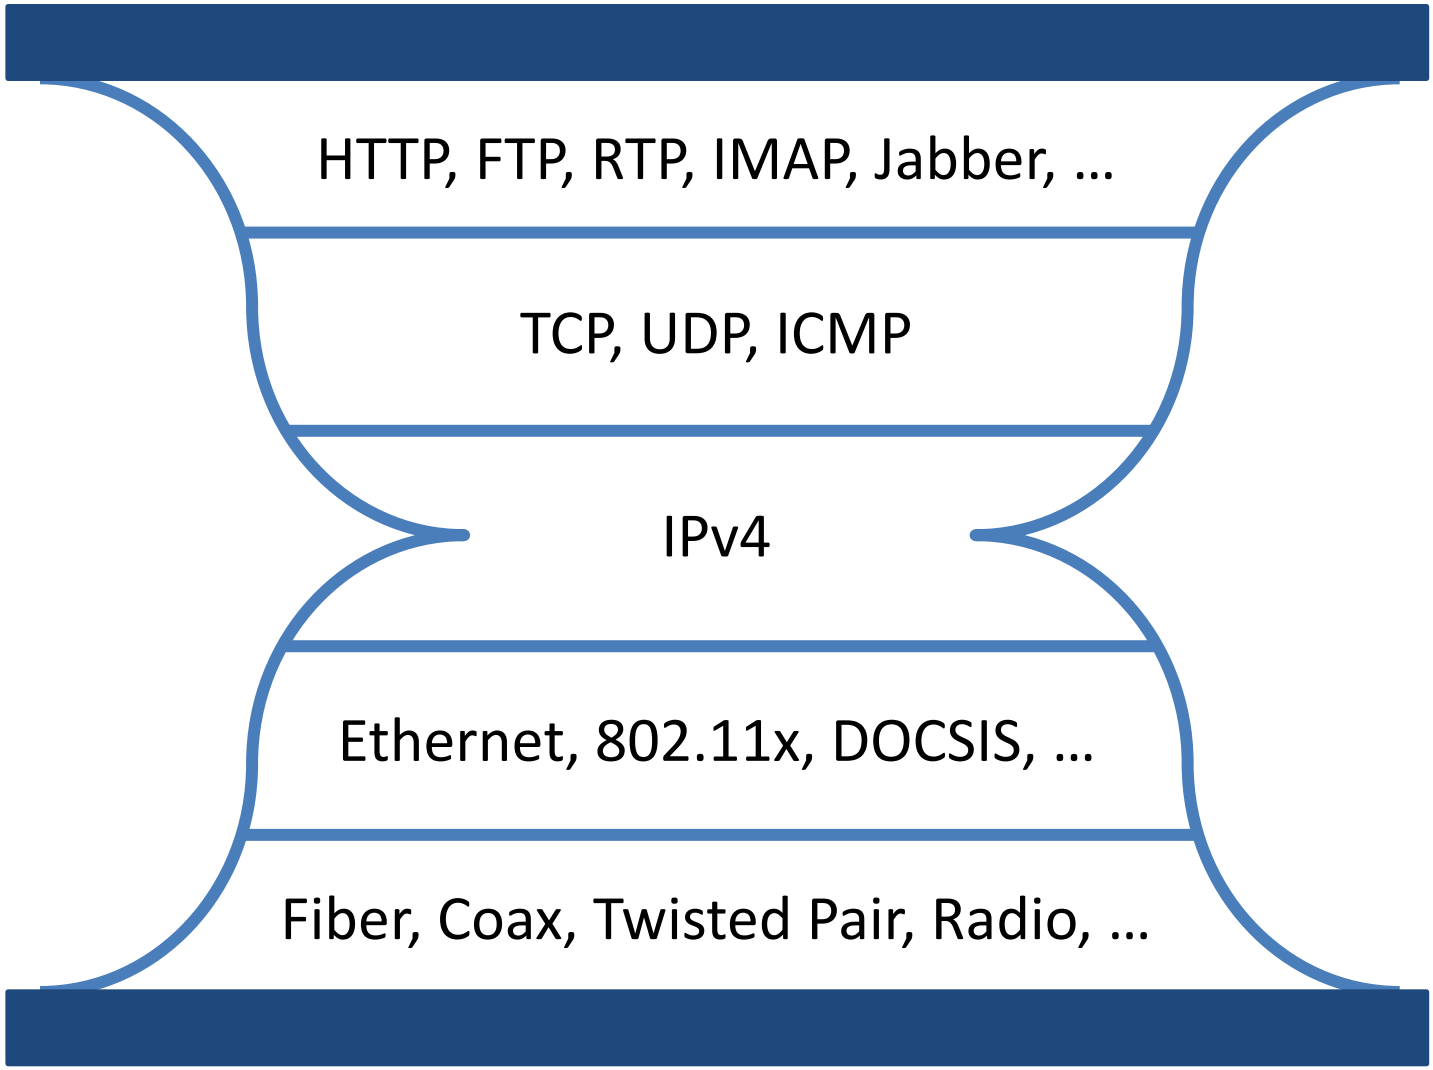

Orthogonal Planes

Control plane: How **Internet paths** are established

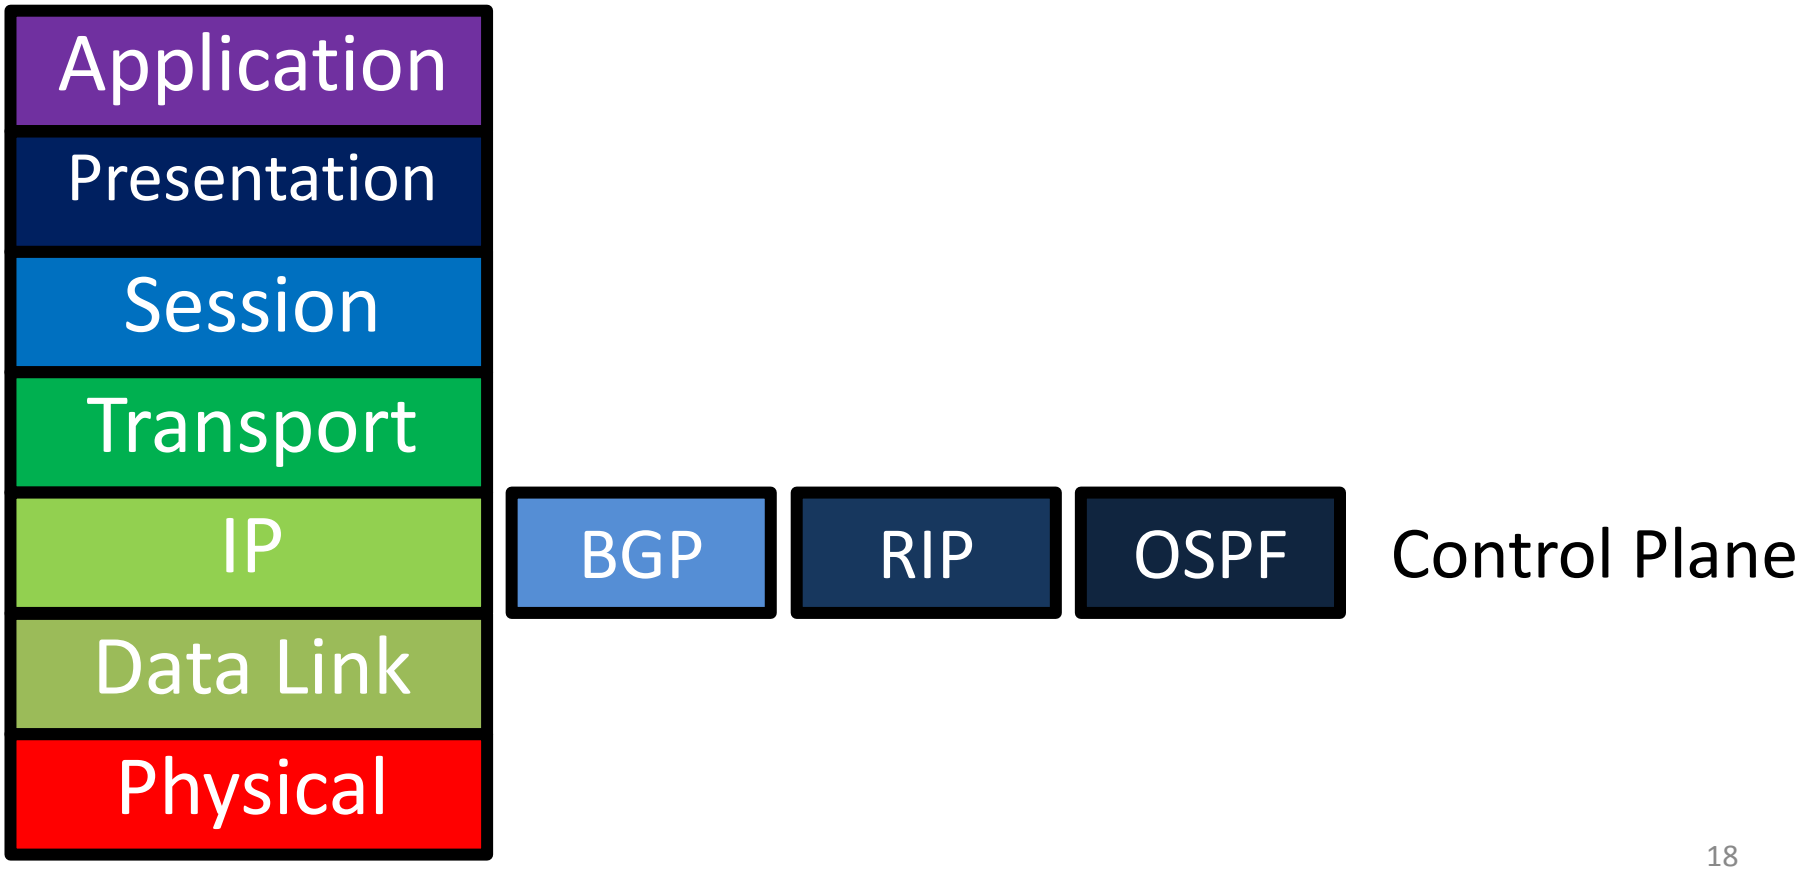

Orthogonal Planes

Data plane: How data is **forwarded** over Internet paths

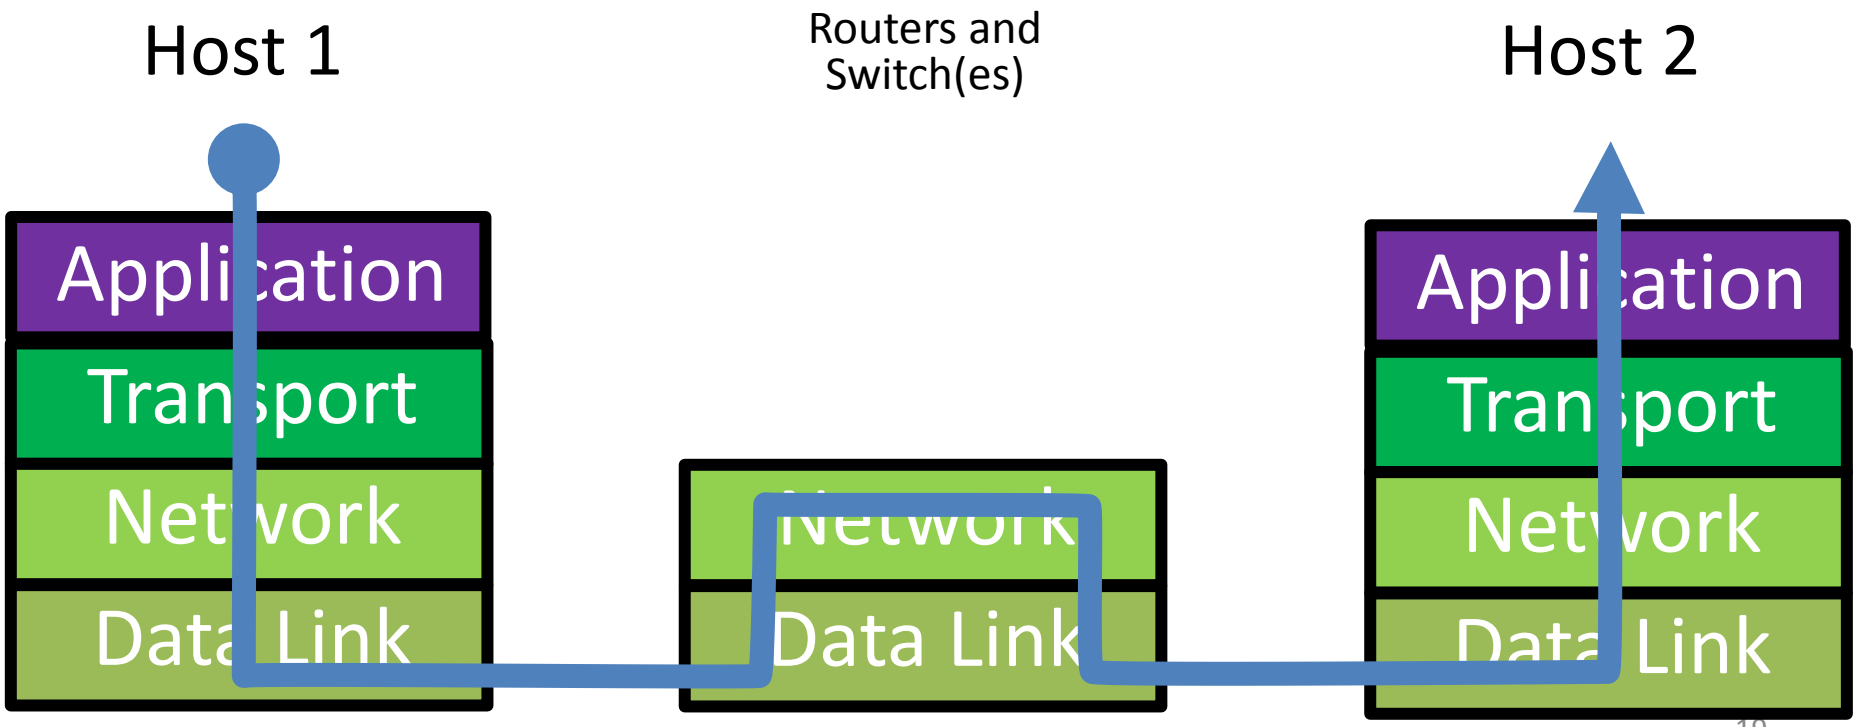

Reality Check

- The layered abstraction is very nice
- Does it hold in reality?

