Communication Networks

Prof. Laurent Vanbever

Communication Networks

Spring 2017

Laurent Vanbever
www.vanbever.eu

ETH Zürich (D-ITET)
February, 20 2017

The Internet

An *exciting* place

8 billion

8 billion

estimated* # of Internet connected devices
in 2016

* Cisco Visual Networking Index 2017

11.6 billion

estimated* # of Internet connected devices
in **2021**

* Cisco Visual Networking Index 2017

~3 exabytes

estimated* daily global IP traffic
in 2016

* Cisco Visual Networking Index 2017

If 

= 1 Gigabyte

~3 exabytes

estimated* daily global IP traffic
in 2016

* Cisco Visual Networking Index 2017

~6 exabytes

estimated* daily global IP traffic
in 2020

* Cisco Visual Networking Index 2017

A *few* Internet services
you *might* have heard of…

The Internet
A *tense* place

Countries get disconnected

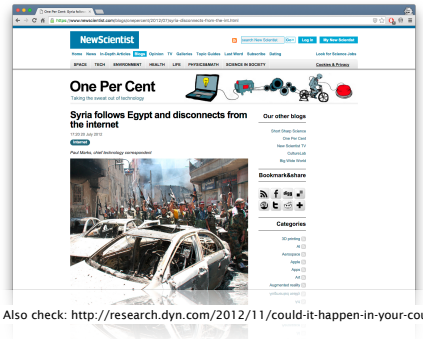

Also check: http://research.dyn.com/2012/11/could-it-happen-in-your-countr/

Internet traffic to/from Egypt in January 2011

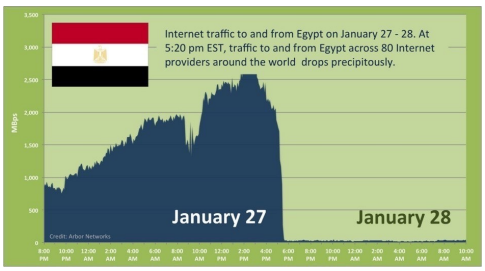

http://huff.to/1KxxoZF

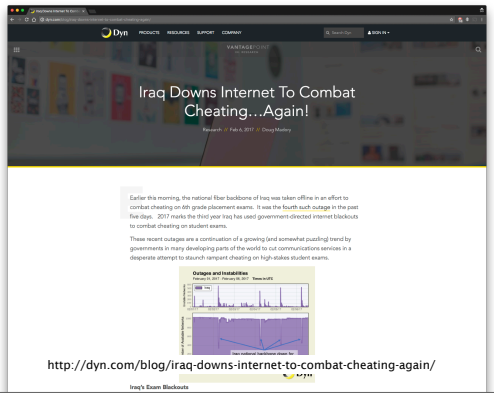

http://dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

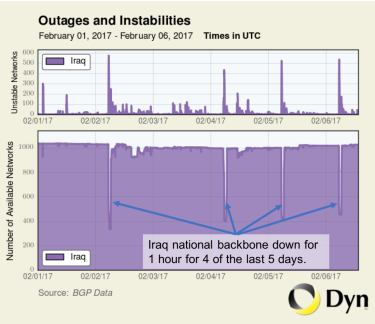

http://dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

http://www.wired.co.uk/article/over-50-internet-shutdowns-2016

Communications get eavesdropped on…

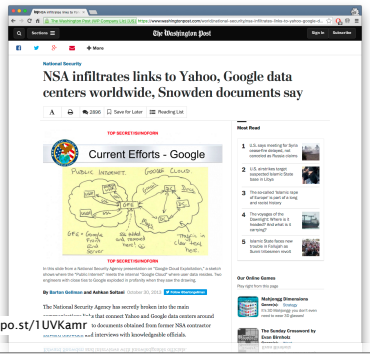

http://wapo.st/1UVKamr

http://wapo.st/1LcAw6p

Some Internet communications
are interfered against or heavily congested

Can ISPs selectively slow down traffic?

The U.S. Federal Communications Commission (FCC)
ordered Comcast to stop interfering with p2p traffic

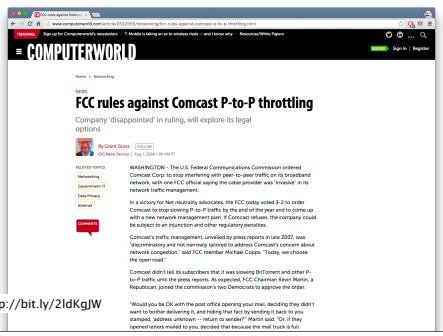

http://bit.ly/2IdKgJW

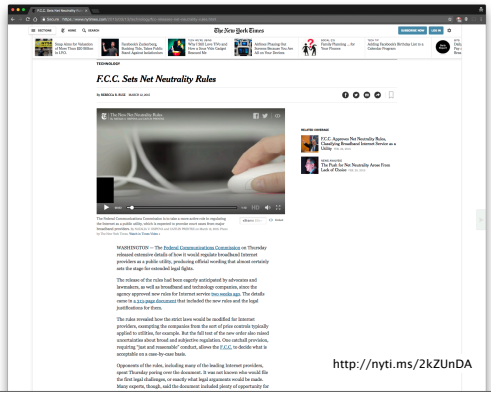

http://nyti.ms/2kZUnDA

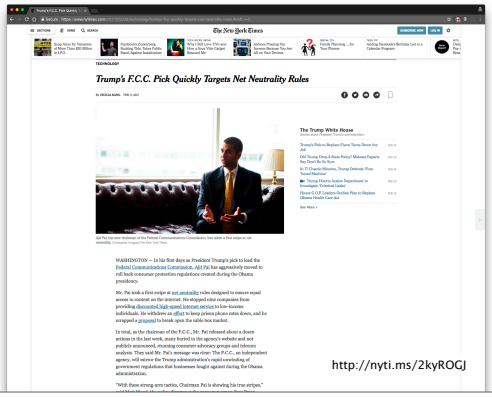

http://nyti.ms/2kyROGJ

Some Internet communications
are interfered against or heavily congested

Who should pay the other for Internet connectivity?

