Communication Networks

Prof. Laurent Vanbever

Communication Networks

Spring 2019

Laurent Vanbever
nsg.ee.ethz.ch

ETH Zürich
February 18 2019

Materials inspired from Scott Shenker & Jennifer Rexford

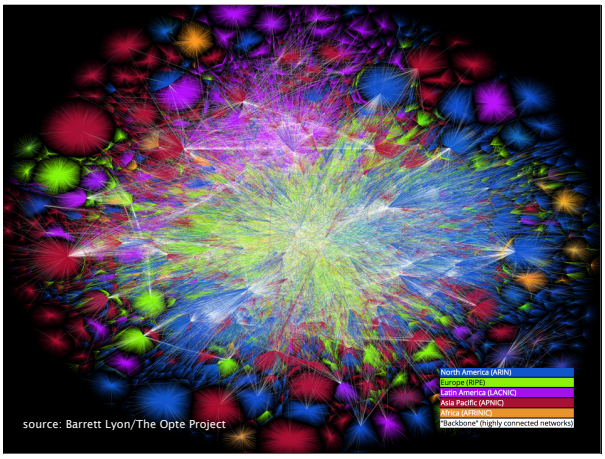

source: Barrett Lyon/The Opte Project

The Internet

An *exciting* place

18 billion

18 billion

estimated* # of Internet connected devices
in 2017

* Cisco Visual Networking Index 2017—2022

28.5 billion

estimated* # of Internet connected devices
in 2022

* Cisco Visual Networking Index 2017—2022

~4 exabytes

estimated* daily global IP traffic
in 2017

* Cisco Visual Networking Index 2017—2022

If  = 1 Gigabyte

~4 exabytes

estimated* daily global IP traffic
in 2017

* Cisco Visual Networking Index 2017—2022

~13 exabytes

estimated* daily global IP traffic
in 2022

* Cisco Visual Networking Index 2017—2022

~75% of all IP traffic

estimated* percentage of video traffic
in 2017

* Cisco Visual Networking Index 2017—2022

Upstream		Downstream		Aggregate	
BitTorrent	18.37%	Netflix	35.15%	Netflix	32.72%
YouTube	13.13%	YouTube	17.53%	YouTube	17.31%
Netflix	10.33%	Amazon Video	4.26%	HTTP - OTHER	4.14%
SSL - OTHER	8.55%	HTTP - OTHER	4.19%	Amazon Video	3.96%
Google Cloud	6.98%	iTunes	2.91%	SSL - OTHER	3.12%
iCloud	5.98%	Hulu	2.68%	BitTorrent	2.85%
HTTP - OTHER	3.70%	SSL - OTHER	2.53%	iTunes	2.67%
Facebook	3.04%	Xbox One Games Download	2.18%	Hulu	2.47%
FaceTime	2.50%	Facebook	1.89%	Xbox One Games Download	2.15%
Skype	1.75%	BitTorrent	1.73%	Facebook	2.01%
	69.32%		74.33%		72.72%

sandvine

Table 1 - Top 10 Peak Period Applications - North America, Fixed Access

http://bit.ly/2GlwI8G

~82% of all IP traffic

estimated* percentage of video traffic
in 2022

* Cisco Visual Networking Index 2017—2022

The Internet

A *tense* place

Countries get disconnected
for political reasons

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

Internet traffic to/from Egypt in January 2011

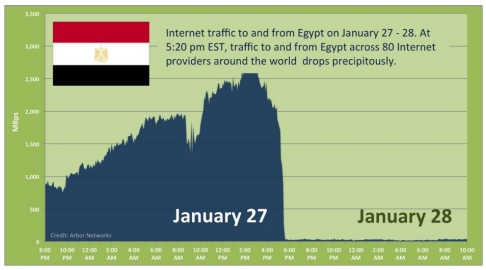

http://huff.to/1KxxoZF

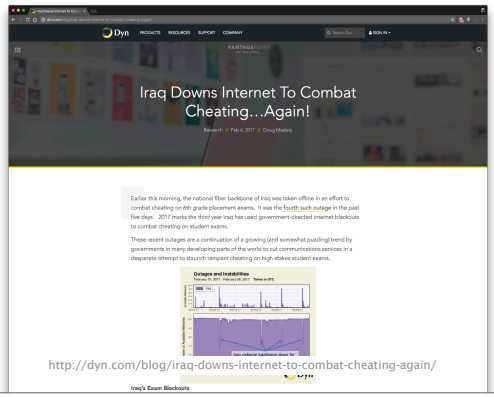

http://dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

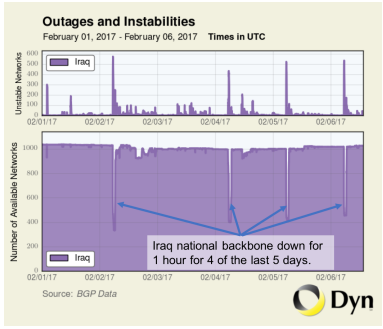

http://dyn.com/blog/iraq-downs-internet-to-combat-cheating-again/

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

https://www.independent.co.uk/news/world/africa/algeria-iraq-shut-down-internet-students-cheating-exams-facebook-a8410341.html