No.

Firewalls

- Analyze application layer headers

Transparent Proxies

- Simulate application endpoints within the network

NATs

- Break end-to-end network reachability

Holding the Internet Together

- Distributed cooperation for resource allocation
  - BGP: what end-to-end *paths* to take (for ~50K ASes)
  - TCP: what *rate* to send over each path (for ~3B hosts)

How do we find a path?

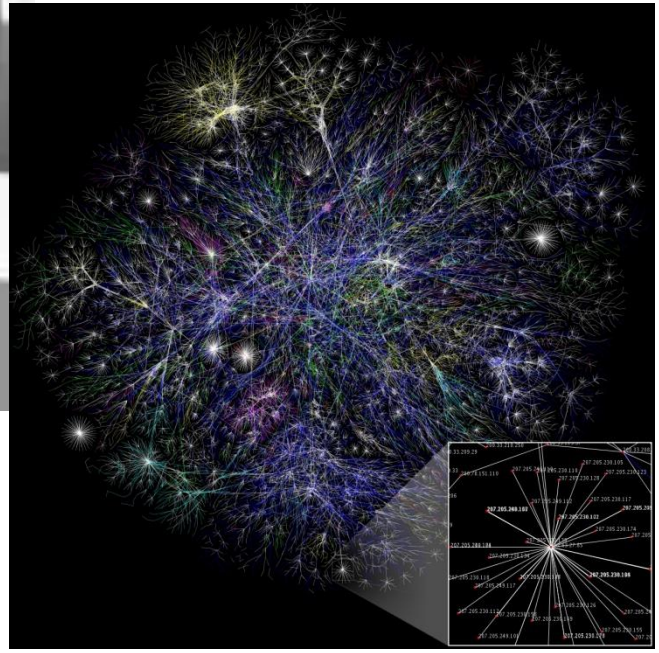

Routing on a Graph

- Goal: determine a "good" path through the network from source to destination
- What is a good path?
  - Usually means the shortest path
  - Load balanced
  - Lowest $$$ cost
- Network modeled as a graph
  - Routers → nodes
  - Link → edges
    - Edge cost: delay, congestion level, etc.

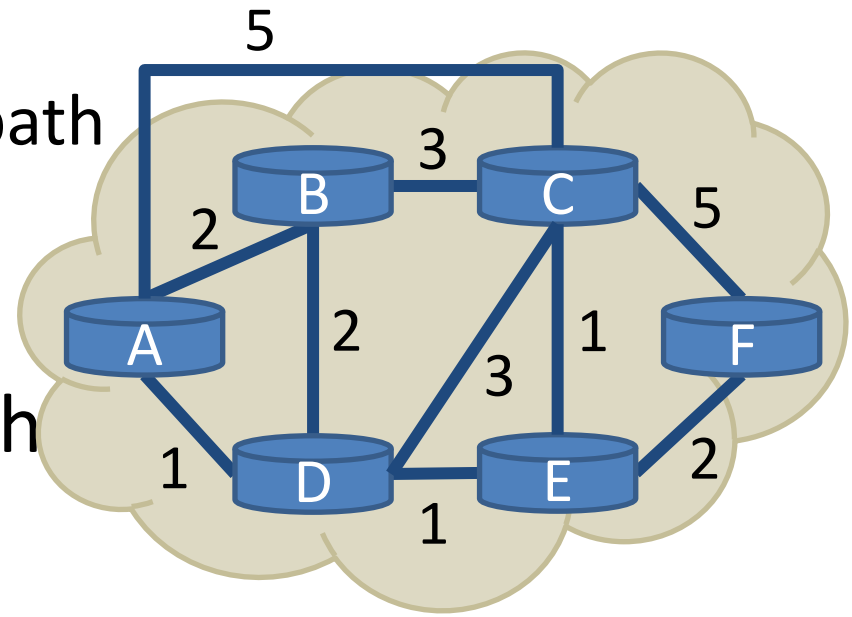

Intra-domain Routing Protocols

- Distance vector
  - Routing Information Protocol (RIP), based on Bellman-Ford
  - Routers periodically exchange reachability info with neighbors
- Link state
  - Open Shortest Path First (OSPF), based on Dijkstra
  - Each network periodically floods neighbor information to all routers
  - Routers locally compute routes

	Link State	Distance Vector
Message Complexity	$O(n^2*e)$	O(d*n*k)
Time Complexity	O(n*log n)	O(n)
Convergence Time	O(1)	O(k)
Robustness	• Nodes may advertise incorrect link costs • Each node computes their own table	• Nodes may advertise incorrect path cost • Errors propagate due to sharing of DV tables

Hierarchical addressing: route aggregation

ISP has an address block; it can further divide this block into sub blocks and assign them to subscriber organizations.

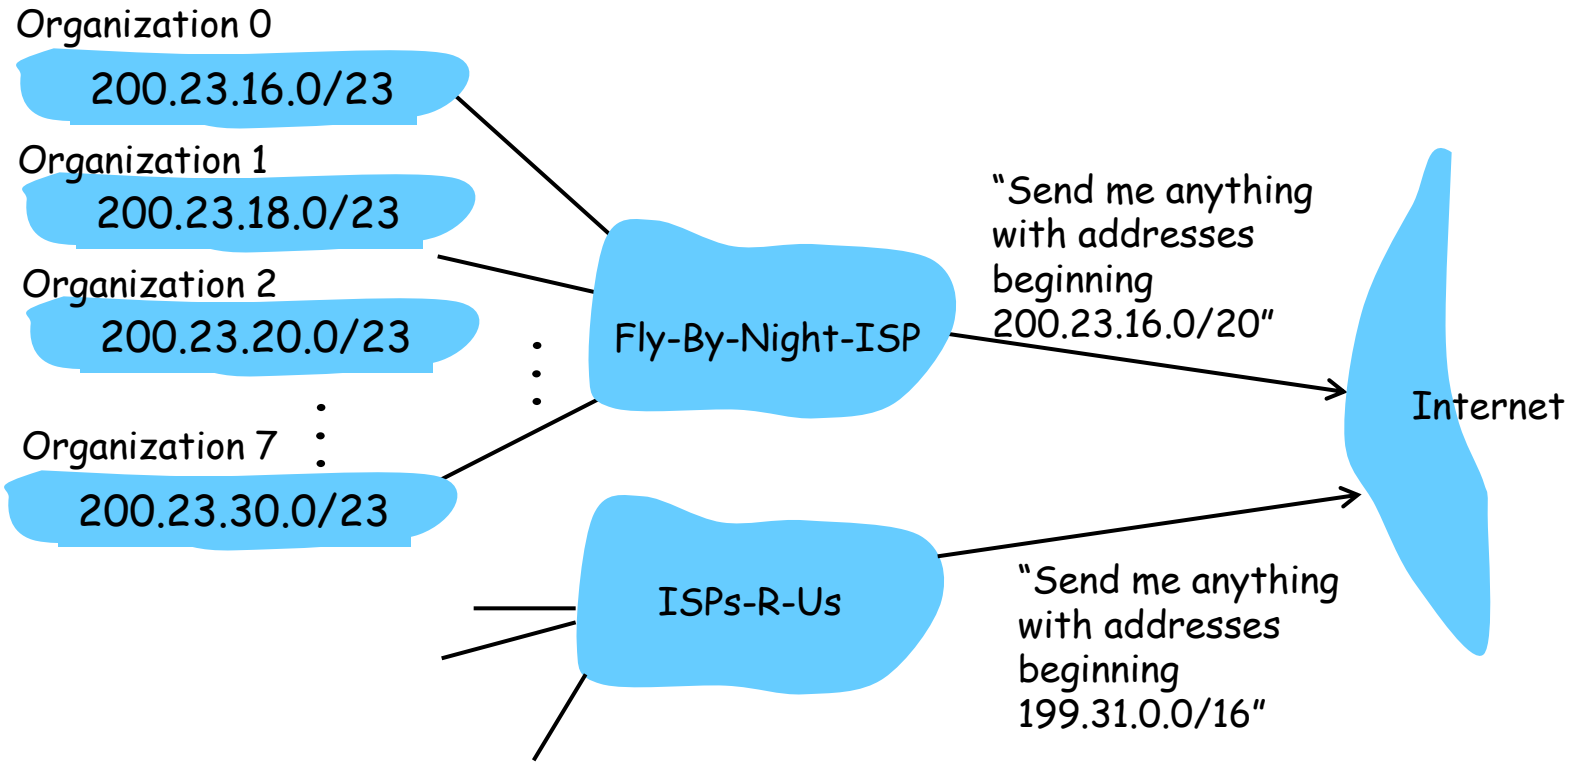

Example CIDR Routing Table

Address	Netmask	Third Byte	Byte Range
207.46.0.0	19	000xxxxx	0 – 31
207.46.32.0	19	001xxxxx	32 – 63
207.46.64.0	19	010xxxxx	64 – 95
207.46.128.0	18	10xxxxxx	128 – 191
207.46.192.0	18	11xxxxxx	192 – 255

