Communication Networks

Prof. Laurent Vanbever

Communication Networks

Spring 2022

Laurent Vanbever
nsg.ee.ethz.ch

21 February 2022
D-ITET

Materials inspired from Scott Shenker & Jennifer Rexford

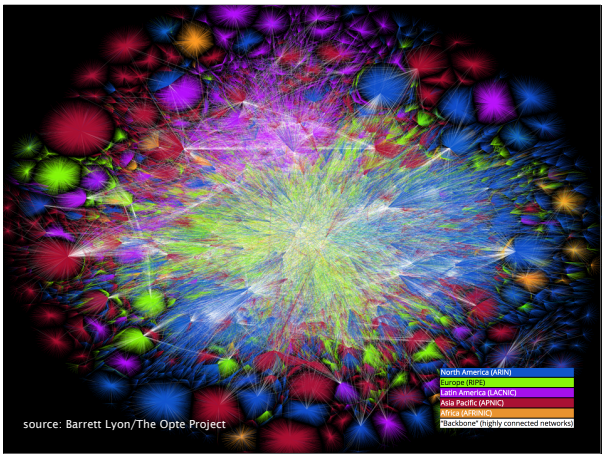

source: Barrett Lyon/The Opte Project

The Internet

An *exciting* place

~22 billion

~22 billion

estimated* # of Internet connected devices
in 2020

* Cisco Visual Networking Index 2018—2023

~30 billion

estimated* # of Internet connected devices
in 2023

* Cisco Visual Networking Index 2018—2023

~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 Internet 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 Internet 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

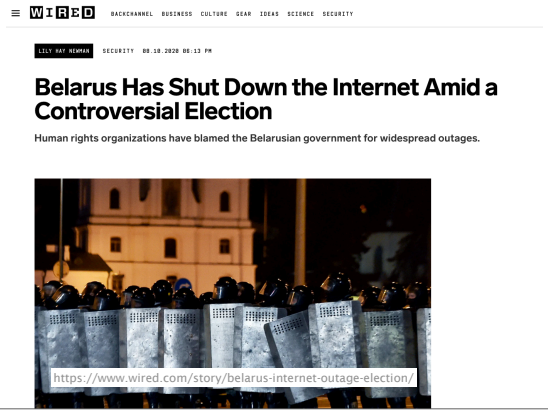

https://www.independent.co.uk/news/world/africa/algeria-iraq-shut-down-internet-students-cheating-exams-facebook-a8410341.html

Internet communications get congested
for economical reasons

Can ISPs selectively slow down traffic?

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

… which it then repealed in 2017

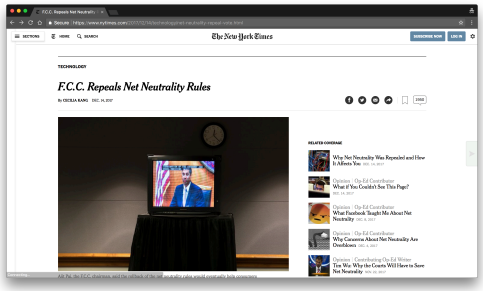

http://nyti.ms/2CkTbRR

Netflix US
@netflix
Follow

We're disappointed in the decision to gut #NetNeutrality protections that ushered in an unprecedented era of innovation, creativity & civic engagement. This is the beginning of a longer legal battle. Netflix stands w/ innovators, large & small, to oppose this misguided FCC order.

10:26 AM - 14 Dec 2017

335,726 Retweets 831,986 Likes

7.1K 336K 832K

… but might restore soon

https://www.forbes.com/sites/waynerash/2021/01/26/net-neutrality-likely-to-return-with-new-fcc-chair/

In Switzerland, network neutrality is enforced by the Swiss Telecommunications Act since 1/1/21

– Art. 12e[41] Offenes Internet

[1] Die Anbieterinnen von Internetzugängen übertragen Informationen, ohne dabei zwischen Sendern, Empfängern, Inhalten, Diensten, Dienstekklassen, Protokollen, Anwendungen, Programmen oder Endgeräten technisch oder wirtschaftlich zu unterscheiden.

[2] Sie dürfen Informationen unterschiedlich übertragen, wenn dies erforderlich ist, um:

a. eine gesetzliche Vorschrift oder einen Gerichtsentscheid zu befolgen;
b. die Integrität oder Sicherheit des Netzes, der über dieses Netz erbrachten Dienste oder der angeschlossenen Endgeräte zu gewährleisten;
c. einer ausdrücklichen Aufforderung der Kundin oder des Kunden nachzukommen; oder
d. vorübergehende und aussergewöhnliche Netzwerküberlastungen zu bekämpfen; dabei sind gleiche Arten von Datenverkehr gleich zu behandeln.