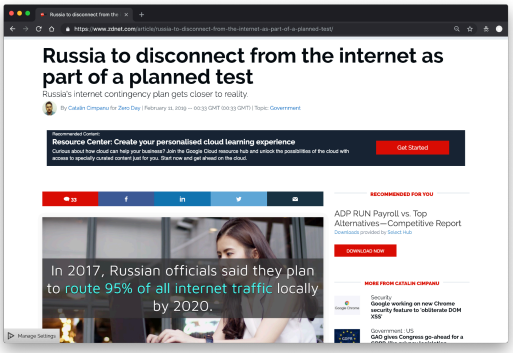

https://www.zdnet.com/article/russia-to-disconnect-from-the-internet-as-part-of-a-planned-test/

Some Internet communications
are interfered against or heavily congested

Can ISPs selectively slow down traffic?

The U.S. Federal Communications Commission (FCC) set network neutrality rules in 2015

http://nyti.ms/2kZUnDA

… which it then repealed in 2017

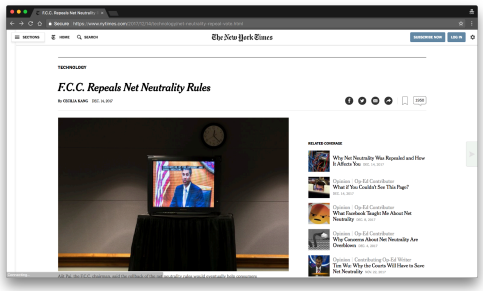

http://nyti.ms/2CkTbRR

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

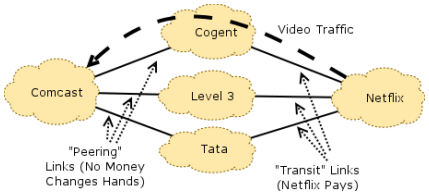

https://freedom-to-tinker.com/blog/feamster/why-your-netflix-traffic-is-slow-and-why-the-open-internet-order-wont-necessarily-make-it-faster/

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

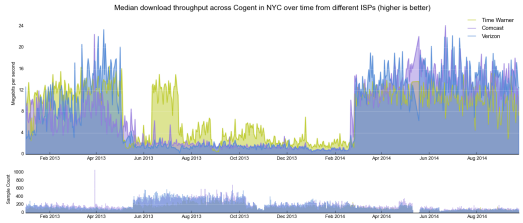

http://bit.ly/1thPzro

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

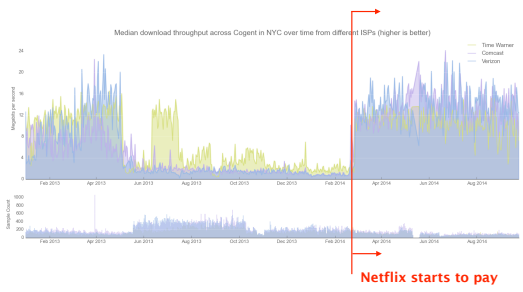

http://arstechnica.com/tech-policy/2014/04/netflix-and-verizon-reach-interconnection-deal-to-speed-up-video/

Closer to us…

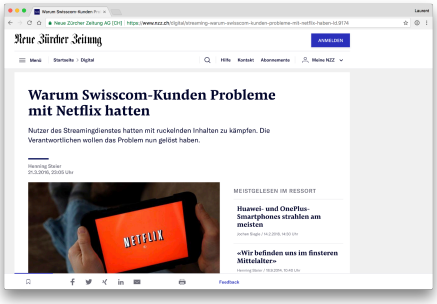

https://www.nzz.ch/digital/streaming-warum-swisscom-kunden-probleme-mit-netflix-haben-ld.9174