Hole in the Routing Table: No coverage for 96 – 127

207.46.96.0/19

Network of networks: BGP and ASes

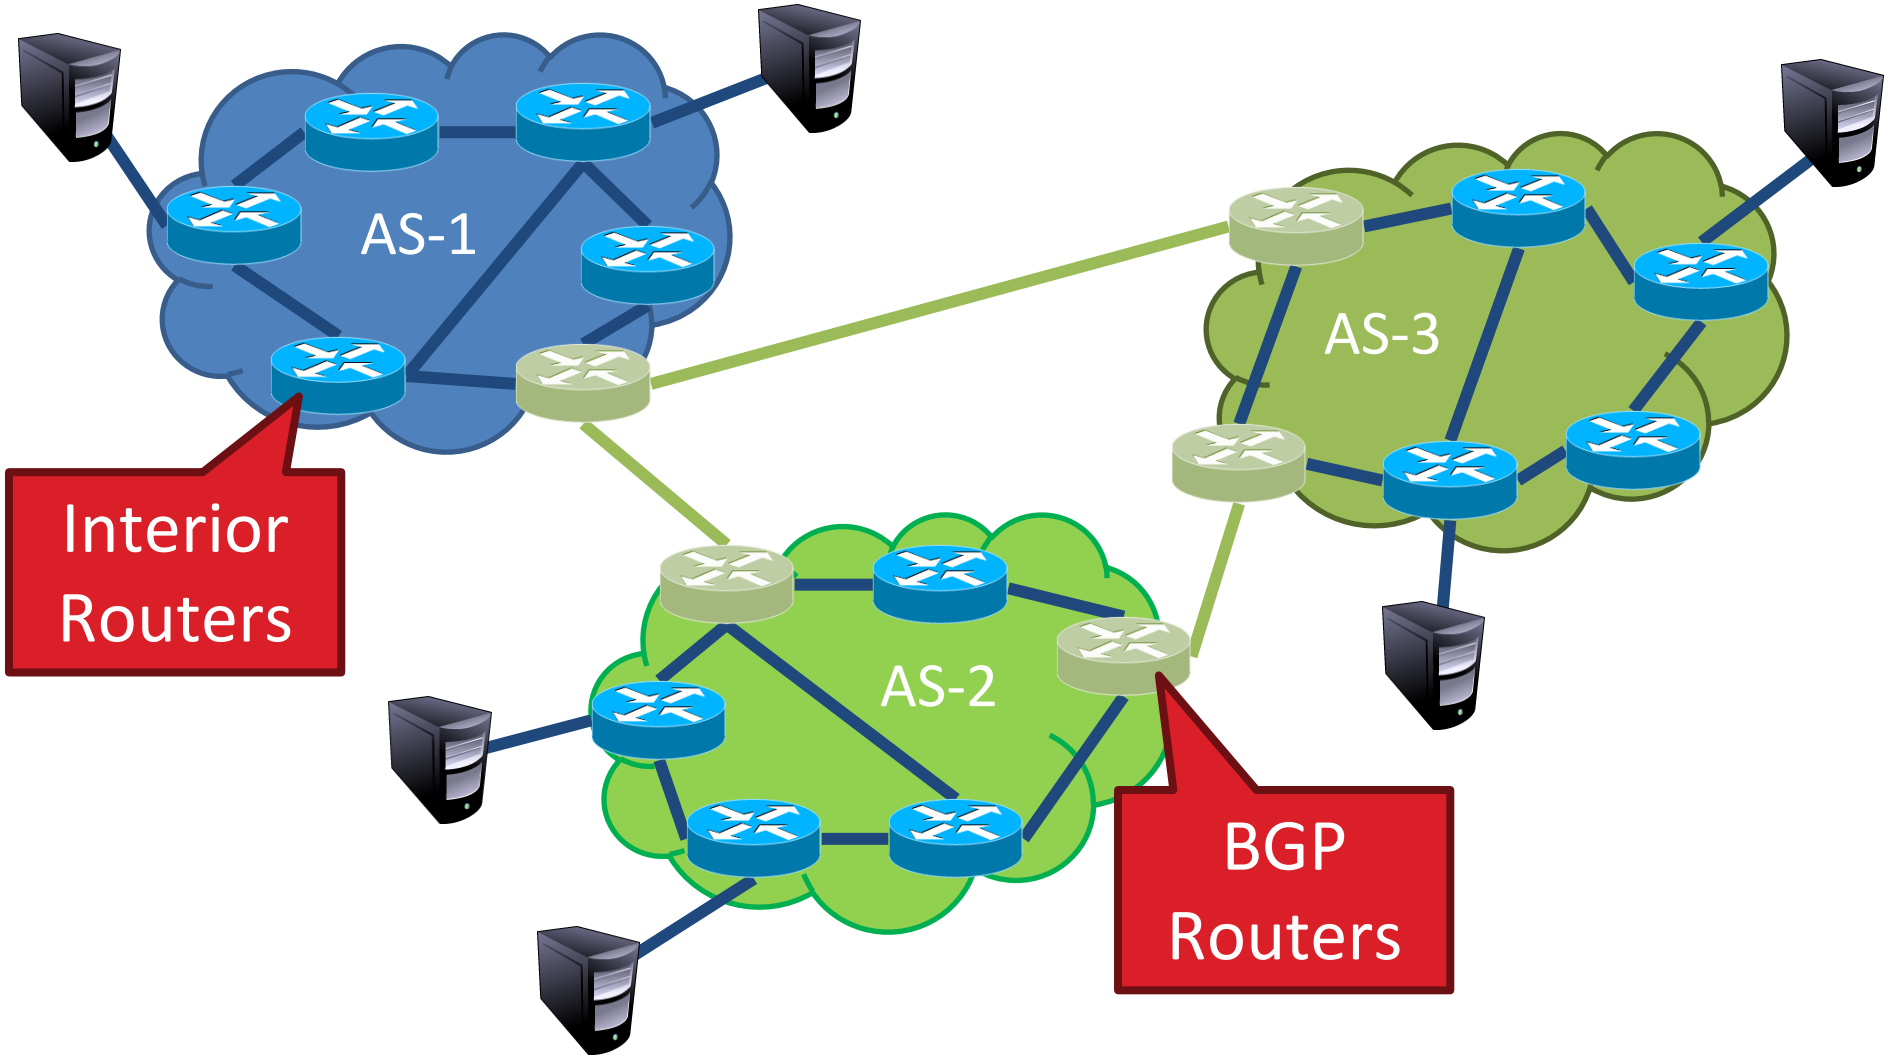

BGP Relationships

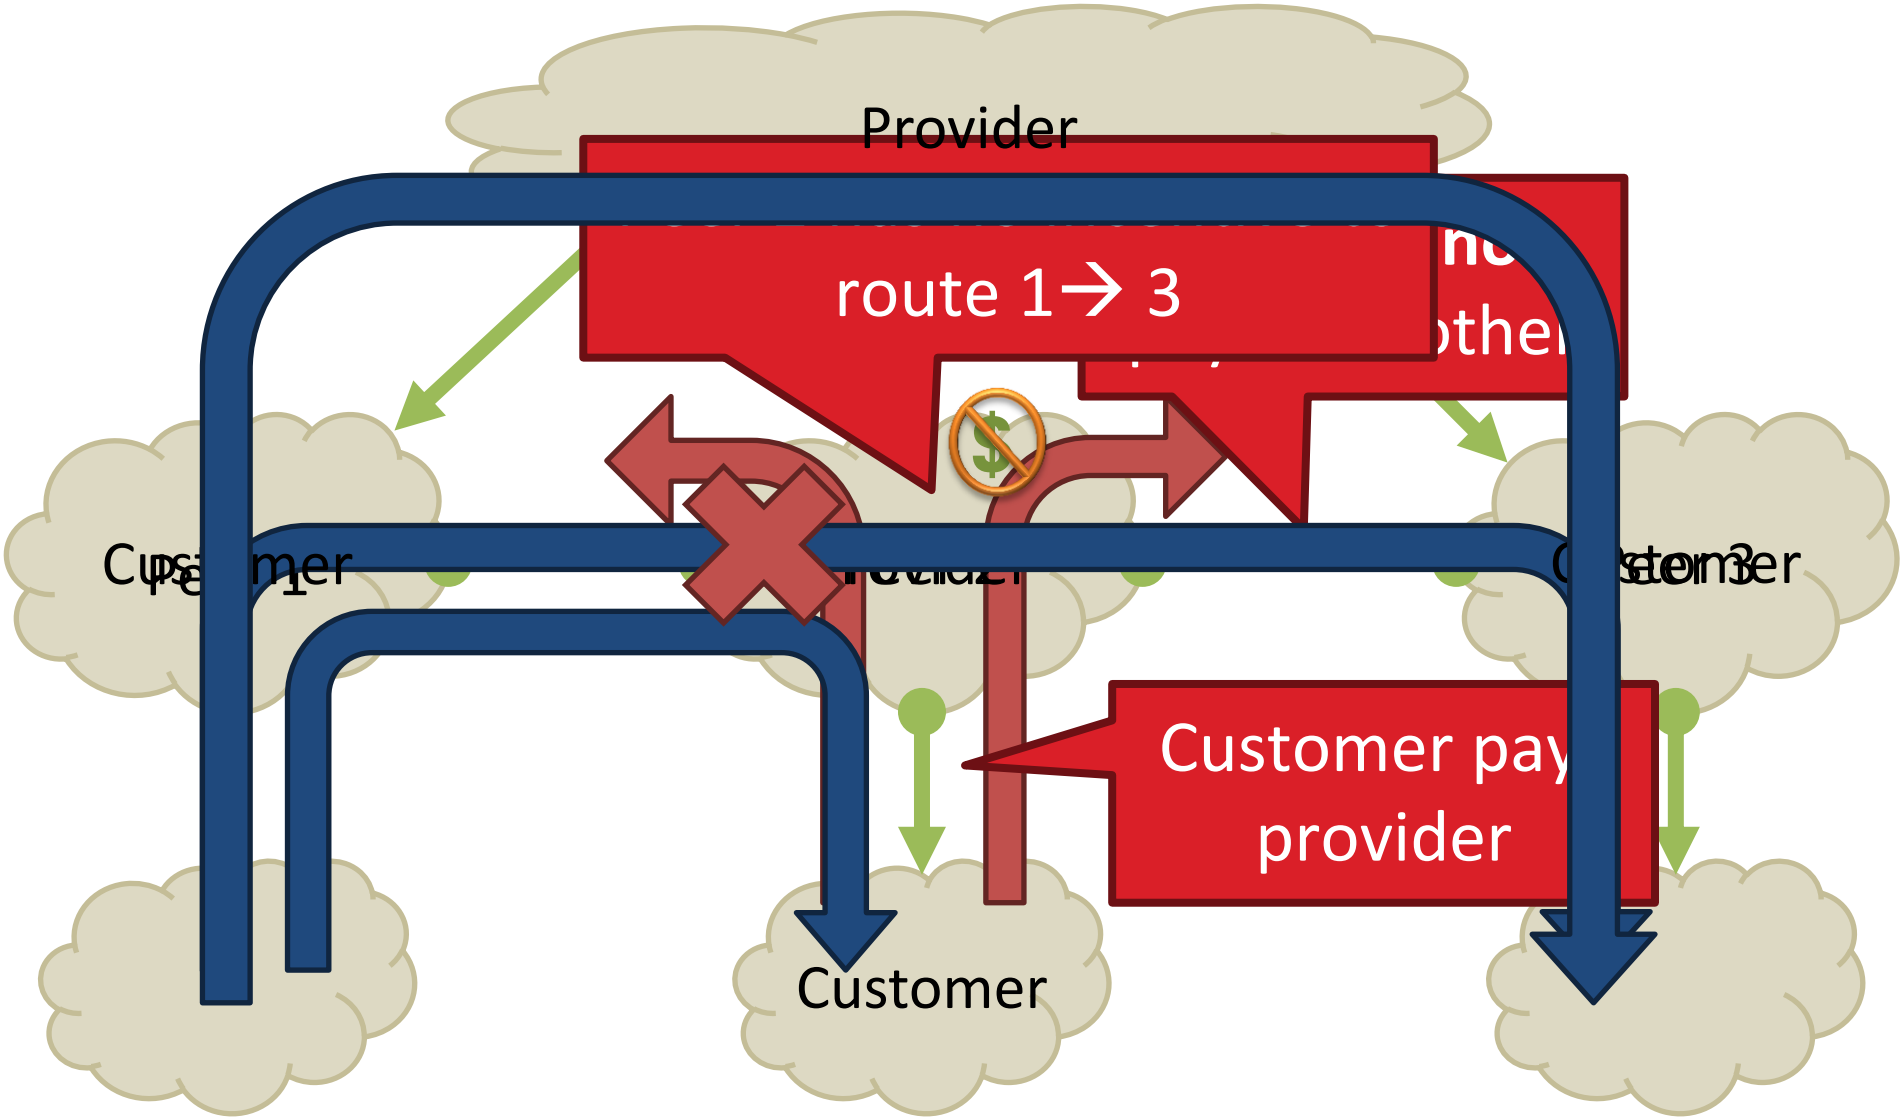

Importing Routes

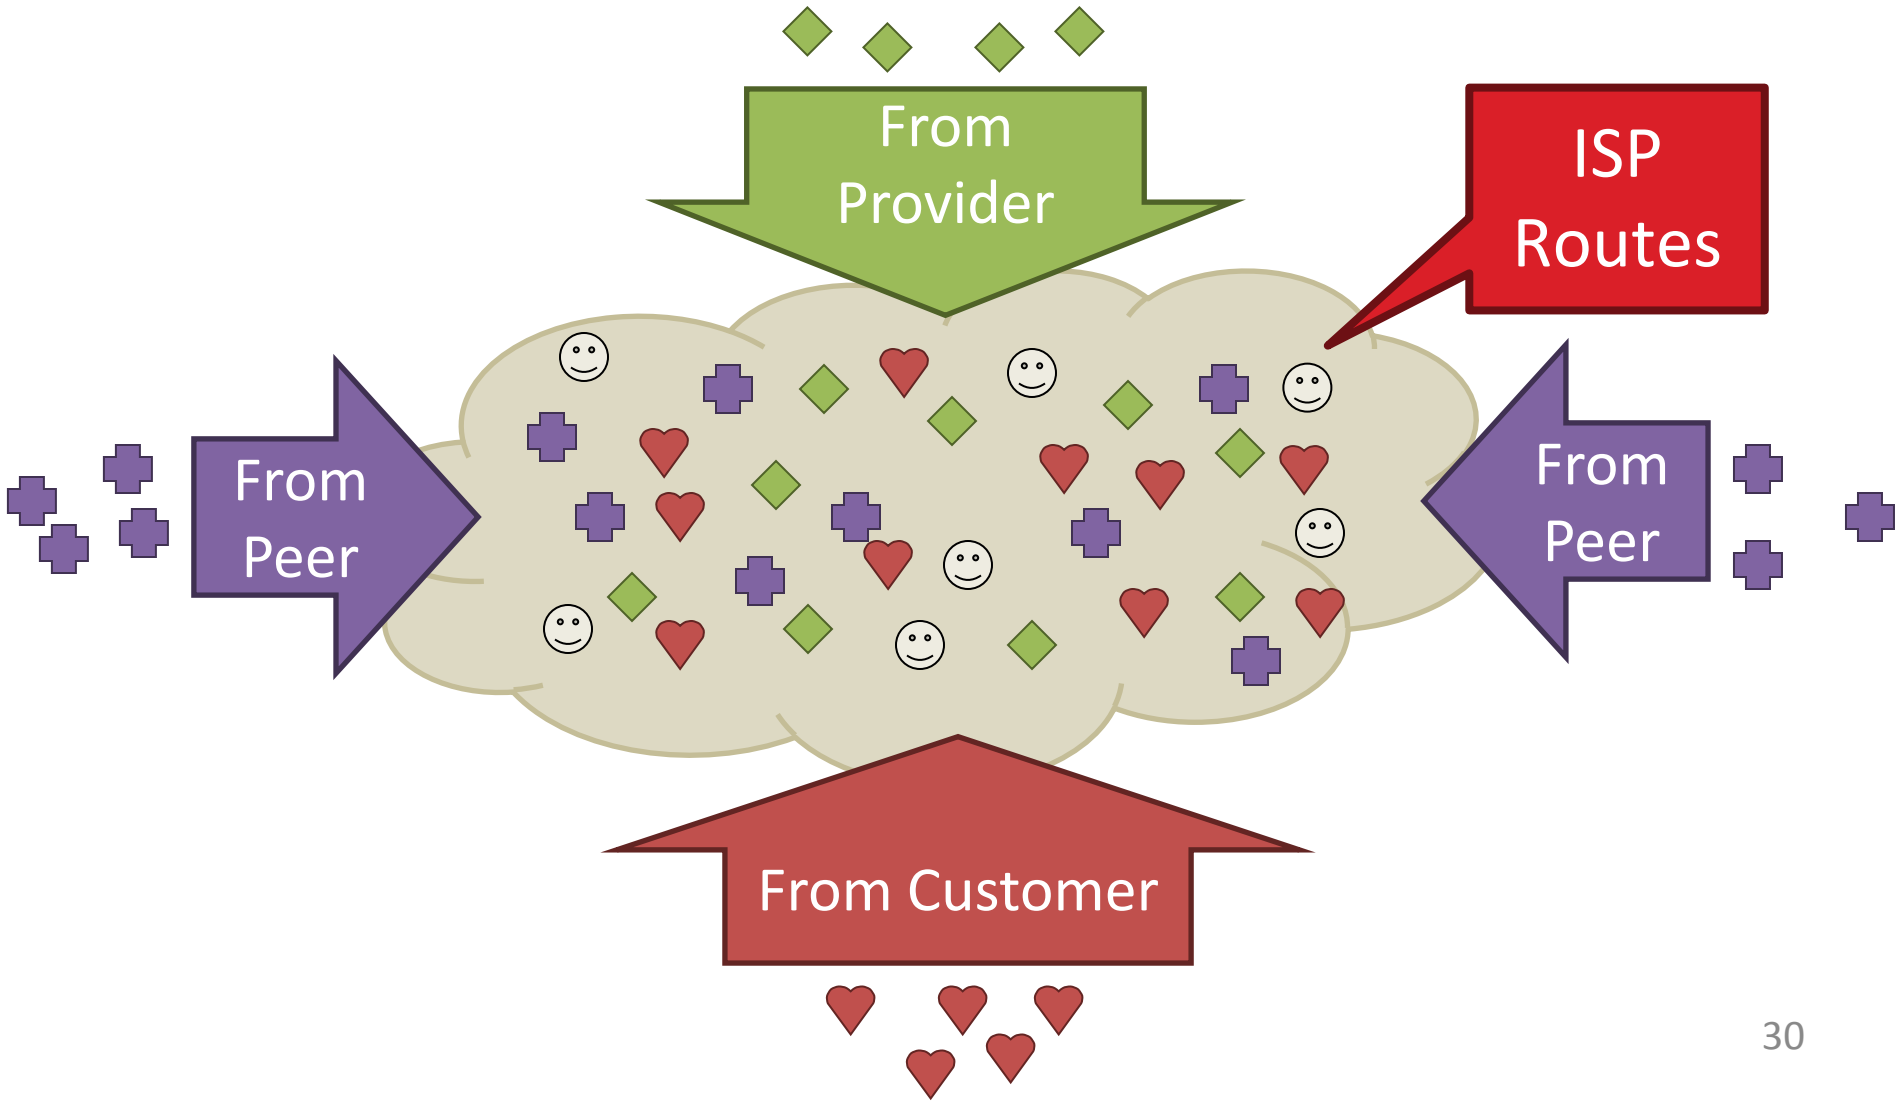

Exporting Routes

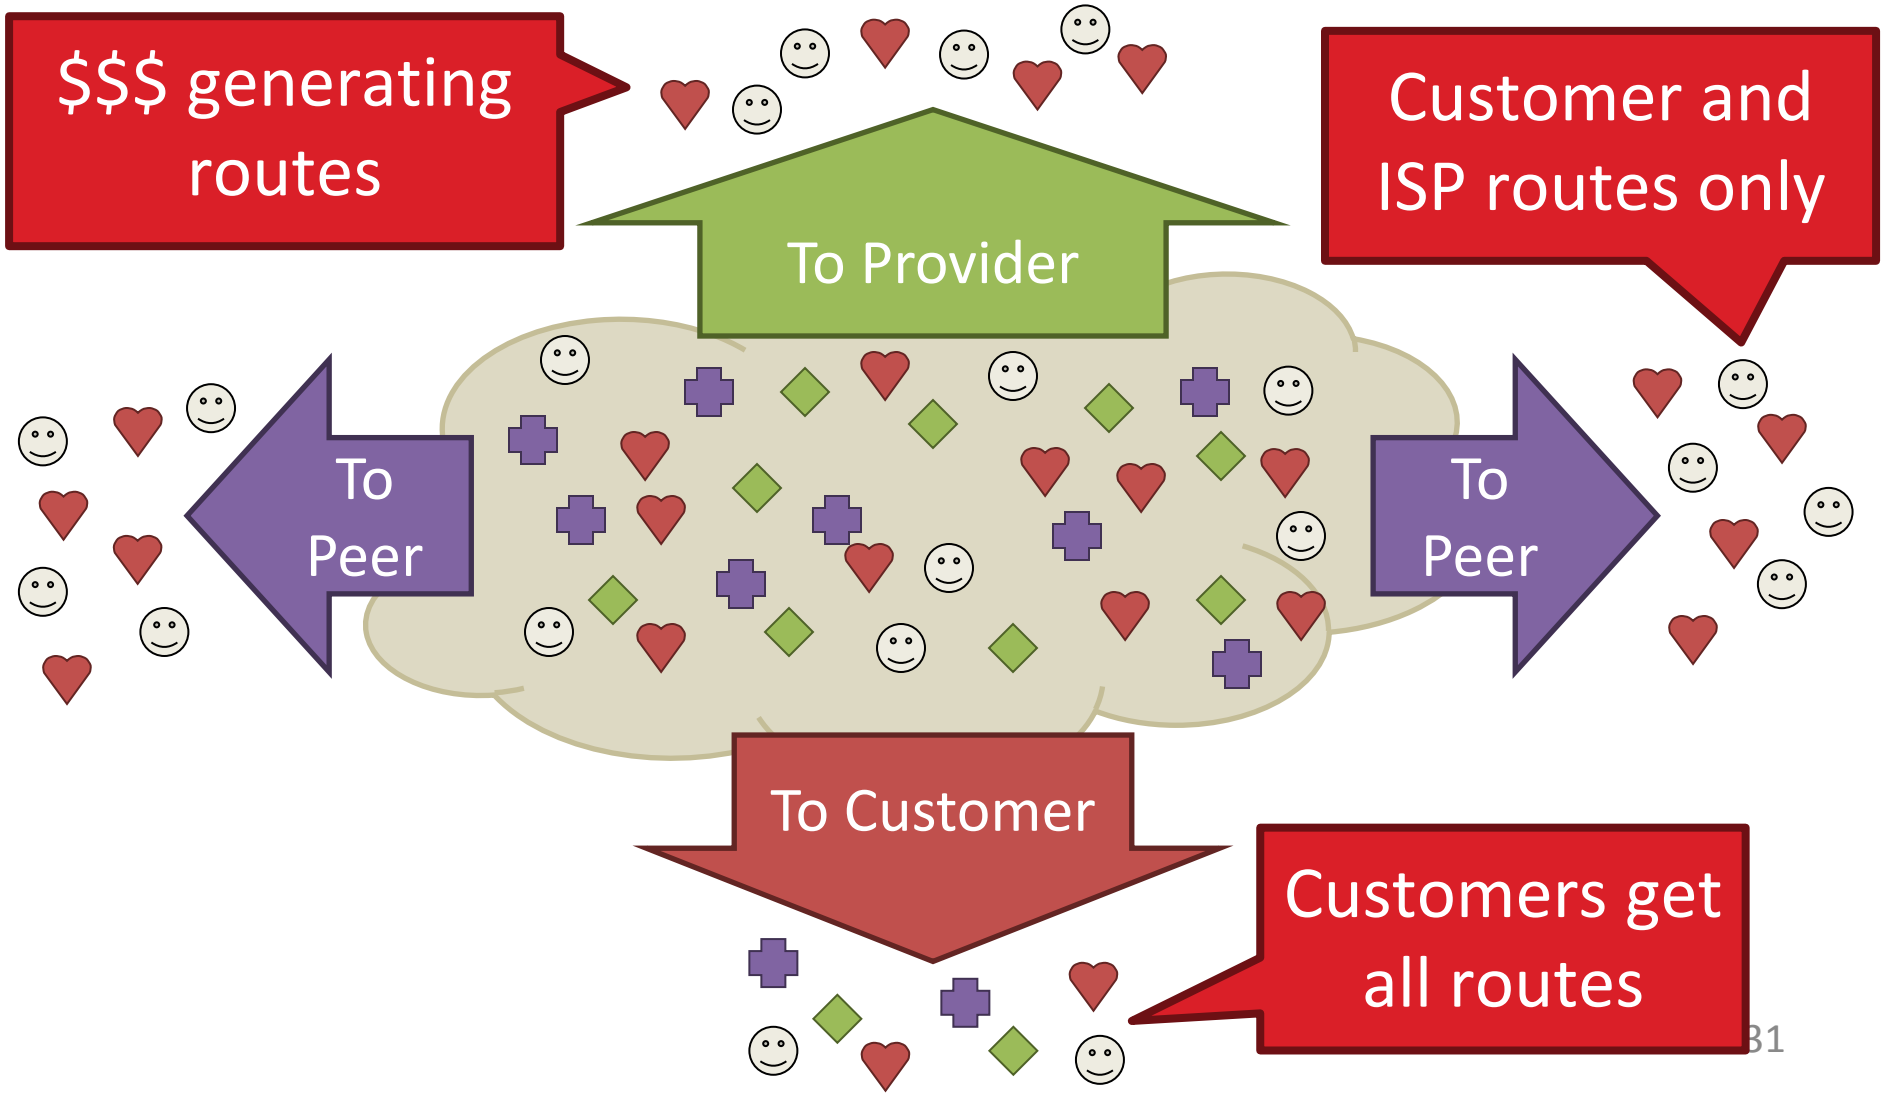

Modeling BGP

- AS relationships
  - Customer/provider
  - Peer
  - Sibling, IXP
- Gao-Rexford model
  - AS prefers to use customer path, then peer, then provider