A primer on the conflict between
Netflix and Comcast

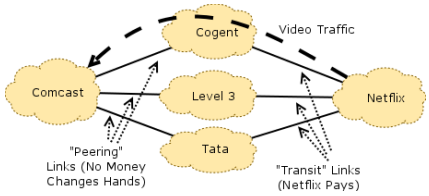

https://freedom-to-tinker.com/blog/feamster/why-your-netflix-traffic-is-slow-and-why-the-open-internet-order-wont-necessarily-make-it-faster/

Due to congestion, throughput across Cogent
to Time Warner, Comcast, Verizon were miserable

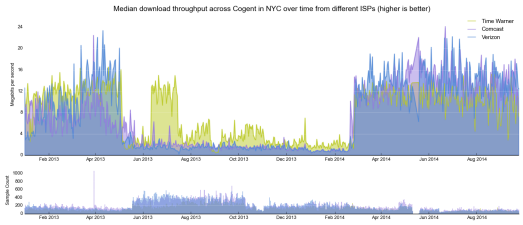

http://bit.ly/1thPzro

Situation massively improved after
Netflix agreed to paid direct connection to the providers

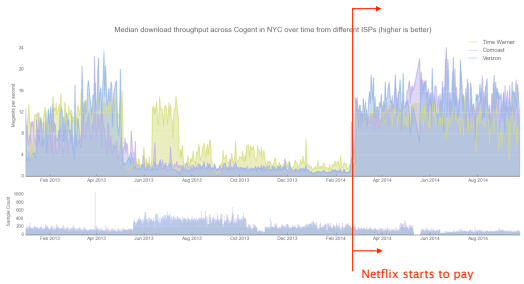

http://arstechnica.com/tech-policy/2014/04/netflix-and-verizon-reach-interconnection-deal-to-speed-up-video/

The Internet infrastructure is
a fragile environment

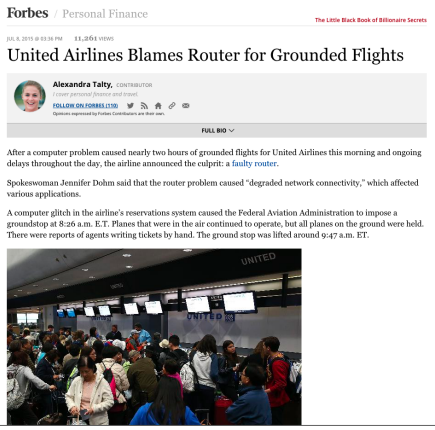

The outage was due to
one faulty Internet device

Facebook, Tinder, Instagram suffer widespread issues

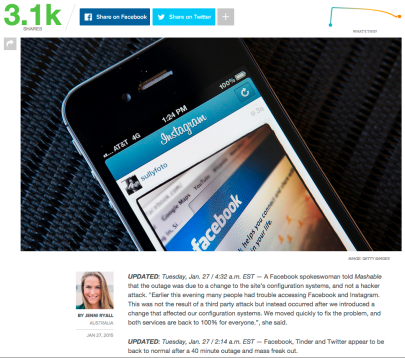

UPDATED: Tuesday, Jan. 27 / 4:32 a.m. EST — A Facebook spokeswoman told Mashable that the outage was due to a change to the site's configuration systems, and not a hacker attack. "Earlier this evening many people had trouble accessing Facebook and Instagram. This was not the result of a third party attack but instead occurred after we introduced a change that affected our configuration systems. We moved quickly to fix the problem, and both services are back to 100% for everyone.", she said.

UPDATED: Tuesday, Jan. 27 / 2:14 a.m. EST — Facebook, Tinder and Twitter appear to be back to normal after a 40 minute outage and mass freak out.

The outage was due to a change to the site's configuration systems

"Human factors are responsible for 50% to 80% of network outages"

Juniper Networks, *What's Behind Network Downtime?*, 2008

The Internet Under Crisis Conditions
Learning from September 11

Committee on the Internet Under Crisis Conditions: Learning from September 11
Computer Science and Telecommunications Board
Division on Engineering and Physical Sciences
NATIONAL RESEARCH COUNCIL

National Research Council. The Internet Under Crisis Conditions: Learning from September 11

The Internet Under Crisis Conditions
Learning from September 11

Committee on the Internet Under Crisis Conditions: Learning from September 11
Computer Science and Telecommunications Board
Division on Engineering and Physical Sciences
NATIONAL RESEARCH COUNCIL

Internet advertisements rates suggest that
The Internet was more stable than normal on Sept 11

The Internet Under Crisis Conditions
Learning from September 11

Committee on the Internet Under Crisis Conditions: Learning from September 11
Computer Science and Telecommunications Board
Division on Engineering and Physical Sciences
NATIONAL RESEARCH COUNCIL

Internet advertisements rates suggest that
The Internet was more stable than normal on Sept 11

Information suggests that operators were watching the news instead of making changes to their infrastucture

Job Snijders
@JobSnijders
Follow

Fun fact: most BGP route leaks happen on Wednesdays, but in the weekend us humans collectively take a break! :-)

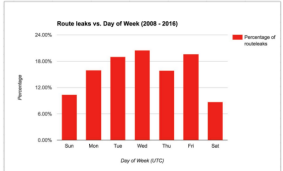

Internet scalability is at risk with
no more IPv4 addresses and a slow IPv6 deployment

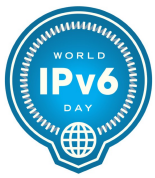

Communication Networks

Course goals

Knowledge

Understand how the Internet works and why

from your network plug…

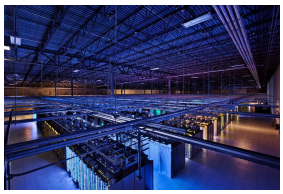

…to Google's data-center

Insight

Key concepts and problems in Networking

Naming Layering Routing Reliability Sharing

Naming Layering Routing Reliability Sharing

How do you address computers, services, protocols?