In February 2018, GitHub was targeted by the (so far) largest DDoS attack…

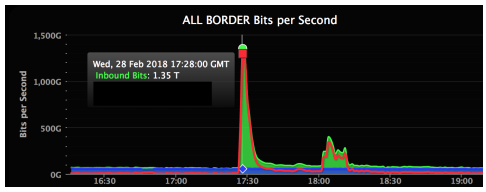

from a normal ~0.1 Tbps to **1.35 Tbps**

Source: Akamai

At the same time, countermeasures improve…

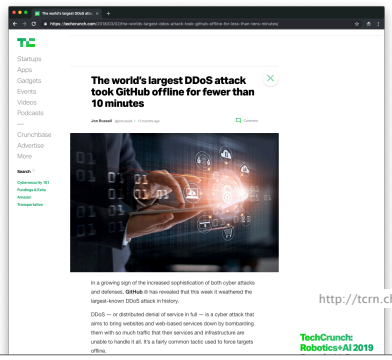

http://tcrn.ch/2tbJmFD

The Internet

A *fragile* place

Despite being absolutely critical,

Internet communications are inherently fragile

November 2017

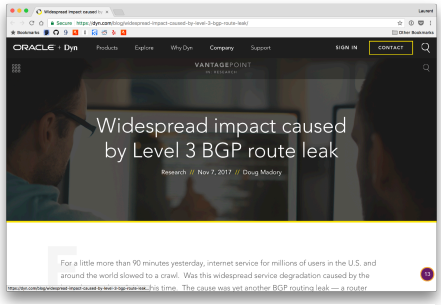

https://dyn.com/blog/widespread-impact-caused-by-level-3-bgp-route-leak/

For a little more than 90 minutes […],

Internet service for millions of users in the U.S. and around the world slowed to a crawl.

The cause was yet another BGP routing leak, a router misconfiguration directing Internet traffic from its intended path to somewhere else.

August 2017

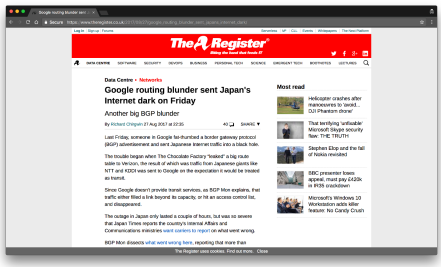

https://www.theregister.co.uk/2017/08/27/google_routing_blunder_sent_japans_internet_dark/

Someone in Google fat-thumbed a
Border Gateway Protocol (BGP) advertisement
and sent Japanese Internet traffic into a black hole.

[…] the result of which was traffic from Japanese giants
like NTT and KDDI was sent to Google
on the expectation it would be treated as transit.

The outage in Japan only lasted a couple of hours,
but was so severe that […] the country's
Internal Affairs and Communications ministries
want carriers to report on what went wrong.

People also often mistakenly destroy
their own infrastructure

Traders work on the floor of the New York Stock Exchange (NYSE) in July 2015. (Photo by Spencer Platt/Getty Images)

DOWNTIME

UPDATED: "Configuration Issue" Halts Trading on NYSE

The article has been updated with the time trading resumed.

A second update identified the cause of the outage as a "configuration issue."

A third update added information about a software update that created the configuration issue.

NYSE network operators identified
the culprit of the 3.5 hour outage,
blaming the incident on a
"network configuration issue"

Forbes / Personal Finance

The Little Black Book of Billionaire Secrets

11,261 VIEWS

United Airlines Blames Router for Grounded Flights

Alexandra Talty, CONTRIBUTOR

FULL BIO

After a computer problem caused nearly two hours of grounded flights for United Airlines this morning and ongoing delays throughout the day, the airline announced the culprit: a faulty router.

Spokeswoman Jennifer Dohm said that the router problem caused "degraded network connectivity," which affected various applications.

A computer glitch in the airline's reservations system caused the Federal Aviation Administration to impose a groundstop at 8:26 a.m. E.T. Planes that were in the air continued to operate, but all planes on the ground were held. There were reports of agents writing tickets by hand. The ground stop was lifted around 9:47 a.m. ET.

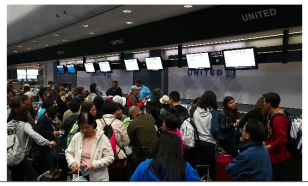

http://bit.ly/2sBJ2jf

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

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

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

Ironically, this means that data networks work better during week-ends…

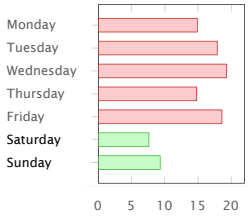

source: Job Snijders (NTT)

"Cost per network outage
can be as high as 750 000$"

Smart Management for Robust Carrier Network Health
and Reduced TCO!, NANOG54, 2012

Communication Networks
Course goals

Knowledge

Understand how the Internet works and why

from your
network plug…

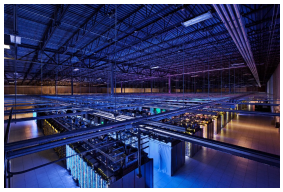

…to Google's data-center

Insights

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?**

Insights

Current research developments

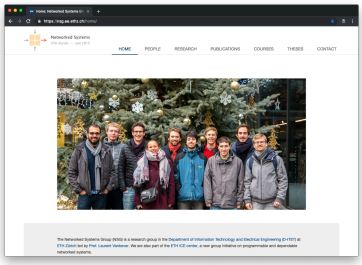

Networked Systems Group

nsg.ee.ethz.ch

Skills

Build, operate and configure networks

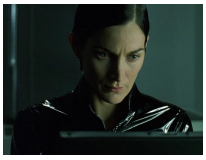

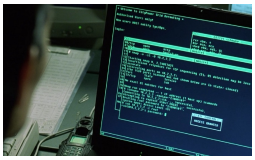

Trinity using a port scanner (`nmap`) in Matrix Reloaded™

Communication Networks

Course organization

Your **dream team** for the semester

Tobias [head]