    - Follow the money!
  - Valley-free routing
  - Hierarchical view of routing (incorrect but frequently used)

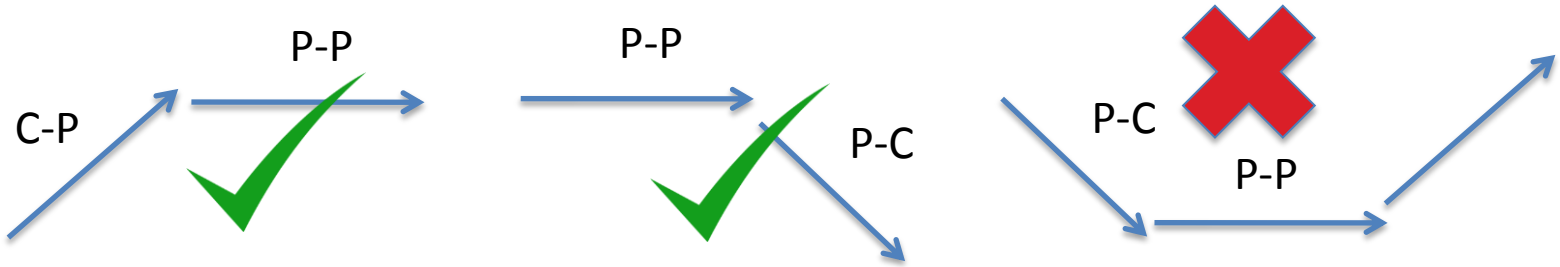

A new Internet model

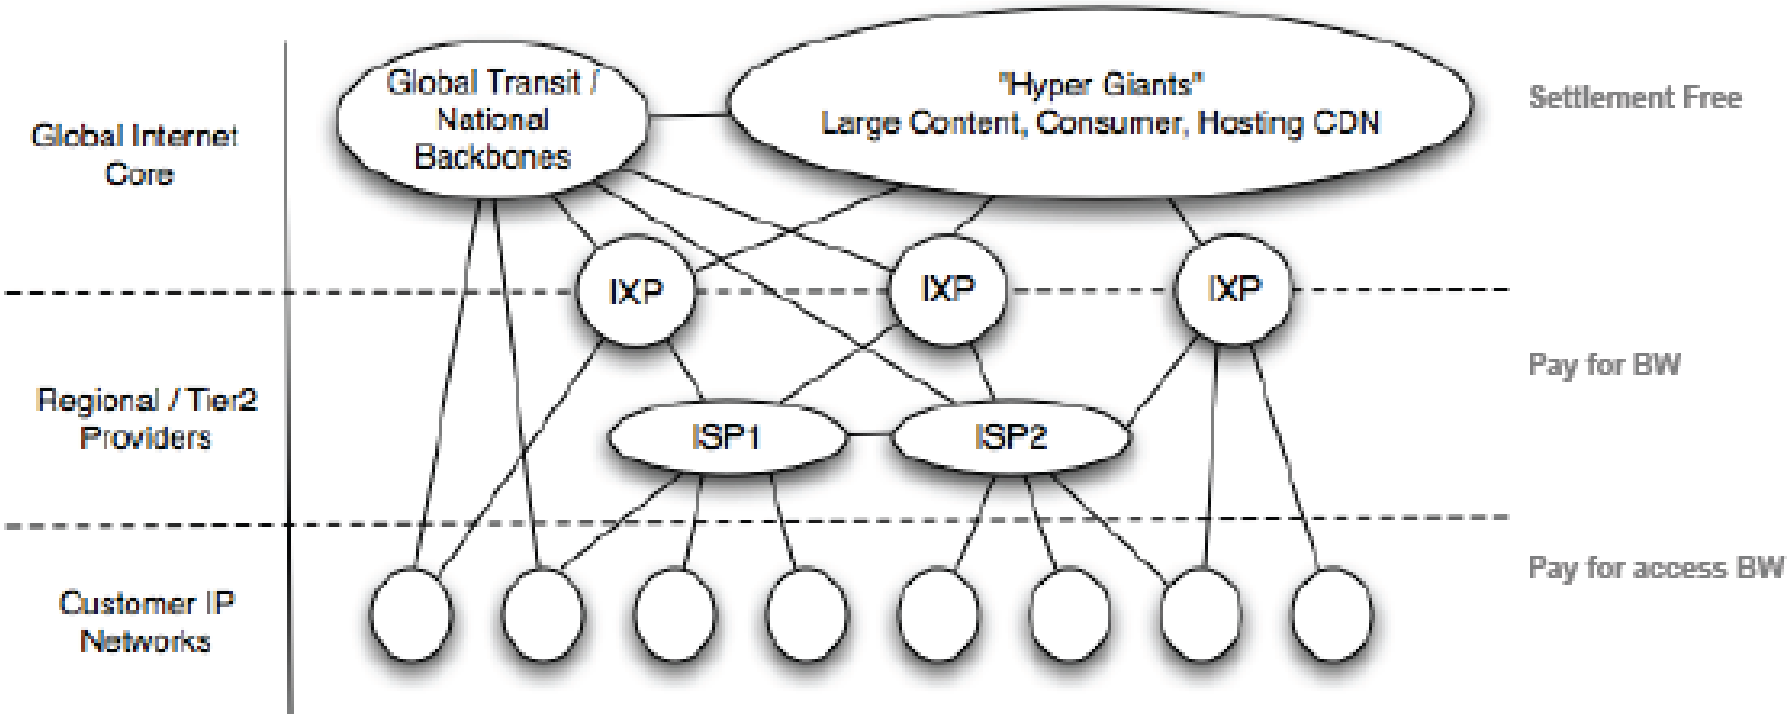

- Flatter and much more densely interconnected Internet
- Disintermediation between content and "eyeball" networks
- New commercial models between content, consumer and transit

How do we **avoid sending too much** for the receiver and network to handle?