[3] Sie dürfen neben dem Zugang zum Internet über denselben Anschluss andere Dienste anbieten, die für bestimmte Inhalte, Anwendungen oder Dienste optimiert sein müssen, um die Qualitätsanforderungen der Kundinnen und Kunden zu erfüllen. Die anderen Dienste dürfen nicht als Ersatz für Internetzugangsdienste nutzbar sein oder angeboten werden, und sie dürfen nicht die Qualität der Internetzugangsdienste verschlechtern.

[4] Behandeln sie Informationen bei der Übertragung technisch oder wirtschaftlich unterschiedlich, so müssen sie die Kundinnen und Kunden sowie die Öffentlichkeit darüber informieren.

[41] Eingefügt durch Ziff. I des BG vom 22. März 2019, in Kraft seit 1. Jan. 2021 (*AS 2020 6159*; BBl 2017 6559).

https://www.fedlex.admin.ch/eli/cc/1997/2187_2187_2187/fr#art_12_e

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

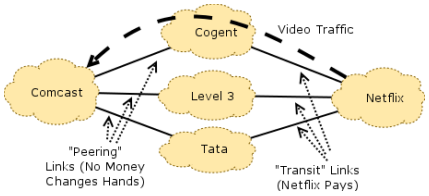

https://freedom-to-tinker.com/blog/feamster/why-your-netflix-traffic-is-slow-and-why-the-open-internet-order-wont-necessarily-make-it-faster/

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

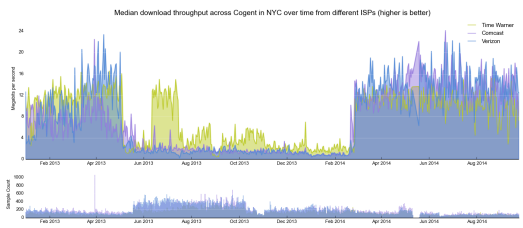

http://bit.ly/1thPzro

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

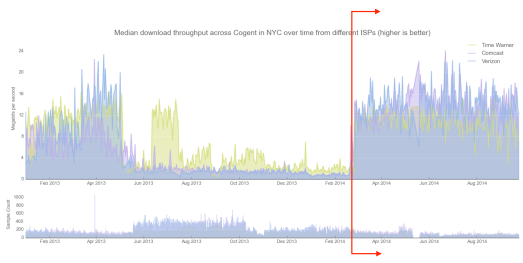

Netflix starts to pay

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

Closer to us…

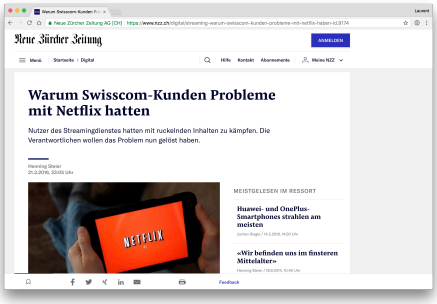

https://www.nzz.ch/digital/streaming-warum-swisscom-kunden-probleme-mit-netflix-haben-ld.9174

Internet infrastructures are regularly targeted by large-scale attacks

In February 2018, GitHub was targeted by a 1.35 Tbps Distributed Denial of Service (DDoS) attack

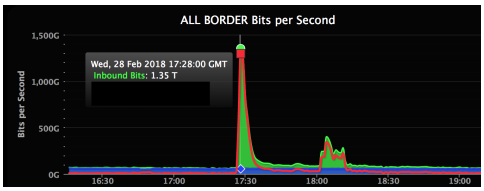

from a normal ~0.1 Tbps to 1.35 Tbps

Source: Akamai

In June 2020, Amazon was targeted by a 2.30 Tbps DDoS attack

In August 2021, Microsoft was targeted by a 2.40 Tbps DDoS attack

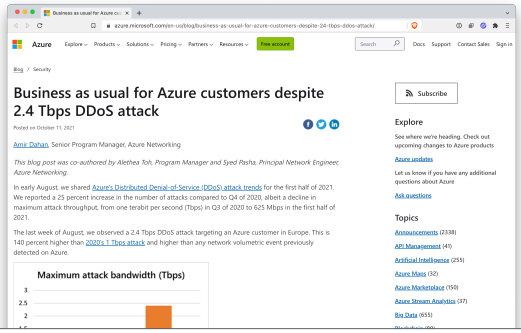

The Internet

A *vital* place during a pandemic

Following the lockdown in March 2020, (wired) networks saw traffic increasing by 15–20%