Alexander

Rüdiger

Roland

Edgar

Ahmed

Thomas

Hendrik

Eric

followed the lecture in '18

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

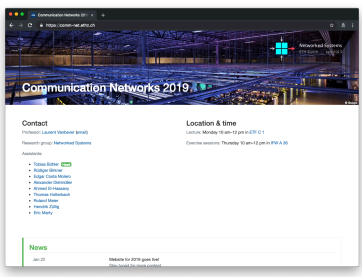

Slides, exercises, projects, extra readings, previous exams, …

The course will be **split in three parts**

Part 1	Part 2	Part 3
Overview	Concepts	Today's Internet
~1.5 lectures	~1.5 lectures	~10 lectures

Your final grade

Exam	Projects
80%	20%
written, open book	continuous performance assessments

Your final grade

Exam	Projects
80%	20%
written, open book	continuous performance assessments

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

#1 Build and operate a real, working "Internet"

#2 Implement an interoperable reliable protocol

Detailed instructions will follow

If you are a repeating student,
let us know if you want to keep your grades!

"Internet Hackathon"
sometime around week 7-8

2016 edition

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

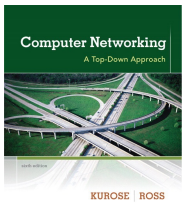

6th edition
ok to use the 5th or 7th

see sections indicated
on comm-net.ethz.ch

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

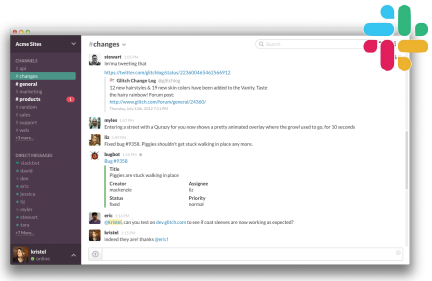

Web, smartphone and desktop clients available

Using Slack is facultative
but very useful nonetheless

Use Slack to

- ask questions
- chat with other students (e.g. your group)
- be informed about course announcements (also on our website)