Sliding Window Example

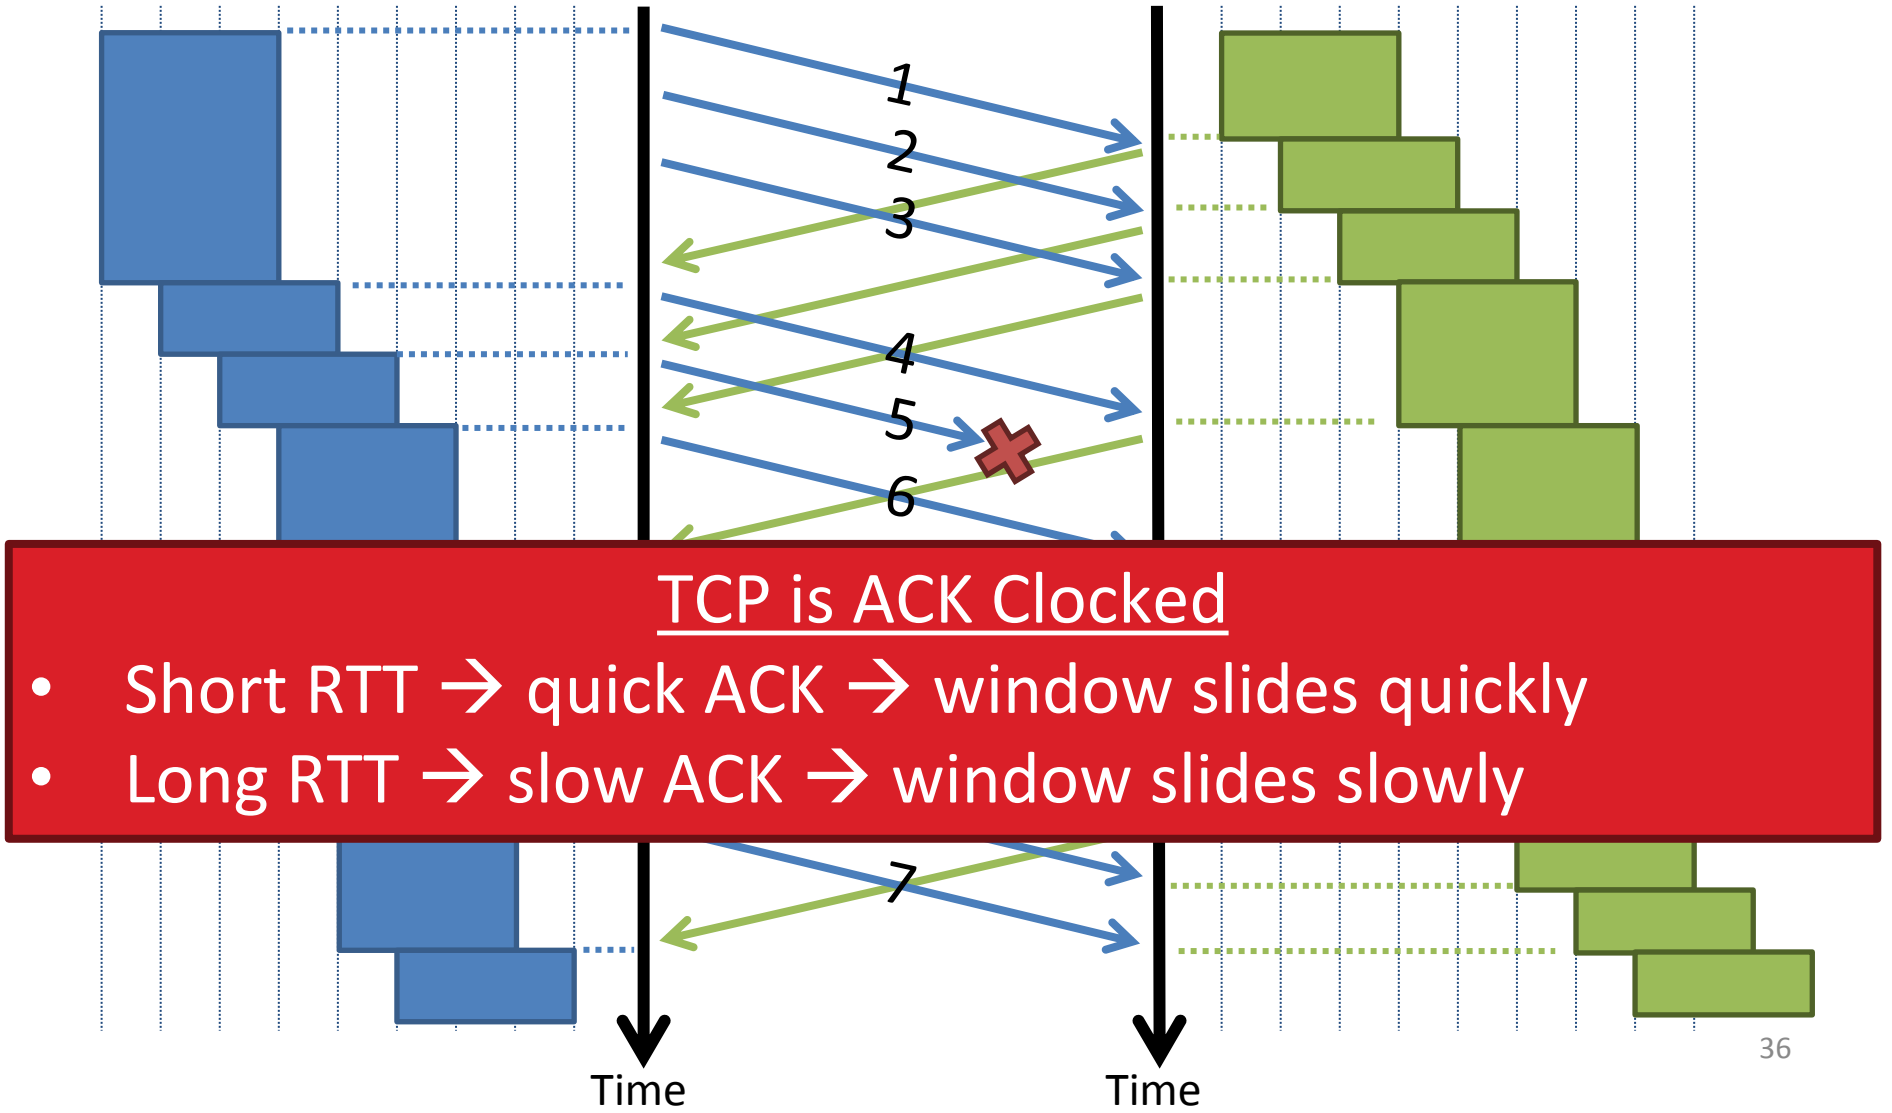

TCP is ACK Clocked

- Short RTT → quick ACK → window slides quickly
- Long RTT → slow ACK → window slides slowly

Congestion Window (*cwnd*)

- Limits how much data is in transit
- Denominated in bytes

1. *wnd* = min(*cwnd*, *adv_wnd*);
2. *effective_wnd* = *wnd* –

   (*last_byte_sent* – *last_byte_acked*);