Naming Layering Routing Reliability Sharing

How do you manage complexity?

Naming Layering Routing Reliability Sharing

How do you go from A to B?

Naming Layering Routing Reliability Sharing

How do you communicate reliably using unreliable mediums?

Naming Layering Routing Reliability Sharing

How do you divide scarce resources among competing parties?

Skill

Build, operate and configure networks

Skill

Build, operate and configure networks

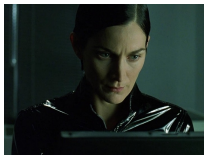

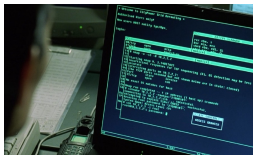

Trinity using a port scanner (`nmap`) in Matrix Reloaded™

Software-Defined Network

Software-Defined Network

enable network programmability

So far, network devices
have been completely locked down

closed software

closed hardware

Cisco™ device

SDN opens up the network devices,
enabling network innovation

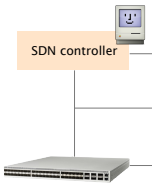

The hype around SDN is huge,
both in the industry and in academia

VMware Acquires Once-Secretive Start-Up Nicira for $1.26 Billion

JULY 23, 2012 AT 1:26 PM PT

VMware, the software company best known for its virtualization technology that forms the backbones of so-called cloud computing today, said it will pay $1.26 billion for Nicira, a networking start-up that has sought to do to networks what VMware has done to computers.

The news comes on the same day that VMware was to report quarterly earnings. And while I don't usually cover VMware's earnings, I may as well mention the results: The company reported revenue for the quarter ended June rose to $1.12 billion, while earnings on a per-share basis were 68 cents. Analysts had been expecting sales of $1.12 billion and earnings of 66 cents.

Nicira had been running in stealth mode for quite awhile; I got to reveal its plans to the world last February.

The deal amounts to a nice payoff for Nicira's investors including Andreessen Horowitz, Lightspeed Venture Partners and NEA, as well as VMware founder Diane Greene and venture capitalist Andy Rachleff.

nicira

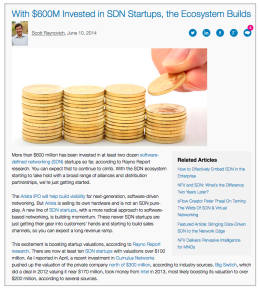

With $600M Invested in SDN Startups, the Ecosystem Builds

Scott Raynovich, June 10, 2014

More than $600 million has been invested in at least two dozen software-defined networking (SDN) startups so far, according to Nyano Report research. You can expect that to continue to climb. With the SDN ecosystem starting to take hold with a broad range of alliances and distribution partnerships, we're just getting started.

The Arista IPO will help build visibility for next-generation, software-driven networking. But Arista is selling its own hardware and is not an SDN pure-play. A new line of SDN startups, with a more radical approach to software-based networking, is building momentum. These newer SDN startups are just getting their gear into customers' hands and starting to build sales channels, so you can expect a long revenue ramp.

This excitement is boosting startup valuations, according to Reyno Report research. There are now at least ten SDN startups with valuations over $100 million. As I reported in April, a recent investment in Cumulus Networks pushed up the valuation of the private company north of $300 million, according to industry sources. Big Switch, which did a deal in 2012 valuing it near $170 million, took money from Intel in 2013, most likely boosting its valuation to over $200 million, according to several sources.

Related Articles

How to Effectively Embed SDN in the Enterprise

NFV and SDN: What's the Difference Two Years Later?

#Flow Creator Peter Phaal On Taming The Wilds Of SDN & Virtual Networking

Featured Article: Bringing Data-Driven SDN to the Network Edge

NFV Delivers Pervasive Intelligence for MNOs

Communication Networks

Course organization

Our website: http://comm-net.ethz.ch
check it out regularly

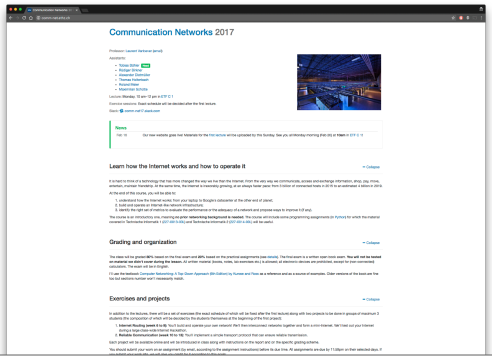

The course will be split into three parts

Part 1	Part 2	Part 3
Overview	Concepts	Today's Internet
2 lectures	2 lectures	~9 lectures

The lectures will be accompanied by exercises,
there will be two sessions per week (to ease scheduling)

available slots	Tue	8-10 10-12 13-15 15-17	vote ***today*** on comm-net.ethz
	Wed	13-15 15-17	
	Thur	13-15 15-17	

Your final grade

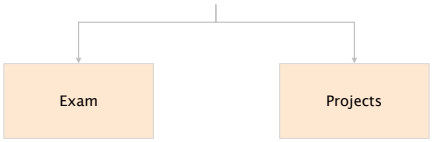

Exam	Projects
80%	20%
written, open book	mandatory

Your final grade

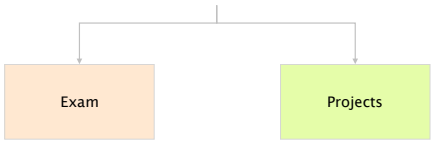

Exam	Projects
80%	20%
written, open book	mandatory

There will be two practical projects,
to be done in group of maximum three students