Register today

> https://join.slack.com/t/comm-net19/signup

Register with your @ethz.ch email

Ping us if you prefer using another one

Use your real name

It greatly facilitates our organization…

We will never use Slack to distribute sensitive data

e.g. your project grades

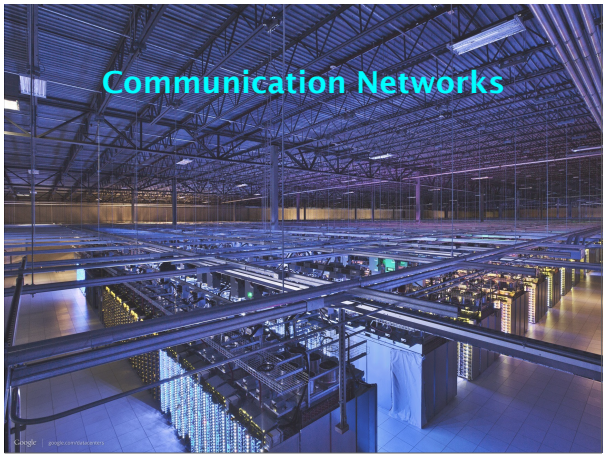

List any

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

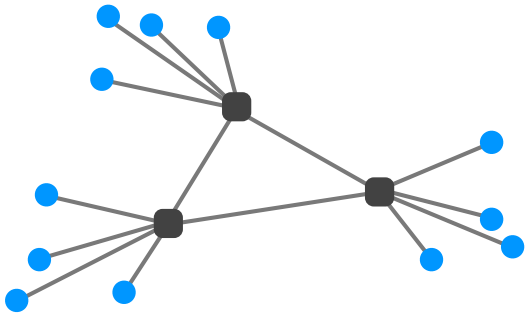

End-systems send & receive data

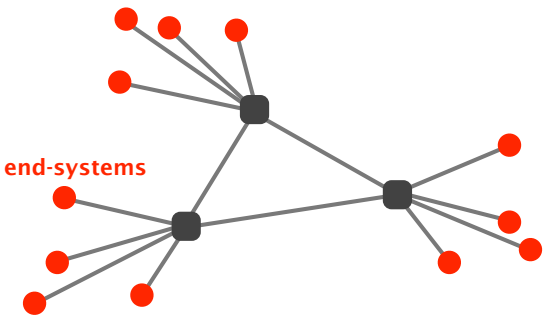

End-systems come in a wide-variety

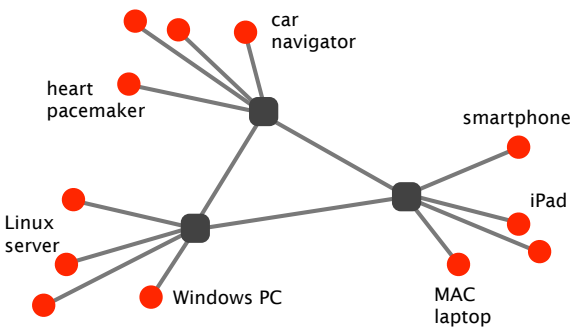

Switches & routers forward data to the destination

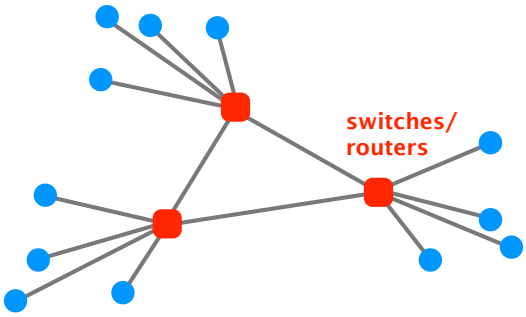

Routers/switches vary in size and usage

Home router

~20 cm
0,5 kg
1 Gbps

Internet core router

>200cm
700kg
>12 Tbps

(>920 Tbps in multi-chassis*)

*https://www.cisco.com/c/en/us/products/collateral/routers/carrier-routing-system/data_sheet_c78-726136.html

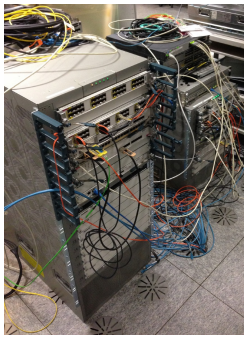

Cisco Nexus 7k
Routers @ETHZ

~25 deployed

Next-generation programmable switches
up to 12.8 Tbps of backplane capacity*

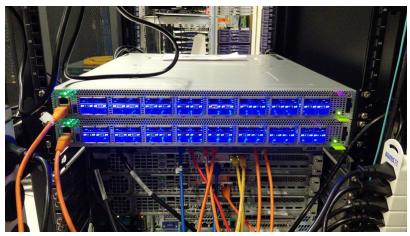

Barefoot Tofino Wedge 100BF-32X

part of our NSG lab

* https://www.barefootnetworks.com/products/brief-tofino-2/

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

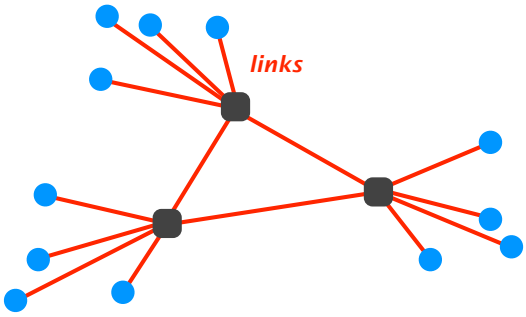

Links, too, vary in size and usage

Copper
ADSL, RJ-45,...

Optical fibers

Wireless link

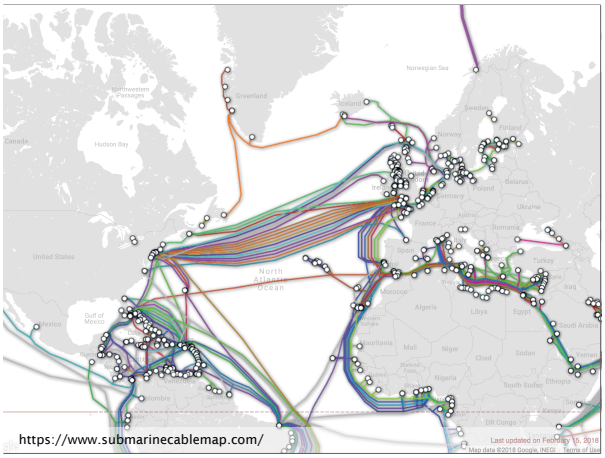

https://www.submarinecablemap.com/

http://marine.orange.com/en/Ships-and-submarine-vehicles/Cable-ships/Le-Rene-Descartes