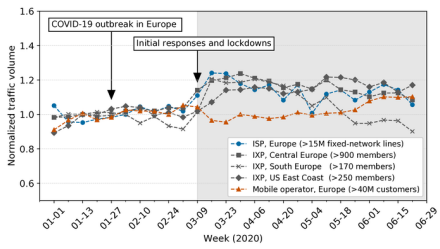

https://labs.ripe.net/Members/oliver_gasser/the-lockdown-effect-implications-of-the-covid-19-pandemic-on-internet-traffic

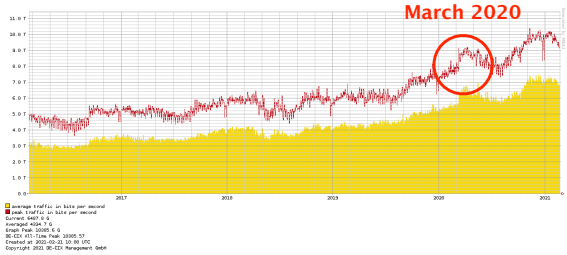

https://www.de-cix.net/en/locations/germany/frankfurt/statistics

Unsurprisingly, we saw a strong increase in
web conferencing, video, and gaming traffic

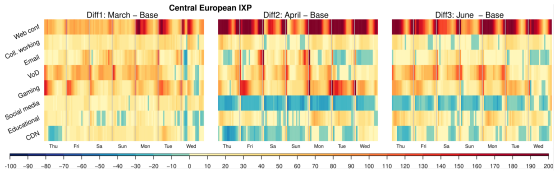

https://labs.ripe.net/Members/oliver_gasser/the-lockdown-effect-implications-of-the-covid-19-pandemic-on-internet-traffic

All in all the Internet performed very well
in these unpreceeding times

Measuring the Internet 13 May 2020

EN FR ES

The Internet Is Resilient Enough to Withstand Coronavirus – But There's a Catch

By David Belson
Former Senior Director, Internet Research and Analysis

Earlier this year, as COVID-19 began to dominate our lives, the world turned to the Internet. This sudden shift to distance learning, working from home, and families sheltering in place drove up online streaming demand, placing additional load on Internet application platforms like Zoom, Netflix, and educational tools such as Kahoot. There was also a dramatic traffic increase across supporting network providers.

[source]

The Internet
A *fragile* place

Despite being absolutely critical,
the Internet infrastructure is inherently fragile

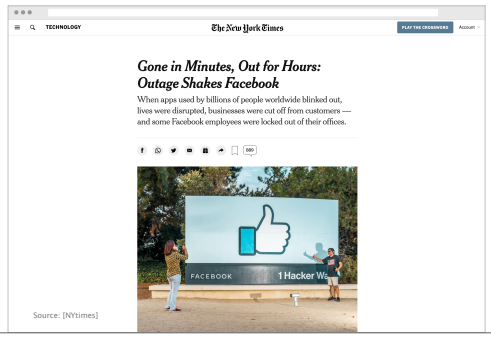

TECHNOLOGY

The New York Times

PLAY THE CROSSWORDS

Account

Gone in Minutes, Out for Hours: Outage Shakes Facebook

When apps used by billions of people worldwide blinked out, lives were disrupted, businesses were cut off from customers — and some Facebook employees were locked out of their offices.

Source: [NYtimes]

Our engineering teams have learned that
configuration changes on the backbone routers
that coordinate network traffic between our data centers
caused issues that interrupted this communication.

This disruption to network traffic had a cascading effect
on the way our data centers communicate,
bringing our services to a halt.

Source: [fb.com]

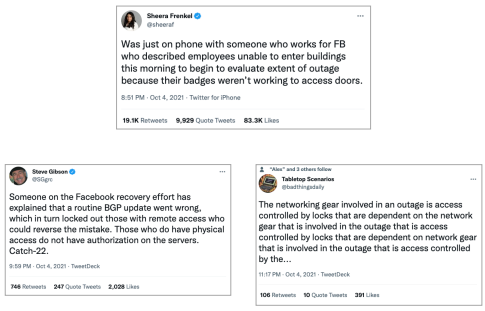

Sheera Frenkel
@sheeraf

Was just on phone with someone who works for FB who described employees unable to enter buildings this morning to begin to evaluate extent of outage because their badges weren't working to access doors.

8:51 PM · Oct 4, 2021 · Twitter for iPhone

19.1K Retweets 9,929 Quote Tweets 83.3K Likes

Steve Gibson
@SGgrc

Someone on the Facebook recovery effort has explained that a routine BGP update went wrong, which in turn locked out those with remote access who could reverse the mistake. Those who do have physical access do not have authorization on the servers. Catch-22.