#1 Build and operate a small Internet

#2 Implement an interoperable reliable protocol

Detailed instructions will follow

"Internet Hackathon"
sometime in week 8

2016 edition

The course follows the textbook
Computer Networking: a Top-Down Approach

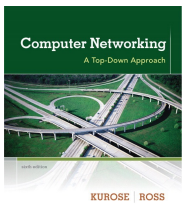

6th edition
ok to use the 5th

see sections indicated
on comm-net.ethz.ch

We'll use Slack (chat client)
to discuss about the course and assignments

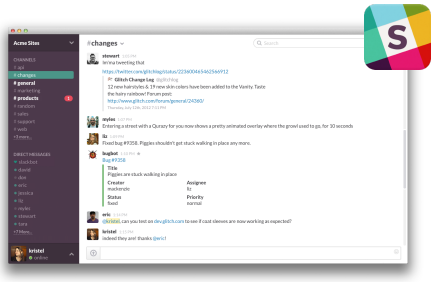

Web, smartphone and desktop clients available

Register today using your *real* name
> https://comm-net17.slack.com/signup

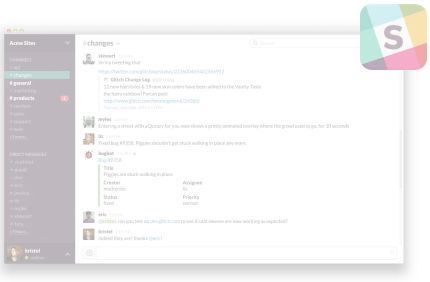

Web, smartphone and desktop clients available

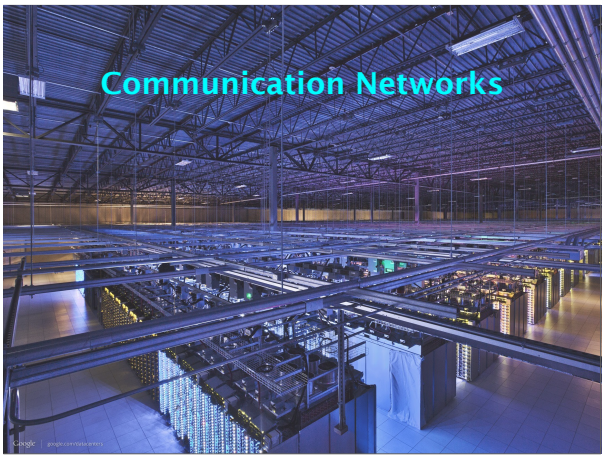

List any
technologies, principles, applications…
used after typing in:

> www.google.ch

and pressing enter in your browser

Communication Networks
Part 1: Overview

#1 What is a network made of?
#2 How is it shared?
#3 How is it organized?
#4 How does communication happen?
#5 How do we characterize it?

Communication Networks
Part 1: Overview

#1 What is a network made of?
How is it shared?
How is it organized?
How does communication happen?
How do we characterize it?

Networks are composed of three basic components

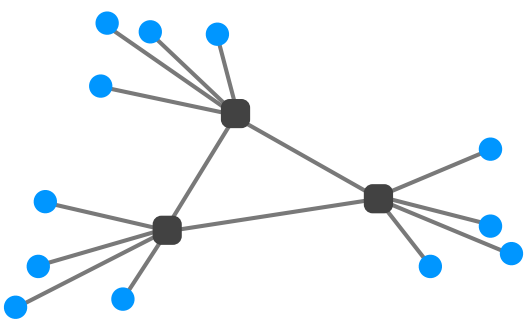

End-systems send & receive data

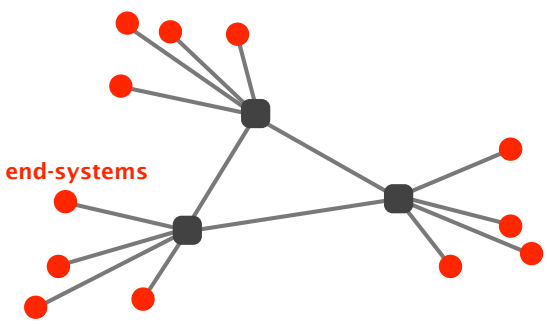

End-systems come in a wide-variety

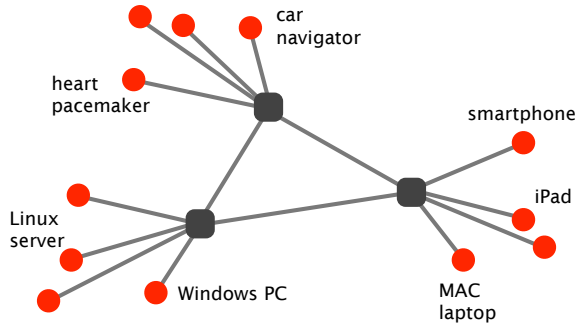

Switches & routers forward data to the destination

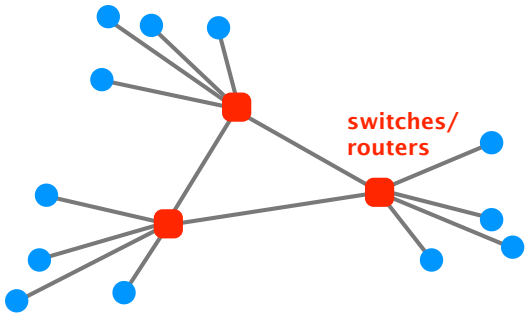

Routers vary in size and usage

Home router

~20 cm
0,5 kg
1 Gbps

Internet core router

>200cm
700kg
1.2 Tbps

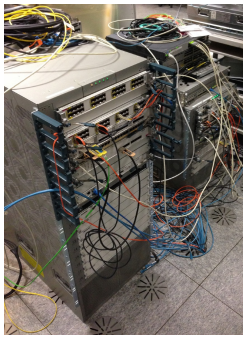

Cisco Nexus 7k
Routers @ETHZ

~25 deployed

Links connect end-systems to switches and switches to each other

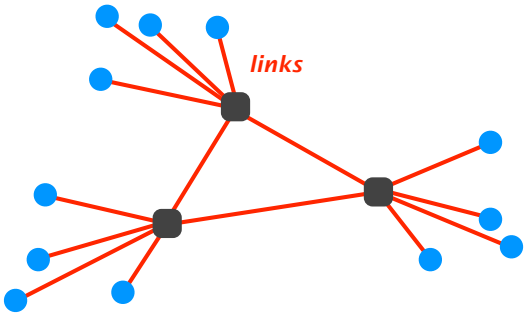

Links, too, vary in size and usage

Copper
ADSL, RJ-45,...