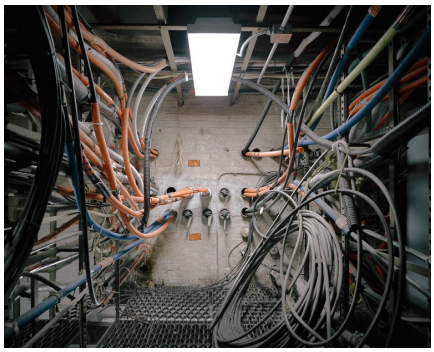

Somewhere in Manhattan… http://www.petergarritano.com/the-internet.html

The *Inter*net is a network of networks

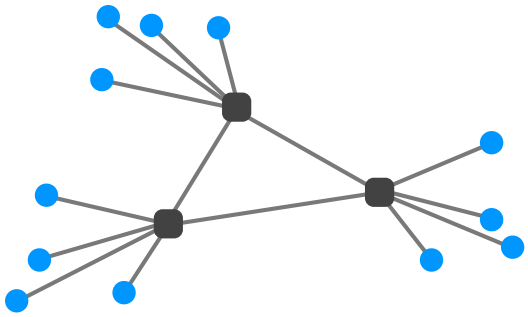

Internet Service Providers

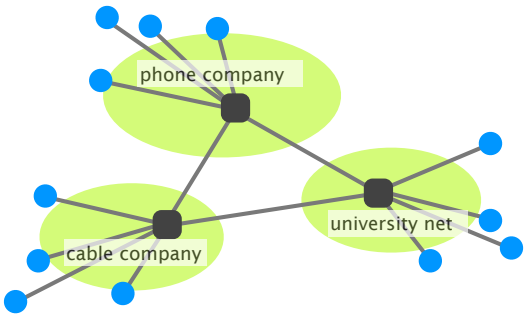

There exists a huge amount of access technologies

Ethernet	most common, symmetric
DSL	over phone lines, asymmetric
CATV	via cable TV, shared
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?

So far, we've been discussing what the "last mile" of the Internet looks like

What about the rest of the network?

3 must-have requirements of a good network topology

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
		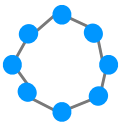	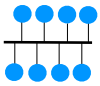
advantages			
disadvantages			

Switched networks provide reasonable and flexible compromise

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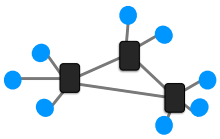
advantages	sharing and per-node capacity can be adapted to fit the network needs
disadvantages	require smart devices to perform: forwarding, routing, resource allocation

Links and switches are shared between flows

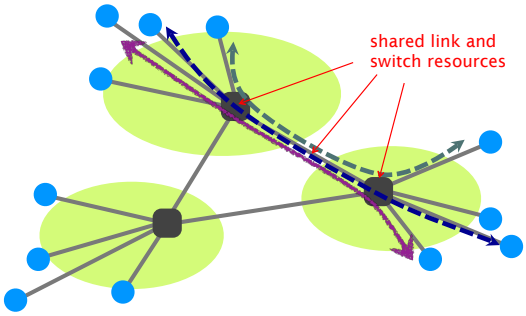

There exist two approaches to sharing: reservation and on-demand

	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

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

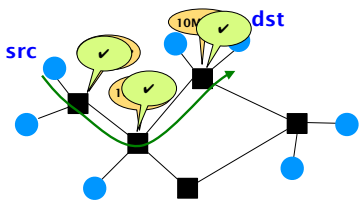

(1) src sends a reservation request for 10Mbps to dst
(2) switches "establish a circuit"
(3) src starts sending data
(4) src sends a "teardown circuit" message

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

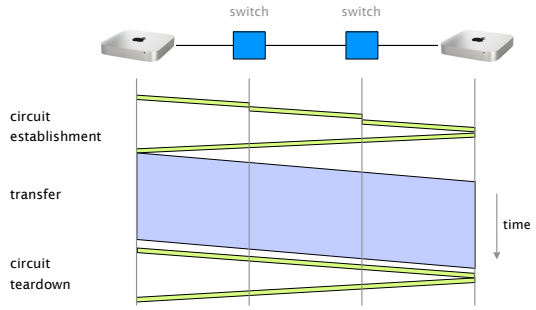

The efficiency of the transfer depends on
how utilized the circuit is once established

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

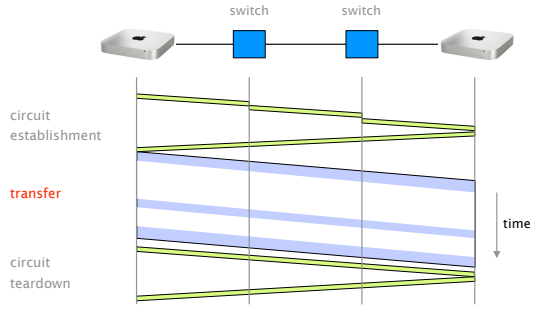

This is another example of poor efficiency.
The circuit is used for a short amount of time

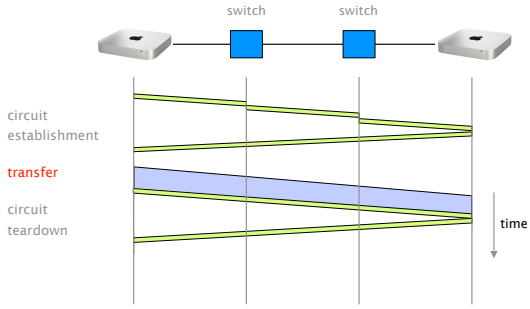

Another problem of circuit switching is that
it doesn't route around trouble

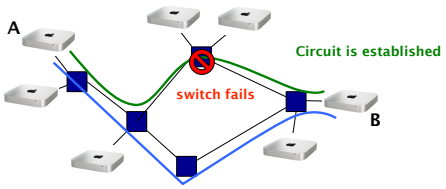

A is forced to signal a new circuit to
restore communication

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

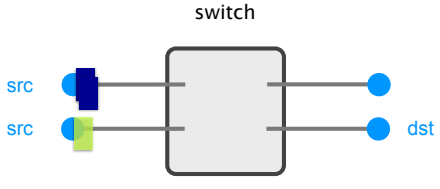

Each packet contains a destination (dst)