9:59 PM · Oct 4, 2021 · TweetDeck

746 Retweets 247 Quote Tweets 2,028 Likes

"Alex" and 3 others follow

Tabletop Scenarios
@battlegalaxy

The networking gear involved in an outage is access controlled by locks that are dependent on the network gear that is involved in the outage that is access controlled by locks that are dependent on network gear that is involved in the outage that is access controlled by the...

11:17 PM · Oct 4, 2021 · TweetDeck

106 Retweets 10 Quote Tweets 391 Likes

Source: Twitter

August 2017

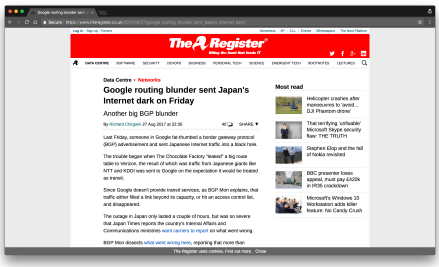

The Register

Data Centre • Networks

Google routing blunder sent Japan's Internet dark on Friday

Another big BGP blunder

By Richard Chirgwin 27 Aug 2017 at 22:35

Most read

Last Friday, someone in Google fat-thumbed a border gateway protocol (BGP) advertisement and sent Japanese Internet traffic into a black hole.

The trouble began when The Chocolate Factory "leaked" a big route table to Verizon, 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.

Since Google doesn't provide transit services, as BGP Mon explains, that traffic either filled a link beyond its capacity, or hit an access control list, and disappeared.

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

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.

In February 2020, a planned maintenance work in Swisscom's network shuts down emergency numbers

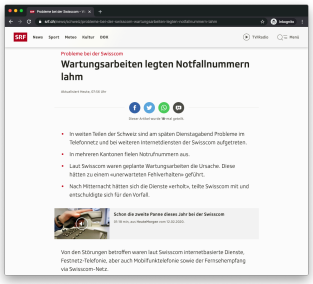

Internet, 4G, TV and telephone network affected as well

https://www.srf.ch/news/schweiz/probleme-bei-der-swisscom-wartungsarbeiten-legten-notfallnummern-lahm

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

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

Communication Networks

Course goals

Knowledge

Understand how the Internet works and why

from your network plug…

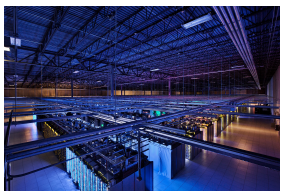

…to mega-scale data-centers

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?**

Skills

Build, operate and configure networks

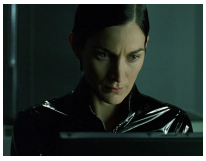

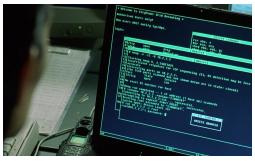

Trinity using a port scanner (`nmap`) in Matrix Reloaded™

Insights

Learn about some of our current research

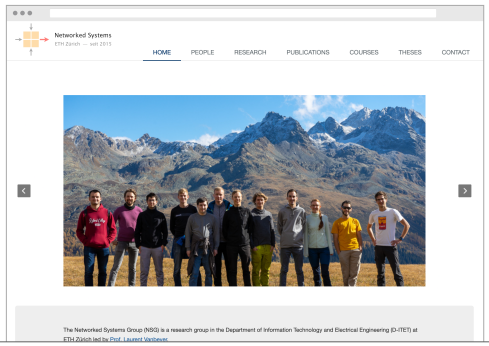

Communication Networks

Course organization

Your **dream team** for the semester

Tobias Coralie Alexander Tibor Thomas

+ **Martin and Nasib** who followed the lecture in previous years

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

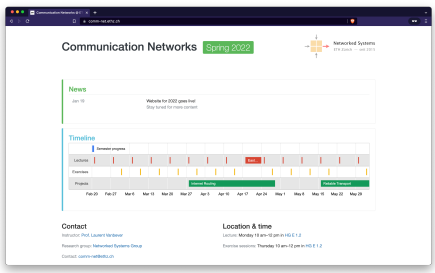

Slides, exercises, projects, extra readings, and previous exams

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

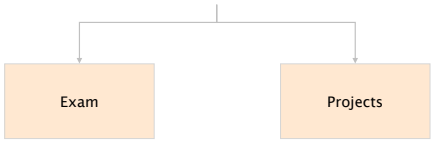

70%

written, open book

30%

continuous
performance assessments

Your final grade

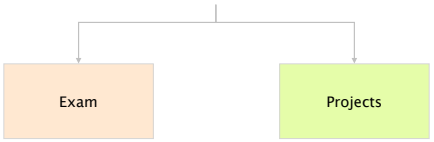

70%

written, open book

30%

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" (20%)