Optical fibers

Wireless link

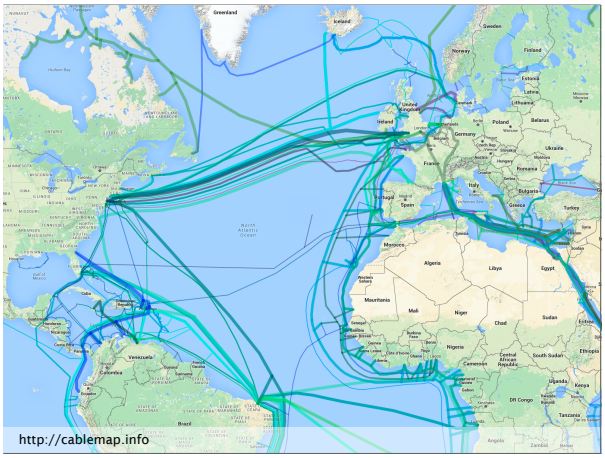

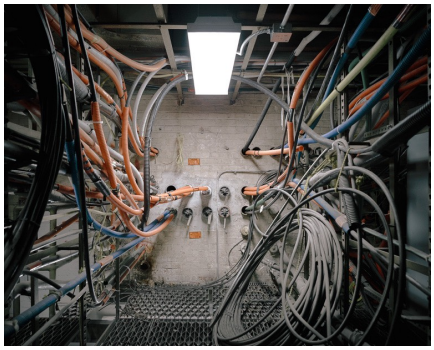

Somewhere in Manhattan... http://www.petergarritano.com/the-internet.html

The *Inter*net is a network of networks

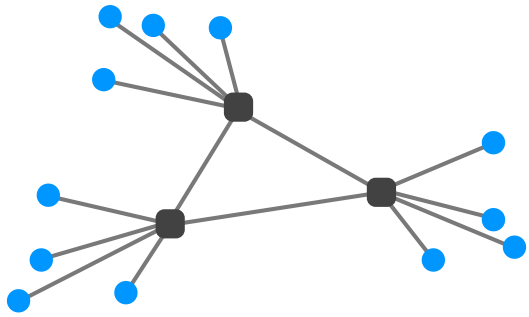

Internet Service Providers

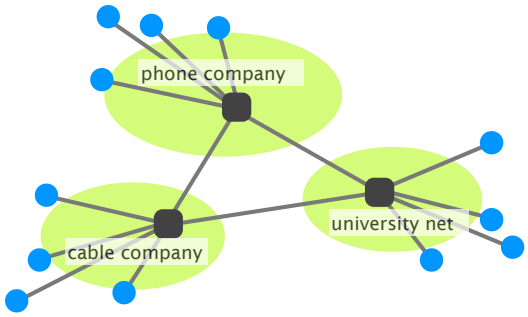

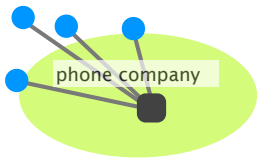

Conceptually, the last mile of the Internet looks like this

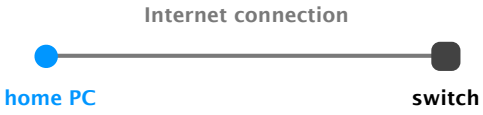

In practice, it looks more like this…

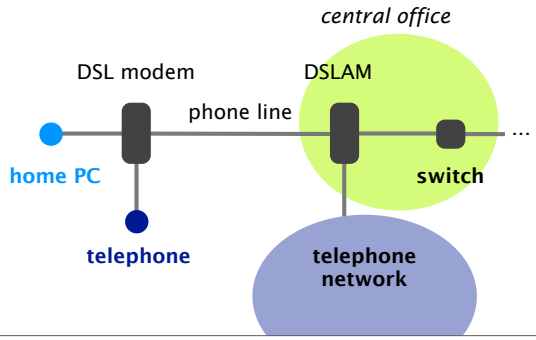

Digital Subscriber Line (DSL) brings high BW to households over phone lines

Digital Subscriber Line (DSL) brings high BW to households over **phone lines**

↓

Why?

Digital Subscriber Line (DSL) brings high BW to households over phone lines

DSL is composed of 3 channels:

- downstream data channel — tens to few hundred Mbps
- upstream data channel — few Mbps to few tens Mpbs
- 2-ways phone channel

DSL is composed of 3 channels:

- downstream data channel — tens to few hundred Mbps
- upstream data channel — few Mbps to few tens Mpbs
- 2-ways phone channel

↓

Why is there such an asymmetry?