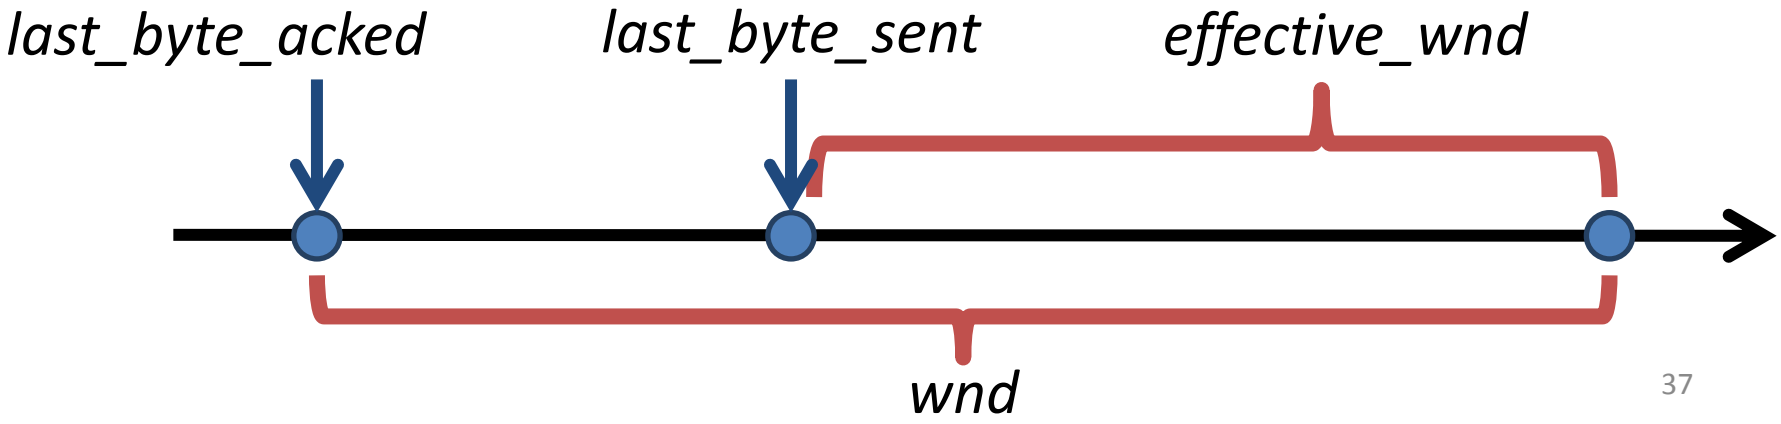

Congestion Avoidance Example

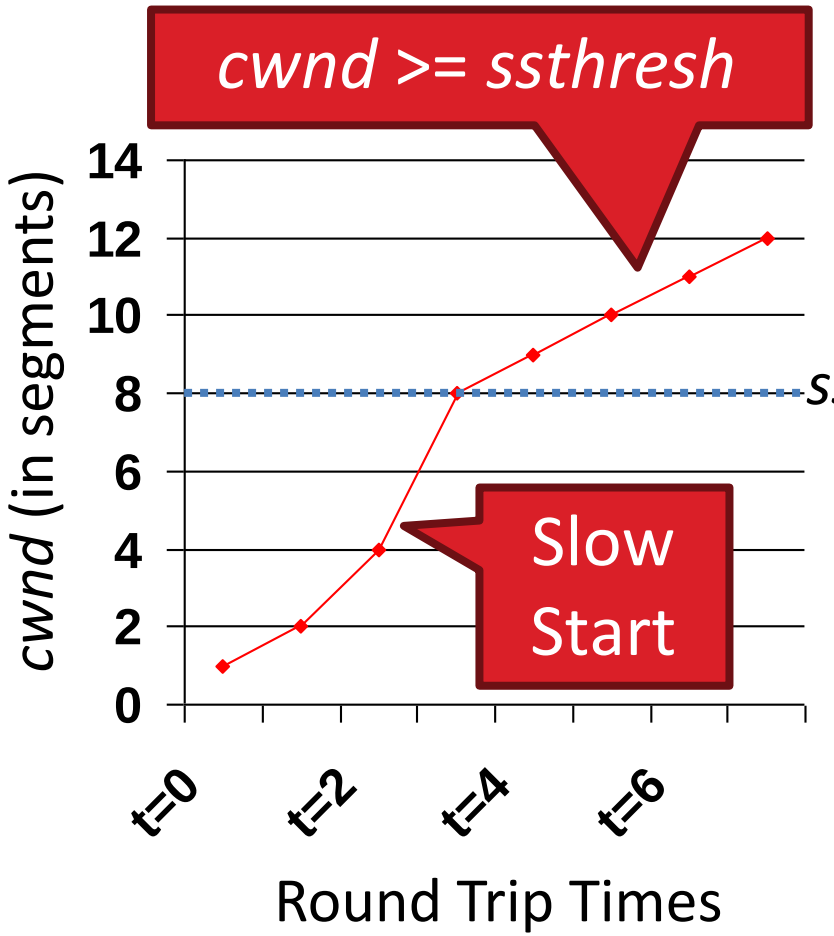

cwnd = 1

cwnd = 2

cwnd = 4

cwnd = 8

cwnd = 9

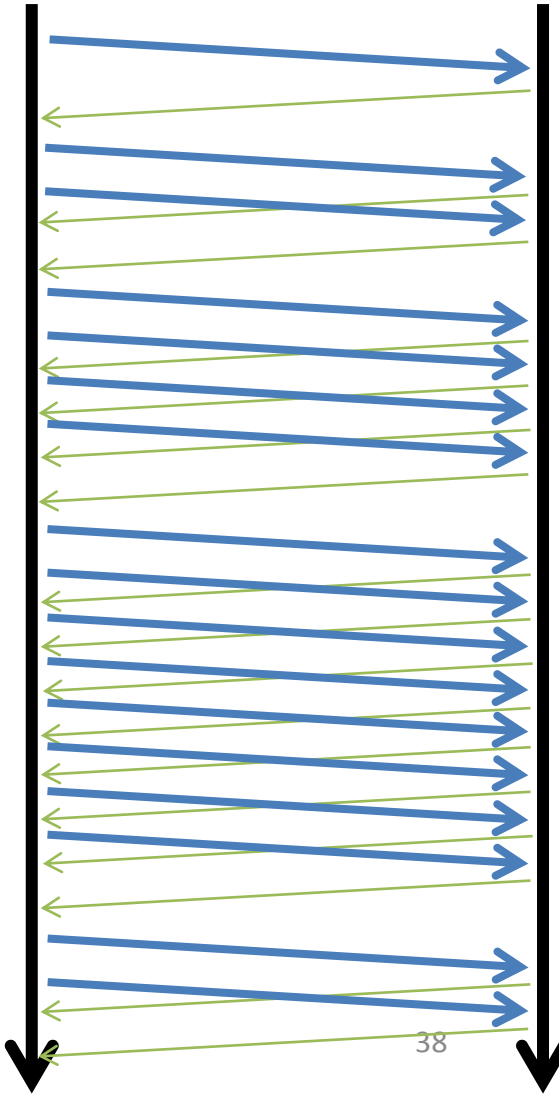

Fast Retransmit and Fast Recovery

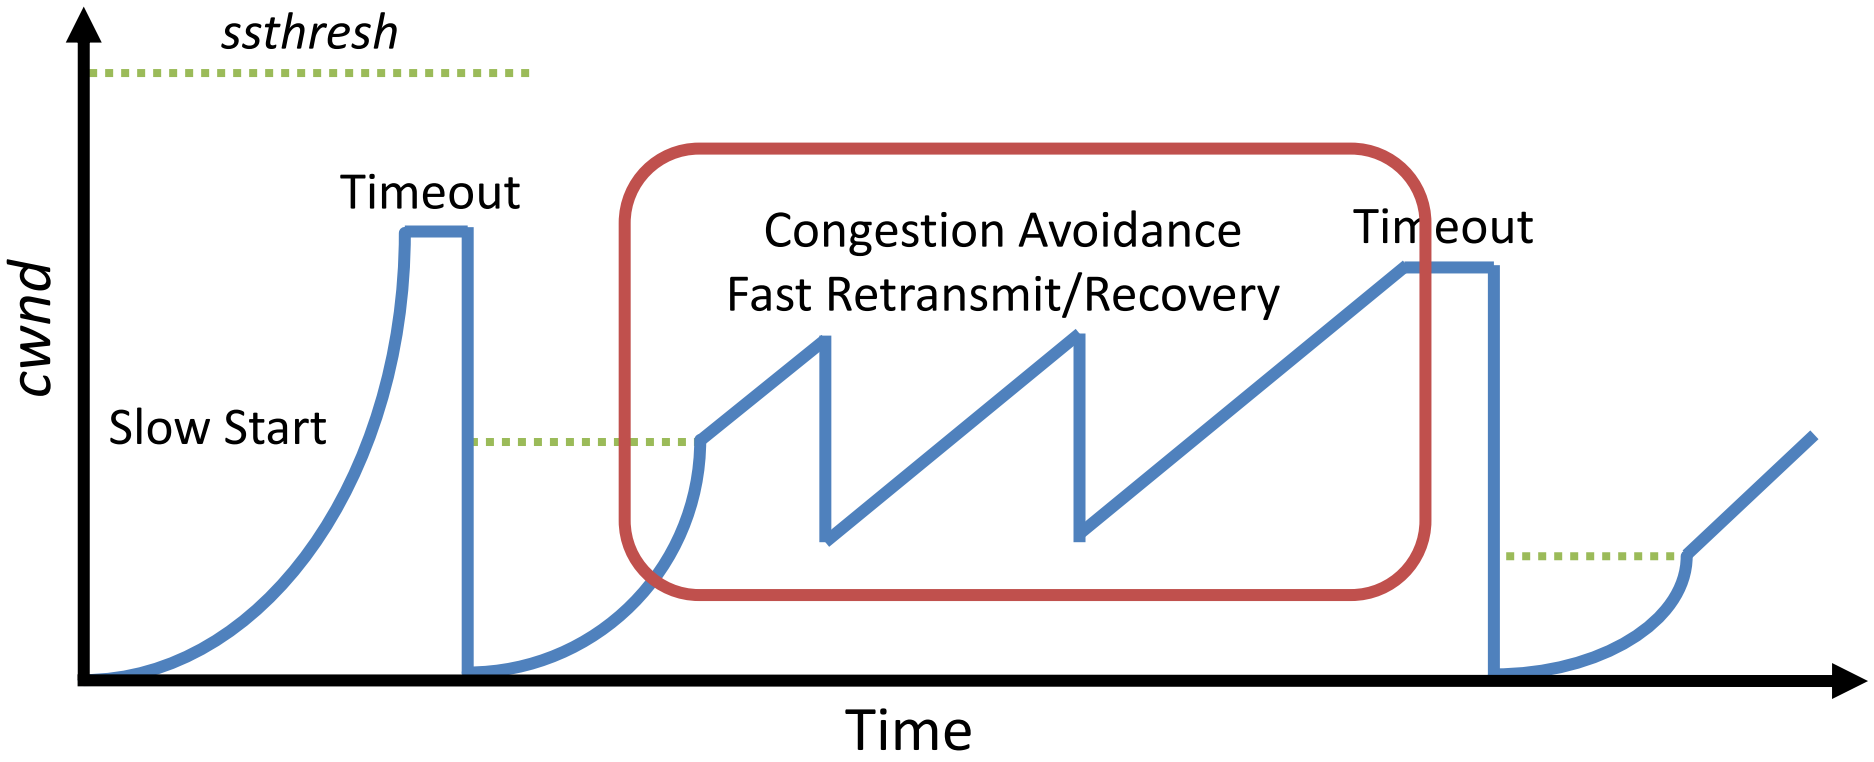

- At steady state, *cwnd* oscillates around the optimal window size
- TCP always forces packet drops

Compound TCP Example

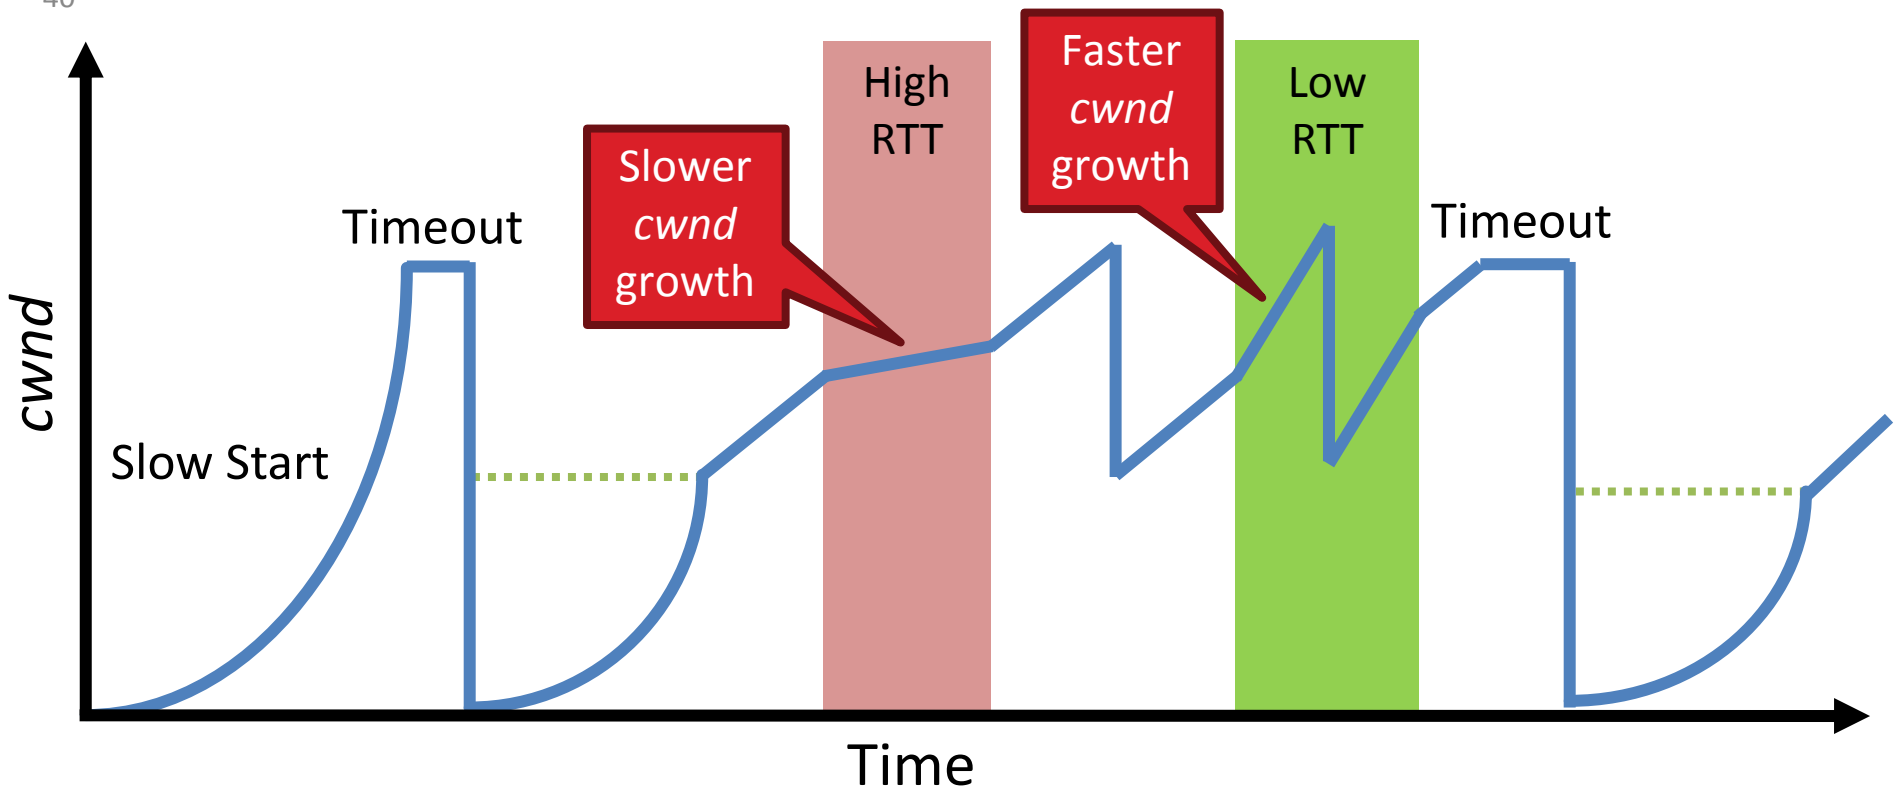

- Aggressiveness corresponds to changes in RTT
- Advantages: fast ramp up, more fair to flows with different RTTs
- Disadvantage: must estimate RTT, which is very challenging