#2 Implement an interoperable reliable protocol (10%)

Detailed instructions will follow

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

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

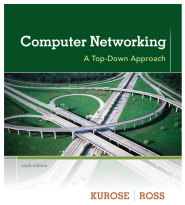

6th edition
using another edition is okay
but numbering might vary

see sections indicated
on comm-net.ethz.ch

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

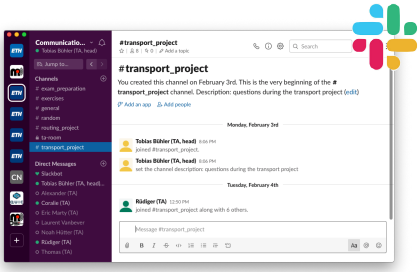

Web, smartphone and desktop clients available

Using Slack is highly recommended
but facultative

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-net22/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 never publish sensitive data on Slack
e.g. your grades

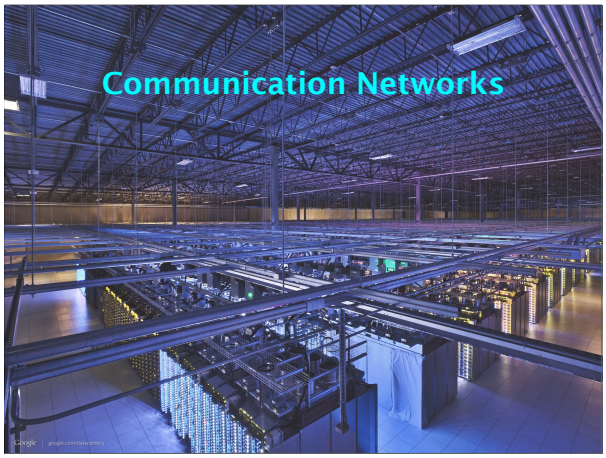

Communication Networks
Part 1: Overview

#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

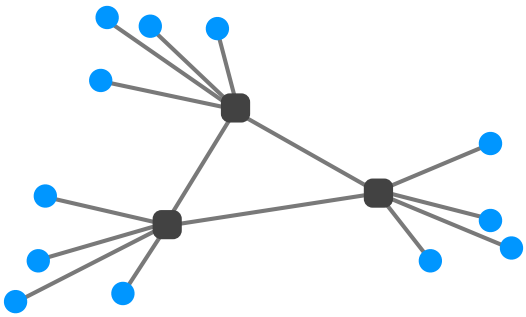

End-systems send & receive data

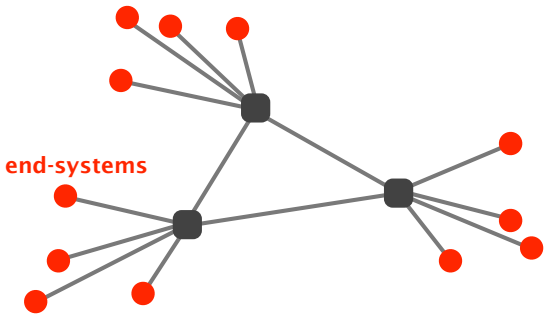

End-systems come in a wide-variety

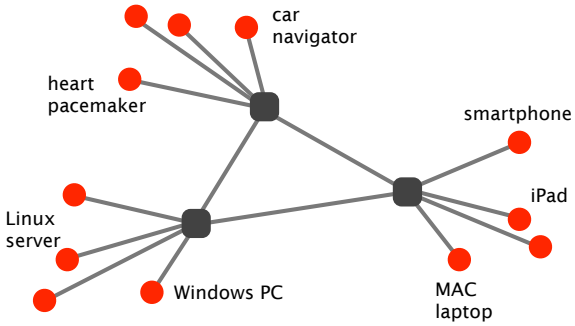

Switches & routers forward data to the destination

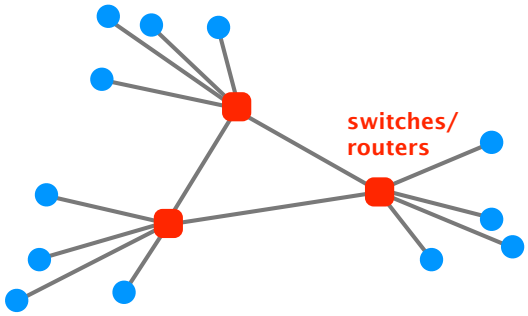

Routers/switches vary in size and usage

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

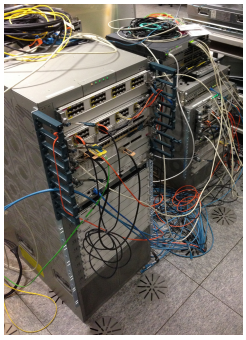

Cisco Nexus 7k

Routers @ETHZ

~25 deployed

Next-generation programmable switches
up to 25.6 Tbps of backplane capacity*

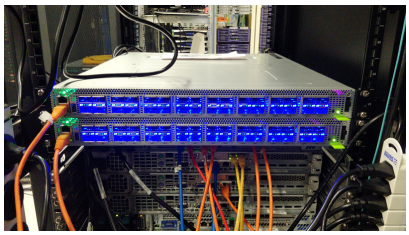

Barefoot Tofino Wedge 100BF-32X part of our NSG lab

* https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/tofino-3-brief.html

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

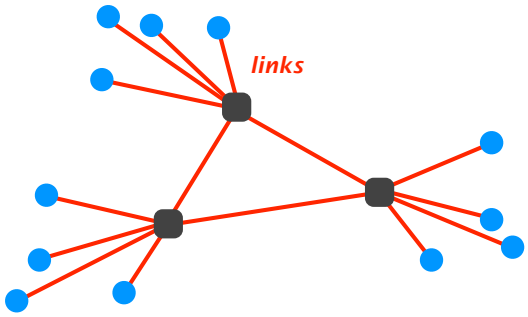

Links, too, vary in size and usage

Copper
ADSL, RJ-45,...

Optical fibers