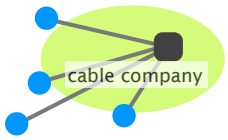

Conceptually, the last mile of the Internet looks like this

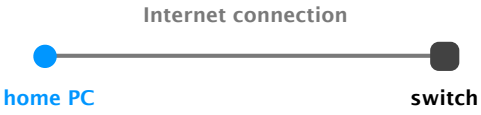

In practice, it looks more like this…

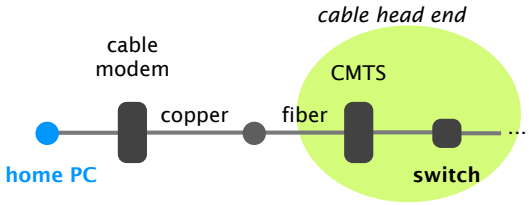

Many households share the same access

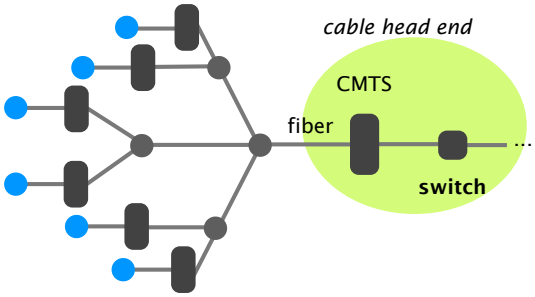

Cable Access Technologies (CATV) brings high BW to the households via cable TV

coaxial copper & fiber

- downstream data channel — tends to hundreds of Mbps
- upstream data channel — tens of Mbps

Unlike ADSL, the medium is shared between households

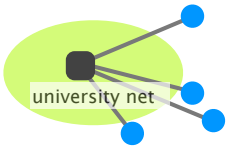

With respect to DSL and cable providers,
enterprise access networks is *much* simpler

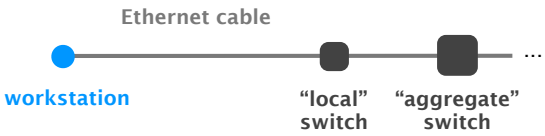

Ethernet is the most widely used
Local Area Network technology

Twisted pair copper

1Gbps, 10 Gbps, 40 Gbps, 100 Gbps, ...
symmetric

ADSL, CATV and Ethernet are only few examples
of access technologies...

Cellular	smart phones
Satellite	remote areas
FTTH	household
Fibers	Internet backbone
Infiniband	High performance computing
	...

Communication Networks

Part 1: Overview

ETH

What is a network made of?

#2 How is it shared?

How is it organized?

How does communication happen?

How do we characterize it?

Up to now, we've seen
what the last mile of the Internet looks like

What about the rest of the network?

3 requirements for a network topology

Should tolerate failures
several paths between each source and destination

Possess enough sharing to be feasible & cost-effective
number of links should not be too high

Provide adequate per-node capacity
number of links should not be too small

Compare these three designs in terms of
sharing, resiliency, and per-node capacity

design	full-mesh	chain	bus
	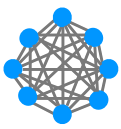	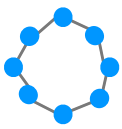	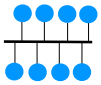
advantages			
disadvantages			

Switched networks provide reasonable and flexible compromise

design	switched

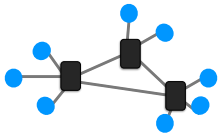

advantages	sharing and per-node capacity can be adapted to fit the network needs
disadvantages	require smart devices to perform: forwarding, routing, resource allocation

Links and switches are shared between flows

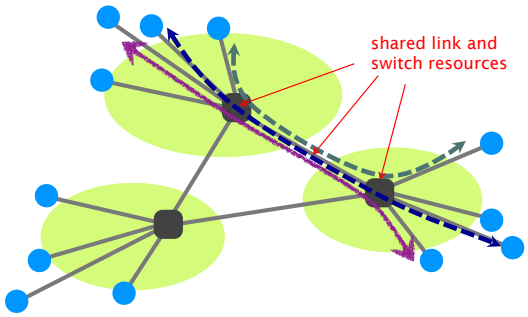

There exist two approaches to sharing: reservation and on-demand

	Reservation	On-demand
principle	reserve the bandwidth you need in advance	send data when you need

Both are examples of statistical multiplexing

	Reservation	On-demand
multiplexing	at the flow-level	at the packet-level

Between reservation and on-demand: Which one do you pick?

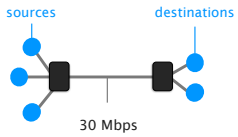

Consider that each source needs 10 Mbps

What do they get with:

- reservation
- on-demand

Assume the following peak demand and flow duration

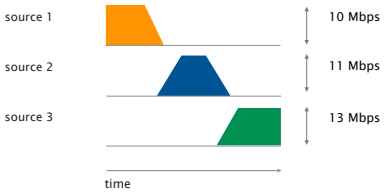

Assume the following peak demand and flow duration

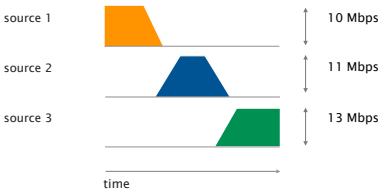

What does each source get with reservation and on-demand?

- first-come first-served
- equal (10 Mbps)

Peak vs average rates

Each flow has Peak rate P
Average rate A

Reservation must reserve *P*, but level of utilization is *A/P*
e.g., *P*=100 Mbps, *A*=10 Mbps, level of utilization=10%

On-demand can achieve higher level of utilizations
depends on degree of sharing and burstiness of flows

Ultimately, it depends on the application

Reservation makes sense when P/A is small
voice traffic has a ratio of 3 or so

Reservation wastes capacity when P/A is big
data applications are bursty, ratios >100 are common

Ultimately, it depends on the application

Reservation makes sense when P/A is small
voice traffic has a ratio of 3 or so

Reservation wastes capacity when P/A is big
data applications are bursty, ratios >100 are common

That's why the phone network used reservations
... and why the Internet does not!