TCP CUBIC Example

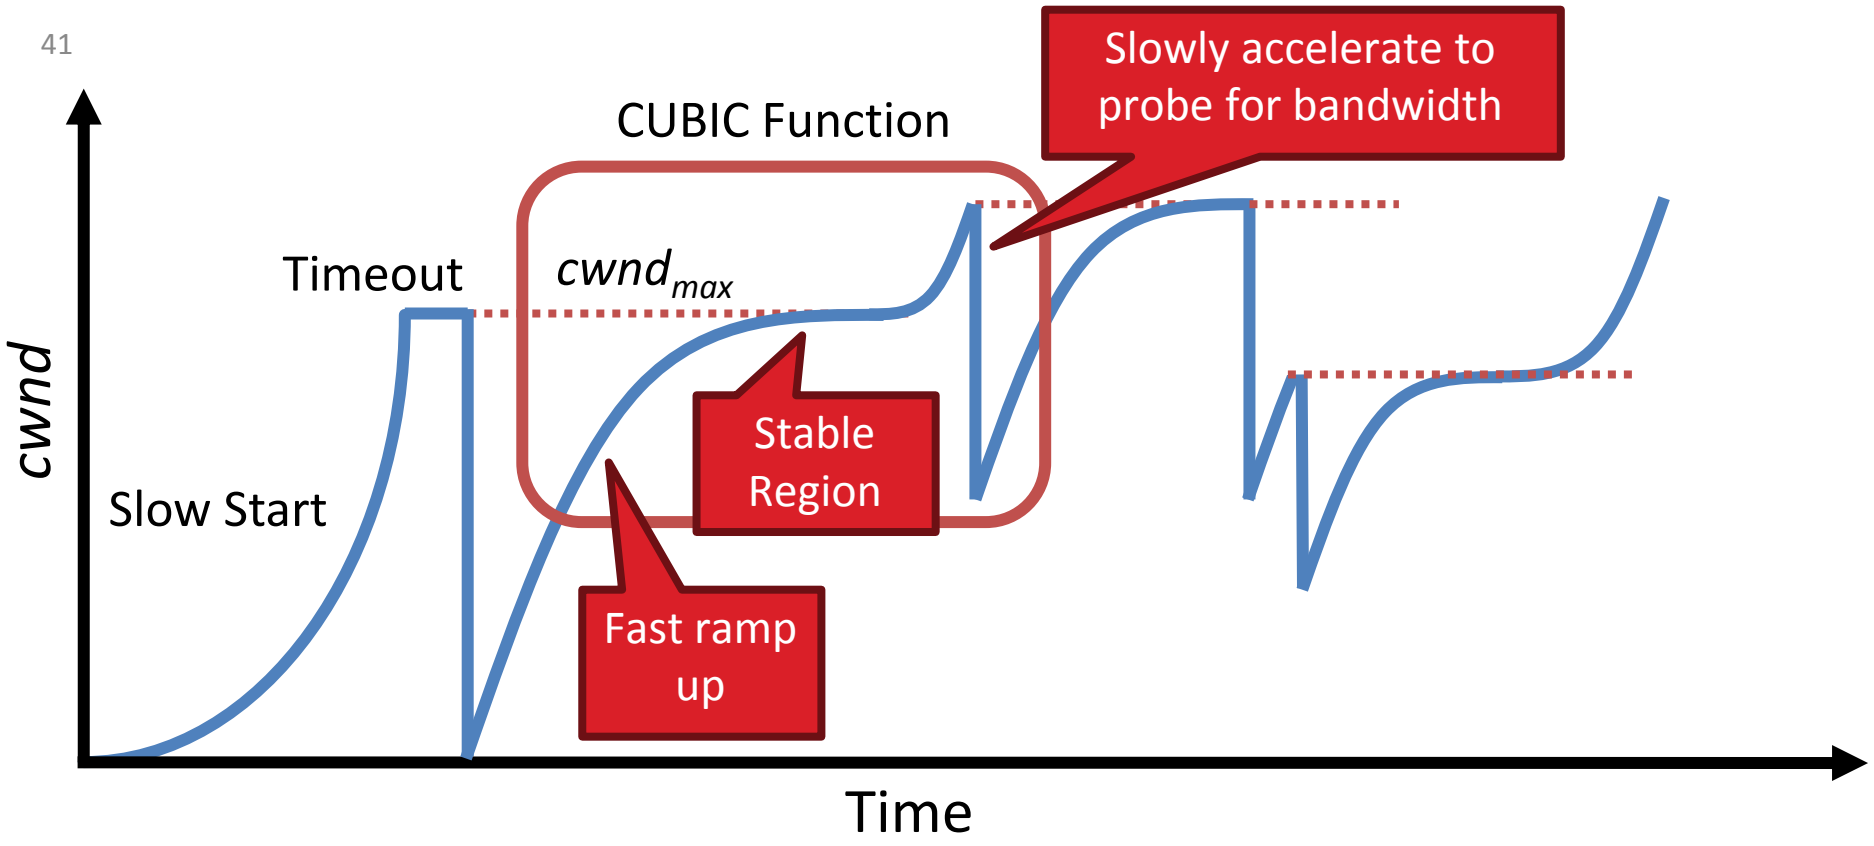

- Less wasted bandwidth due to fast ramp up
- Stable region and slow acceleration help maintain fairness
  - Fast ramp up is more aggressive than additive increase
  - To be fair to Tahoe/Reno, CUBIC needs to be less aggressive

Issues with TCP

- The vast majority of Internet traffic is TCP
- However, many issues with the protocol
  - Lack of fairness
  - Synchronization of flows
  - Poor performance with small flows
  - Really poor performance on wireless networks
  - Susceptibility to denial of service

Multipath TCP

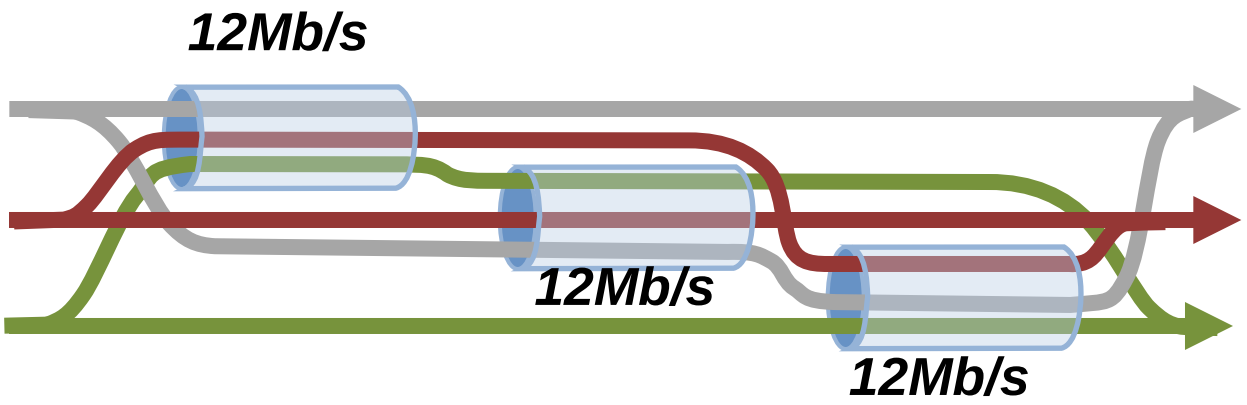

Each flow has a choice of a 1-hop and a 2-hop path.

How should split its traffic?

The Internet topology

15412	12041	p2c
15412	12486	p2c
15412	12880	p2c
15412	13810	p2c
15412	15802	p2c
15412	17408	p2c
15412	17554	p2c
15412	17709	p2c
15412	18101	p2c
15412	19806	p2c
15412	19809	p2c
15413…		