Wireless link

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?

A good network topology fulfills
at least three requirements

Tolerate failures
>1 path should exist between each node

Allow sharing to be feasible & cost-effective
links should not be too high

Provide ample capacity
links should not be too small

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

design	full-mesh	chain	bus
	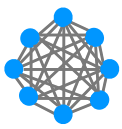	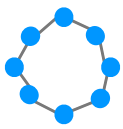	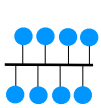
advantages			
disadvantages			

Switched networks provide
reasonable and flexible compromise

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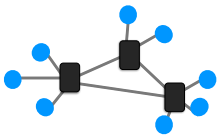
advantages	sharing and per-node capacity can be adapted to fit the network needs
disadvantages	require smart devices to perform: forwarding, routing, resource allocation

Links and switches are shared between flows

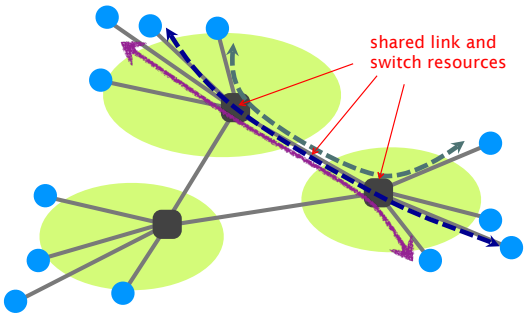

There exist two approaches to sharing: reservation and on-demand

	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?

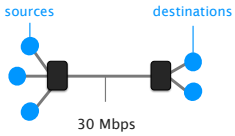

Consider that each source needs 10 Mbps

What do they get with:

- reservation
- on-demand

Assume the following peak demand and flow duration

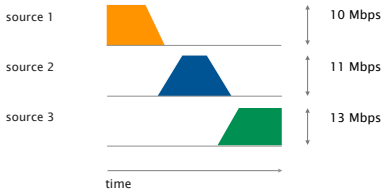

Assume the following peak demand and flow duration

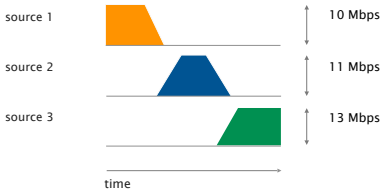

What does each source get with reservation and on-demand?

- 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*
P=100 Mbps, *A*=10 Mbps, level of utilization=10%

On-demand can usually achieve higher level of utilization
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

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!