In practice, the two approaches are implemented using circuit-switching or packet-switching, respectively

	Reservation	On-demand
implem.	circuit-switching	packet-switching

	Reservation	On-demand
implem.	circuit-switching	packet-switching

Circuit switching relies on the Resource Reservation Protocol

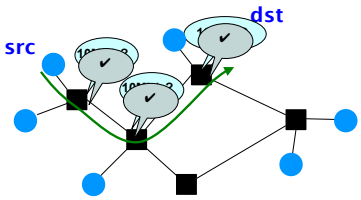

(1) src sends a reservation request for 10Mbps to dst
(2) switches "establish a circuit"
(3) src starts sending data
(4) src sends a "teardown circuit" message

The Resource Reservation Protocol establishes a circuit within each switch

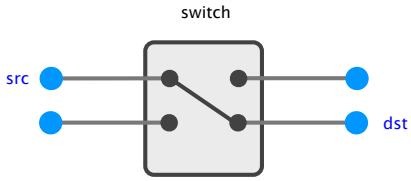

There exist many kinds of circuits

Time-based multiplexing

- divide time in slots
- allocate one slot per circuit

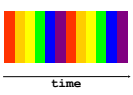

Frequency-based multiplexing

- divide spectrum in frequency bands
- allocate one band per circuit

Let's walk through example of data transfer using circuit switching

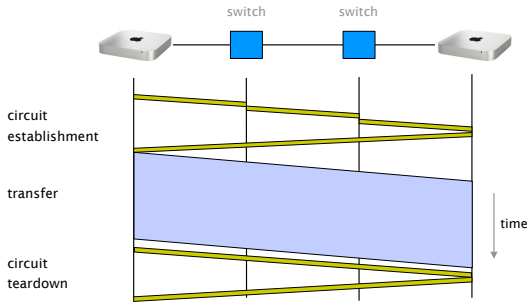

The efficiency of the transfer depends on
how utilized the circuit is once established

This is an example of poor efficiency.
The circuit is mostly idle due to traffic bursts

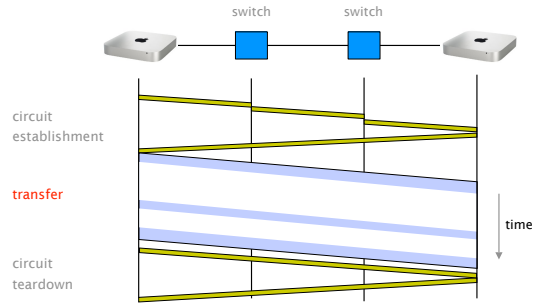

This is another example of poor efficiency.
The circuit is used for a short amount of time

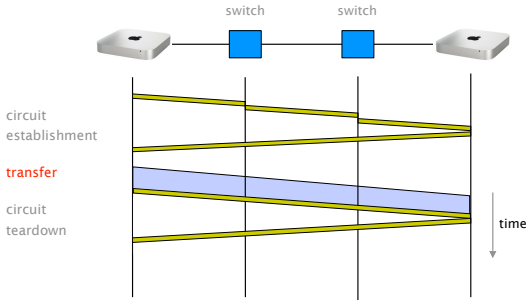

Another problem of circuit switching is that
it doesn't route around trouble

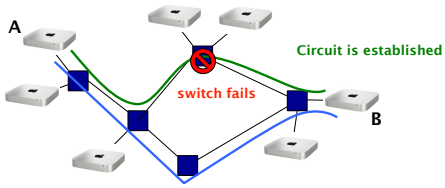

A is forced to signal a new circuit to
restore communication

Pros and cons of circuit switching

advantages	disadvantages
predictable performance	inefficient if traffic is bursty or short
simple & fast switching once circuit established	complex circuit setup/teardown which adds delays to transfer
	requires new circuit upon failure

What about packet switching?

Reservation	On-demand
circuit-switching	packet-switching

In packet switching,
data transfer is done using independent packets

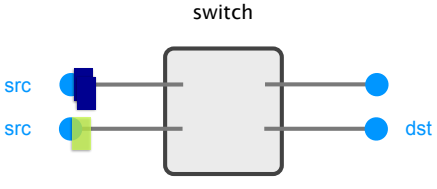

Each packet contains destination (dst)

Since packets are sent without global coordination,
they can "clash" with each other

To absorb transient overload,
packet switching relies on buffers

To absorb transient overload,
packet switching relies on buffers

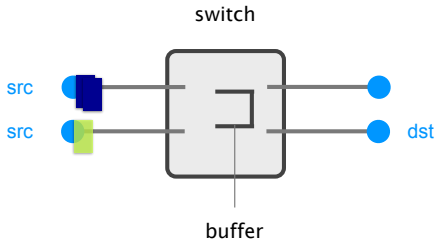

Packet switching routes around trouble

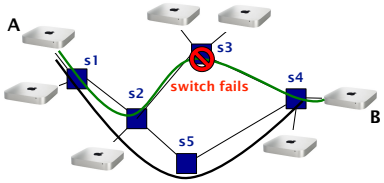

route recomputed
on the fly by s2

Pros and cons of packet switching

advantages	disadvantages
efficient use of resources	unpredictable performance
simpler to implement than circuit switching	requires buffer management and congestion control
route around trouble	

Packet switching beats Circuit switching
with respect to *resiliency* and *efficiency*

Internet ♥ packets

Packet switching will be our focus for the rest of the course

Communication Networks

Part 1: Overview

ETH

What is a network made of?

How is it shared?

#3 How is it organized?

How does communication happen?

How do we characterize it?