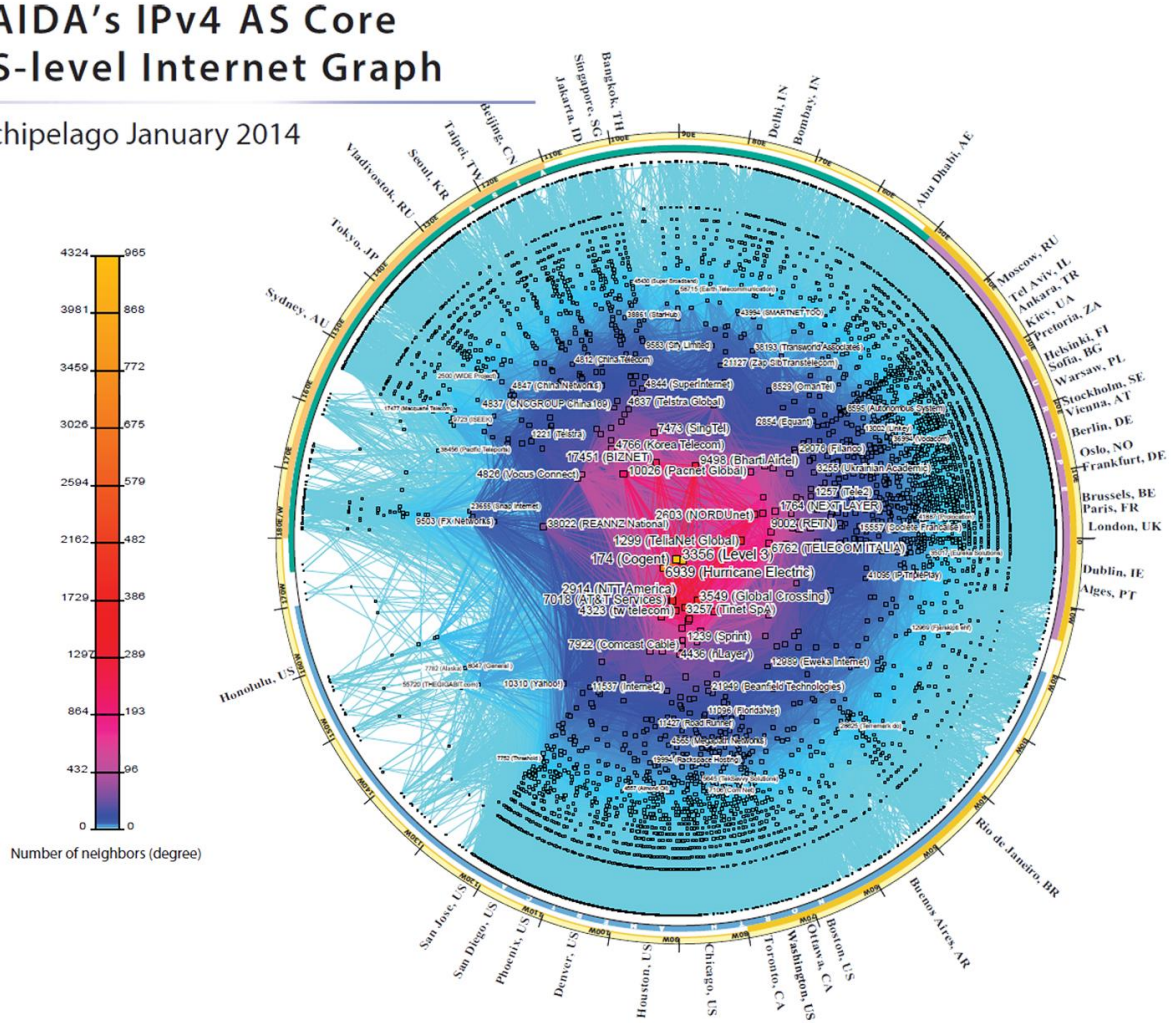

Social networks

- Social networks are graphs of people

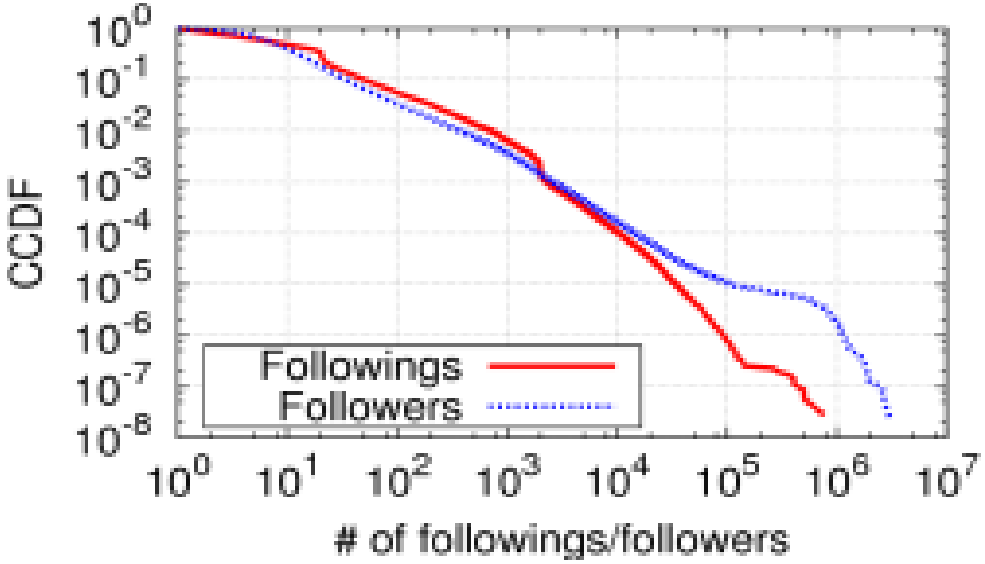

Figure 1: Number of followings and followers

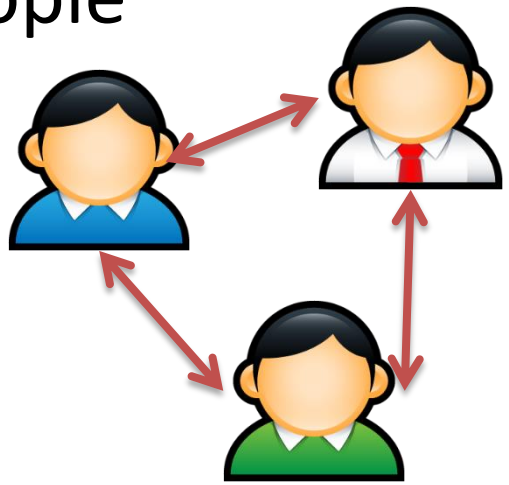

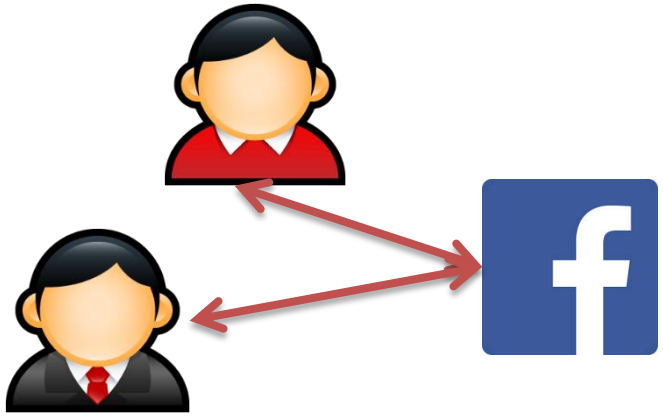

Poisson vs self similar

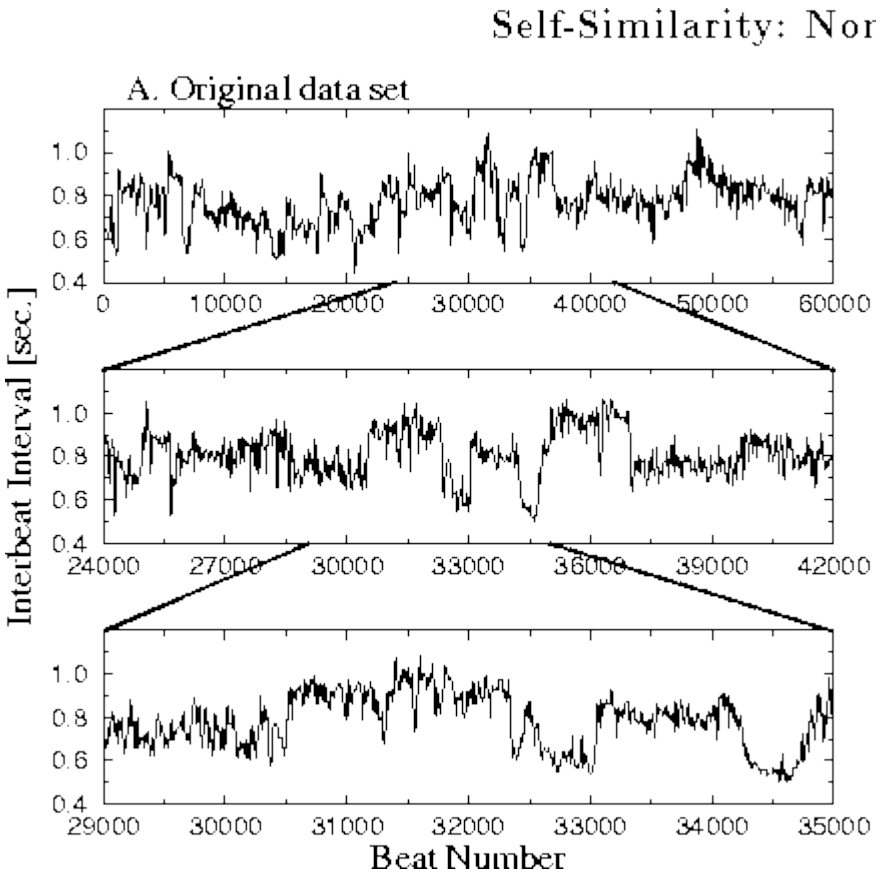

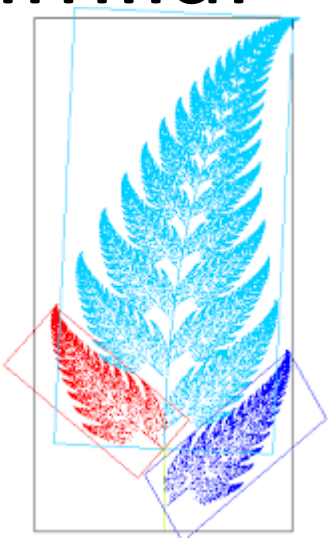

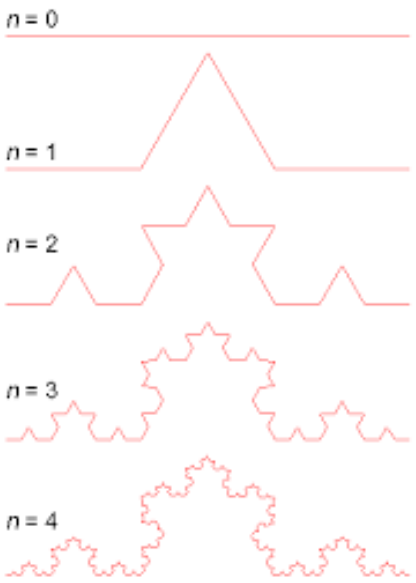

Poisson vs self similar

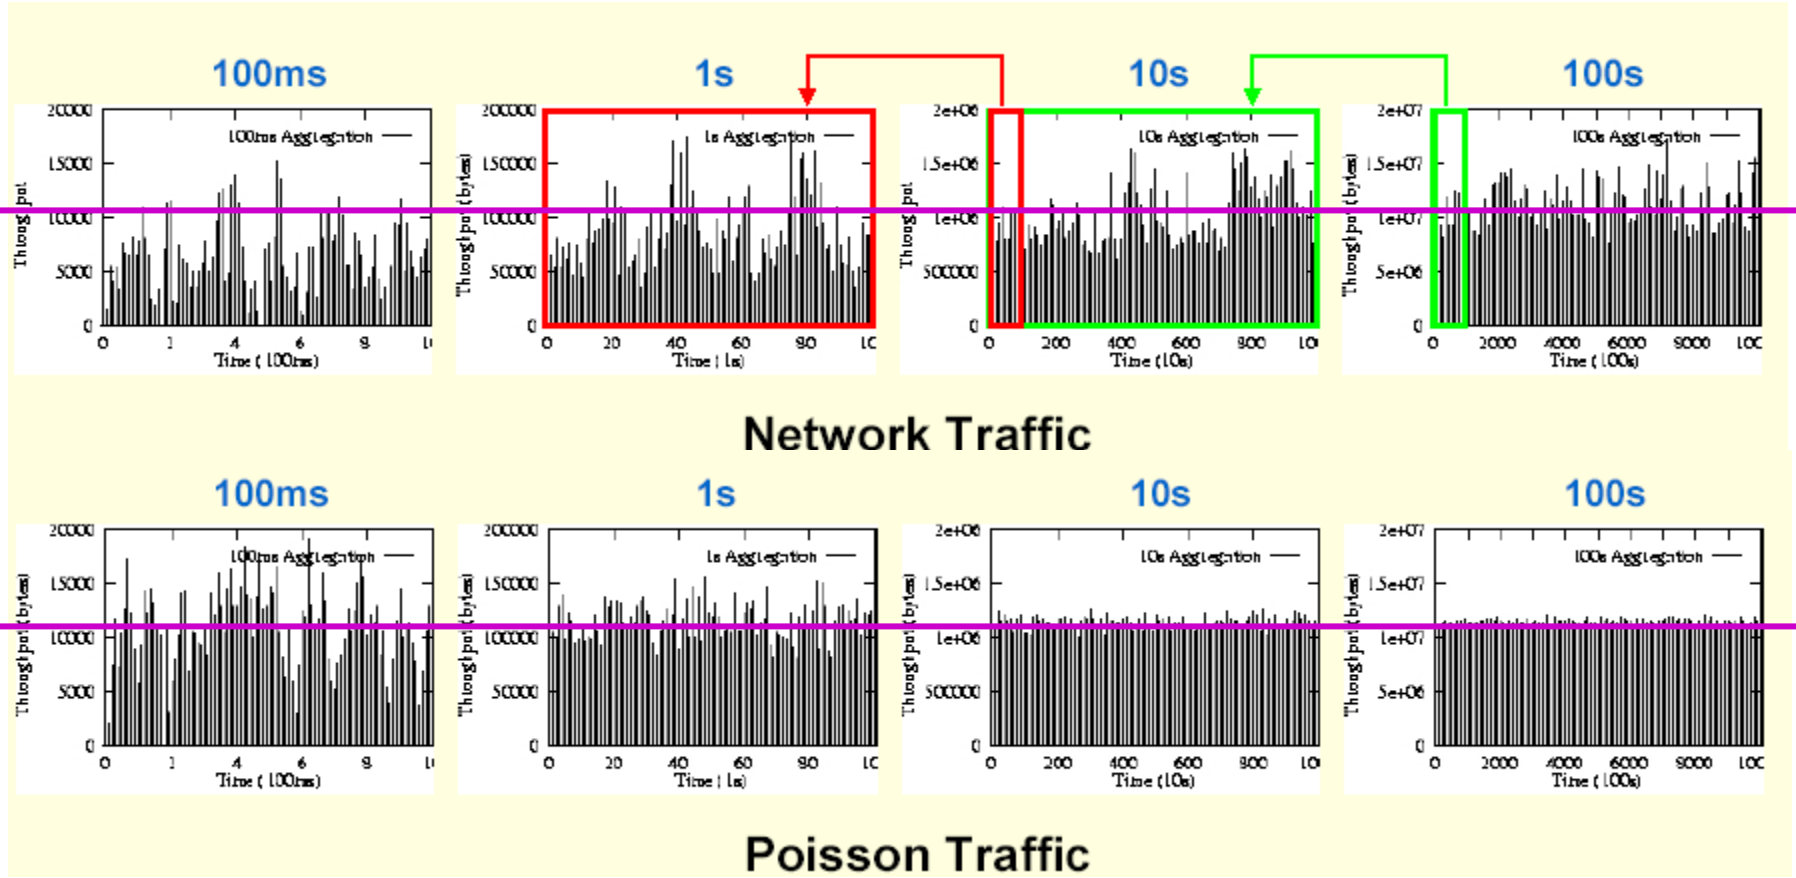

- Others have shown that traffic is non-stationary, and may well approximated as Poisson on shorter time scales

Other topics covered in class

- Web and web server loads
- Wireless performance
- HAS streaming and content popularity
- Future, Content/information centric networking, and Middleboxes
- SDN and Network virtualization

… the last topics/papers looking towards the future …

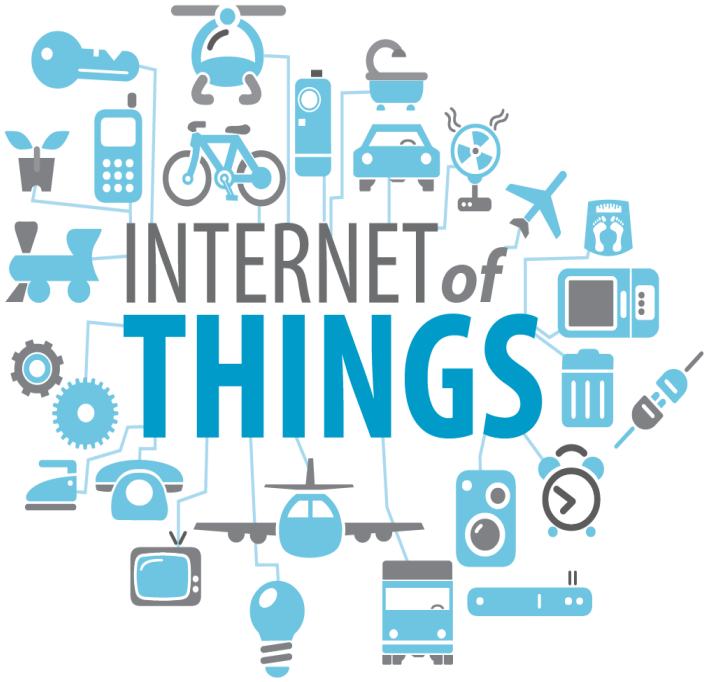

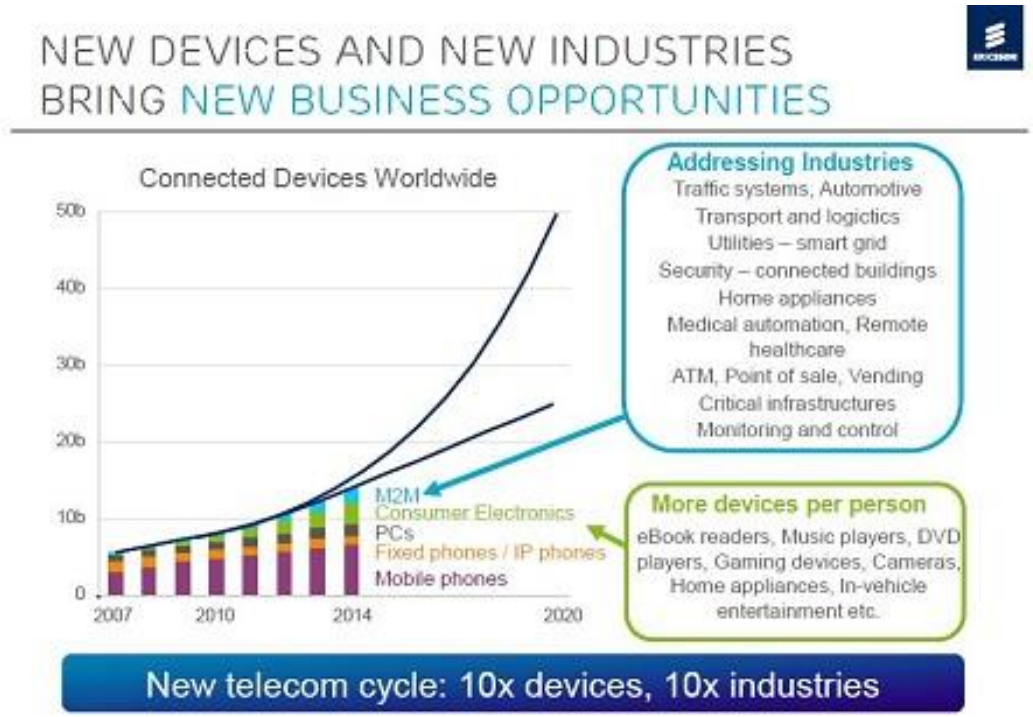

The 2020 vision

- Everything that can be connected will be connected
  - 50B devices (perhaps more like 500B …)
- IoT and smart cities
  - Machine-to-machine
- High-definition 3D streaming to heterogeneous clients

The exam

- Friday June 5, 2015
- Closed book
- Some “example” questions online
  - For this course and offering of the course, somewhat different approach …
- Bonus points from project and participation will be assigned during the exam (not before)
  - See website for details

… more exam …

- Read all instructions carefully
- Please explain how you derived your answers. Your final answers should be clearly stated (and should typically include a figure or table).
- Write answers legibly; no marks will be given for answers that cannot be read easily.
- Where a discourse or discussion is called for, be concise and precise.
- No assistance: closed book, closed notes, and no electronics ...

… yet more exam …

- If necessary, state any assumptions you made in answering a question. However, remember to read the instructions for each question carefully and answer the questions as precisely as possible. Solving the *wrong question may result in* deductions! It is better to solve the *right question incorrectly, than the wrong* question correctly.
- Please use English. (If needed, feel free to bring a dictionary from an official publisher. Hardcopy, not electronic!! Also, your dictionary is not allowed to contain any notes; only the printed text by the publisher.)