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

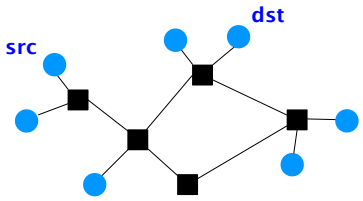

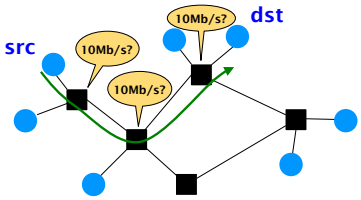

(1) src sends a reservation request for 10Mbps to dst
(2) switches "establish a circuit"
(3) src starts sending data
(4) src sends a "teardown circuit" message

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

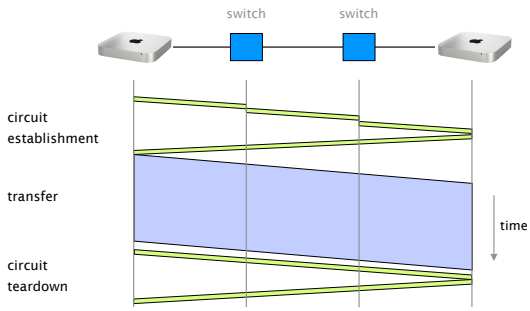

The efficiency of the transfer depends on
how utilized the circuit is once established

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

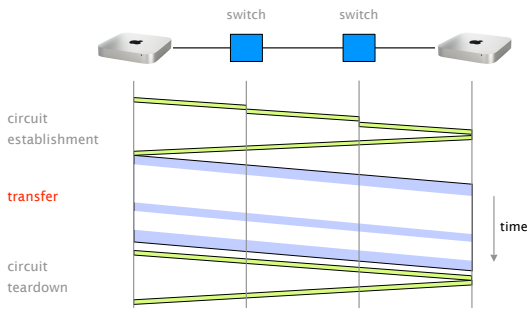

This is another example of poor efficiency.
The circuit is used for a short amount of time

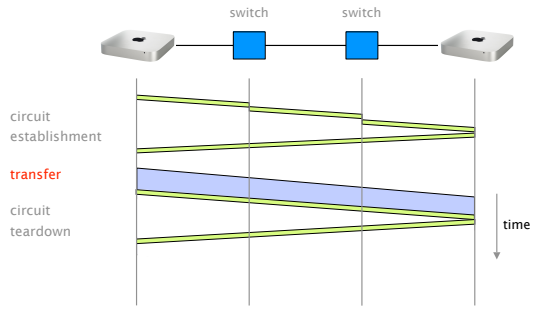

Another problem of circuit switching is that
it doesn't route around trouble

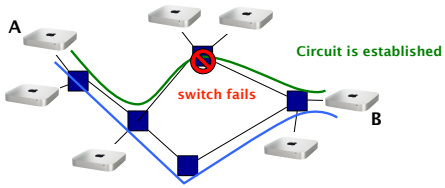

A is forced to signal a new circuit to restore communication

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

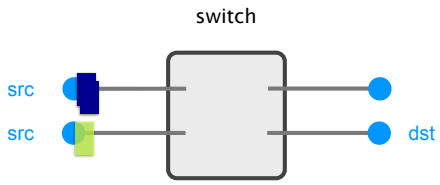

Each packet contains a destination (dst)

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

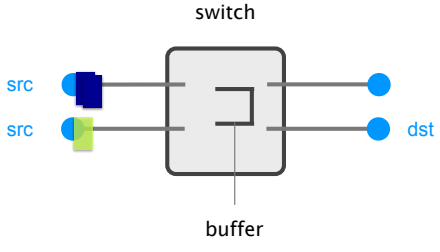

Packet switching routes around trouble

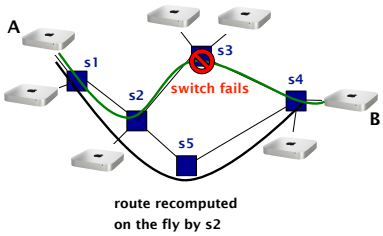

Pros and cons of packet switching

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?