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

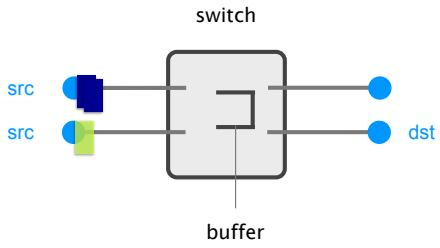

Packet switching routes around trouble

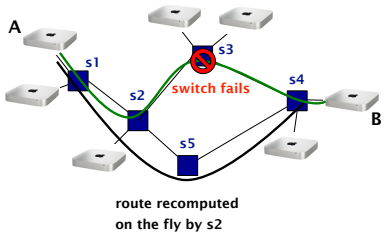

Pros and cons of packet switching

advantages	disadvantages
efficient use of resources	unpredictable performance
simpler to implement	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

- 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…

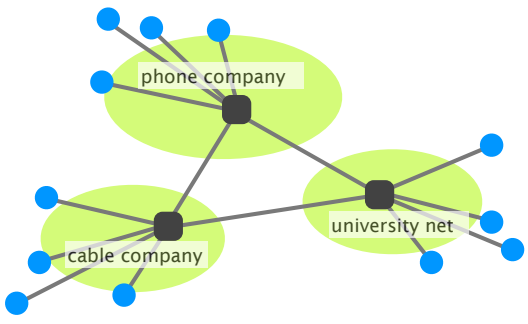

The real Internet is a "tad" more complex

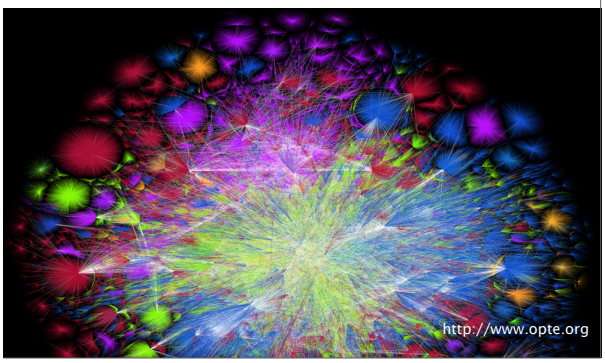

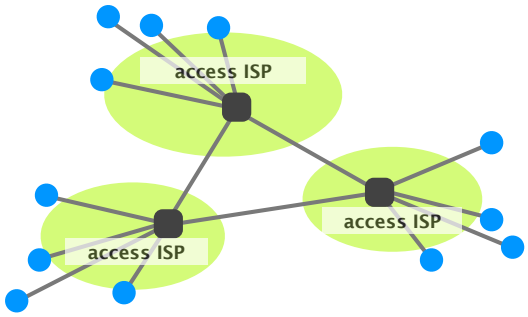

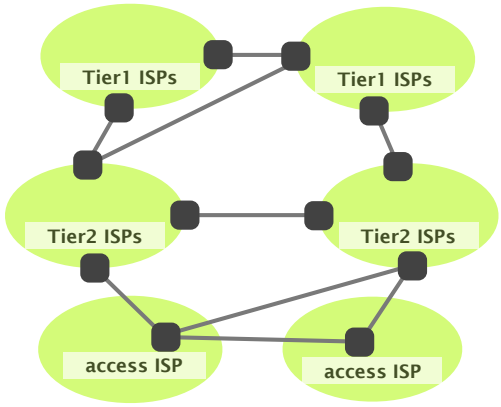

The Internet has a hierarchical structure

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

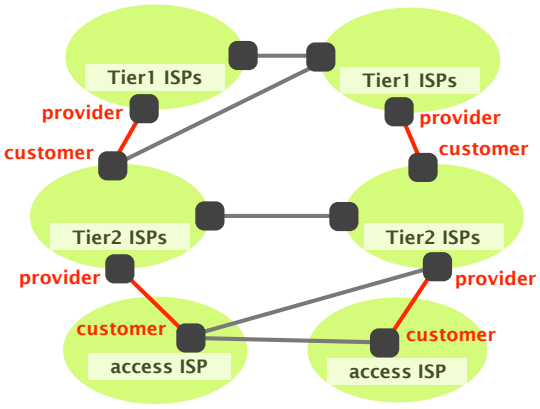

The distribution of networks in Tiers
is extremely skewed towards Tier-3s

	total	~60,000 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%

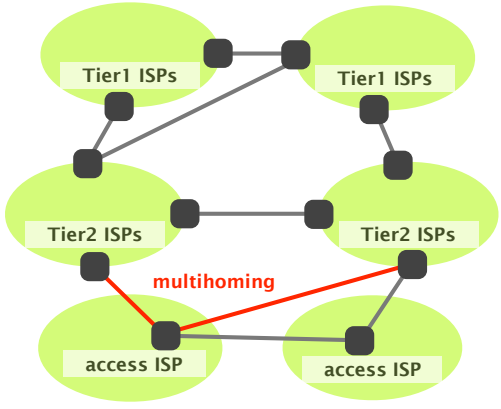

Some networks have an incentive to connect directly,
to reduce their bill with their own provider