So far, this is our vision of the Internet…

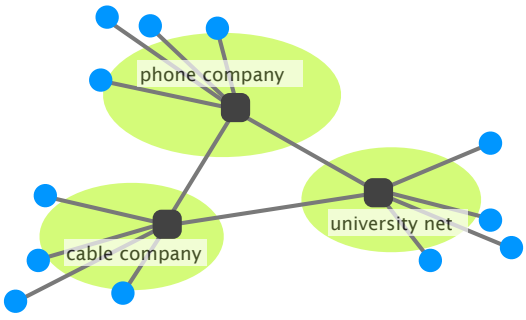

In practice, the Internet is a "bit" more complex

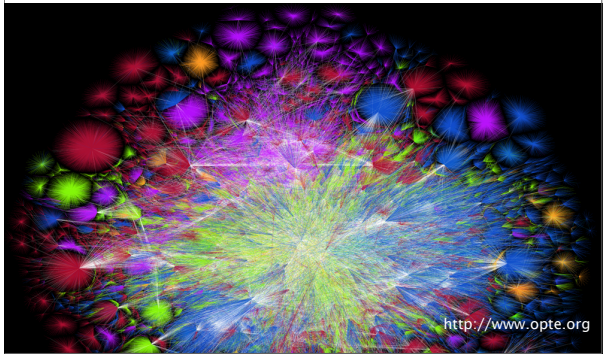

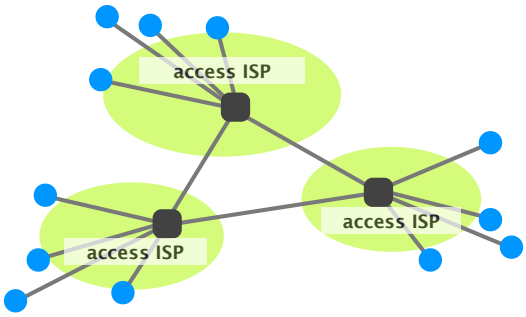

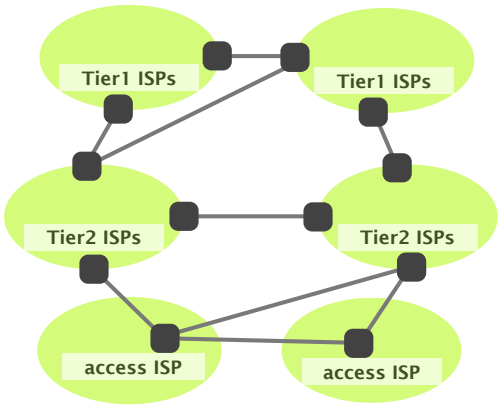

The Internet has a hierarchical structure

Tier-1 international	have no provider
Tier-2 national	provide transit to Tier-3s have at least one provider
Tier-3 local	do not provide any transit have at least one provider

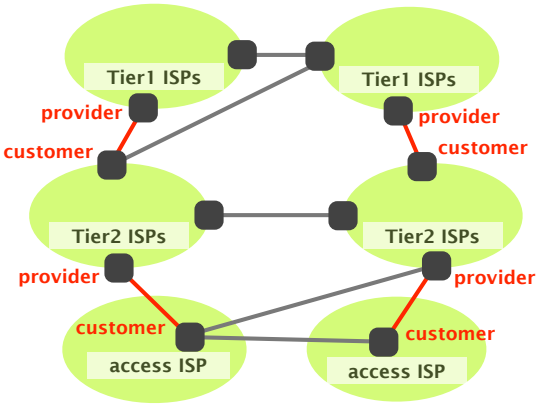

The distribution of networks in Tiers
is extremely skewed towards Tier-3s

	total	~50,000 networks
Tier-1 international	have no provider	~12
Tier-2 national	provide transit to Tier-3s have at least one provider	1,000s
Tier-3 local	do not provide any transit have at least one provider	85-90%

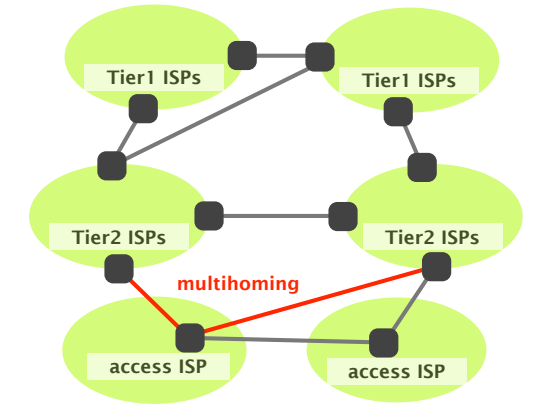

Some networks have an incentive to connect directly,
to reduce their bill wrt their own provider

This is known as "peering"

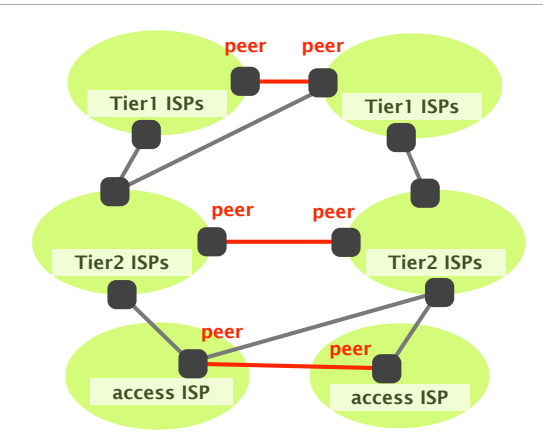

Interconnecting each network to its neighbors one-by-one is not cost effective

Physical costs
of provisioning or renting physical links

Bandwidth costs
a lot of links are not necessarily fully utilized

Human costs
to manage each connection individually

Internet eXchange Points (IXPs) solve these problems by letting *many* networks connect in one location

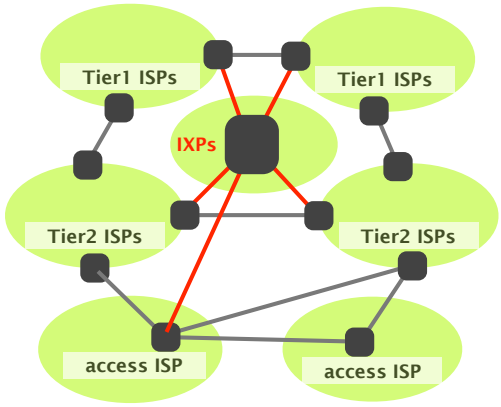

Communication Networks

Part 1: Overview

#1	What is a network made of?
#2	How is it shared?
#3	How is it organized?
#4	How does communication happen?
#5	How do we characterize it?