This is known as "peering"

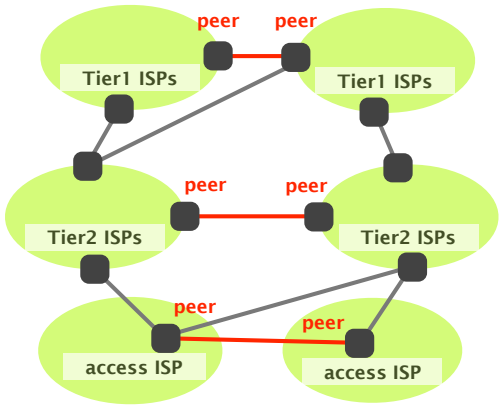

Interconnecting each network to its neighbors one-by-one is not cost effective

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

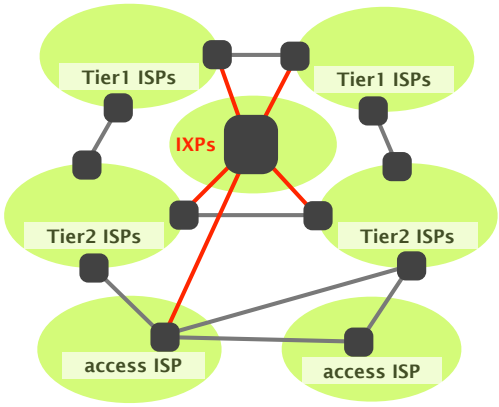

A brief overview of Internet history

The Internet history starts in the late 50's, with people willing to communicate differently

Telephone network is *the* communication system
entirely based on circuit switching

People start to want to use networks for other things
defense, (not personal) computers, …

… but knew that circuit-switching will not make it
too inefficient for bursty loads and not resilient

From this wish arose three crucial questions

Paul Baran RAND	How can we design a more resilient network? lead to the invention of packet switching
Len Kleinrock UCLA	How can we design a more efficient network? (also) lead to the invention of packet switching
Bob Kahn DARPA	How can we connect all these networks together? lead to the invention of the Internet as we know it

The 60s saw the creation of packet switching and the Advanced Research Projects Agency Network

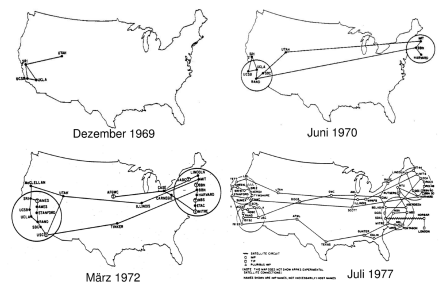

The first message ever exchanged on the Internet was “lo”

Oct. 29 1969 — Leonard Kleinrock @UCLA tries to log in a Stanford computer

UCLA: We typed the L... Do you see it?

Yes! We see the L — Stanford

We typed the O... Do you see it?

Yes! We see the O

We typed the G. system crashes

http://ftp.cs.ucla.edu/csd/first_words.html

The 70s saw the creation of Ethernet, TCP/IP and the e-mail

1971	Network Control Program predecessor of TCP/IP
1972	Email & Telnet
1973	Ethernet
1974	TCP/IP paper by Vint Cerf & Bob Kahn

In the 80s, TCP/IP went mainstream

1983	NCP to TCP/IP Flag day Domain Name Service (DNS)
1985	NSFNet (TCP/IP) succeeds to ARPANET
198x	Internet meltdowns due to congestion
1986	Van Jacobson saves the Internet (with congestion control)

The 90s saw the creation of the Web as well as the Internet going commercial

1989	Arpanet is decommissioned
	Birth of the Web Tim Berners Lee (CERN)
1993	Search engines invented (Excite)
1995	NSFNet is decommissioned
1998	Google reinvents search

The new millennium brings the Web 2.0, focus on user-generated content

1998	IPv6 standardization
2004	Facebook goes online
2006	Google buys YouTube
2007	Netflix starts to stream videos
2007	First iPhone Mobile Internet access

Fast Internet access everywhere, every device needs an Internet connection

2009	Mining of the Bitcoin genesis block
	Fast mobile Internet access: 4G/LTE
	Internet of Things (IoT) boom Cars & refrigerators in the Internet
2018	Only 26% of the Alexa Top 1000 websites reachable over IPv6 http://www.worldipv6launch.org/measurements/
Soon?	Fully encrypted transport protocols For example QUIC

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?

No exercise session this Thursday

Next Monday on Communication Networks

Routing concepts