The *Inter*net is a network of networks

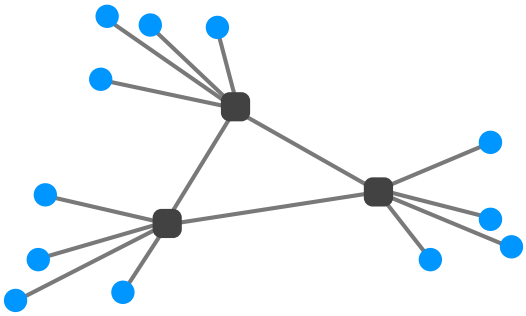

Internet Service Providers

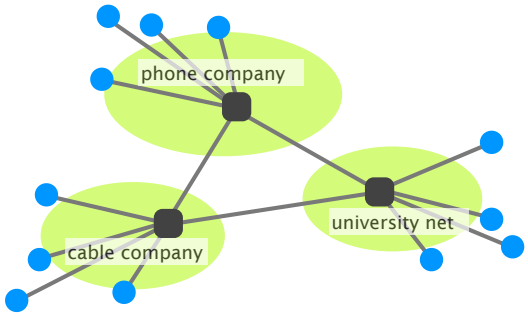

So far, this is our vision of the Internet…

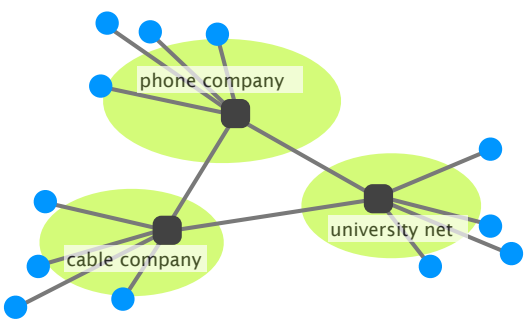

The real Internet is a "tad" more complex

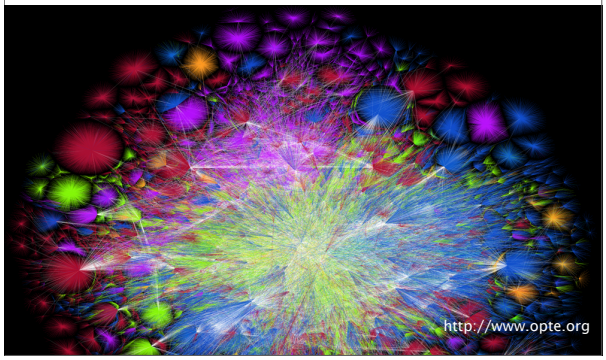

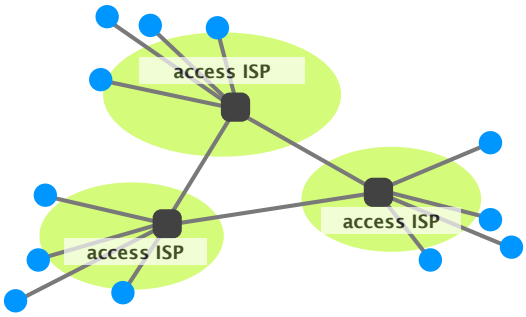

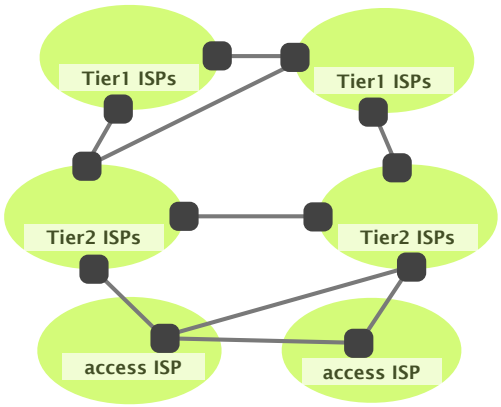

The Internet has a hierarchical structure

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

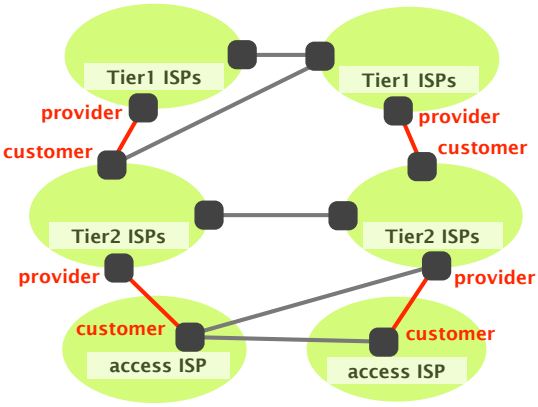

The distribution of networks in Tiers
is extremely skewed towards Tier-3s

	total	~70,000 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%

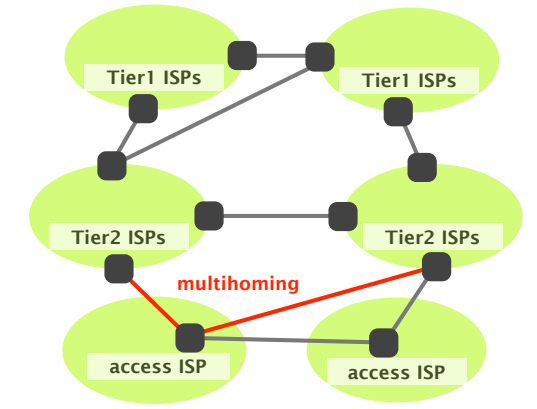

Some networks have an incentive to connect directly,
to reduce their bill with their own provider

This is known as "peering"

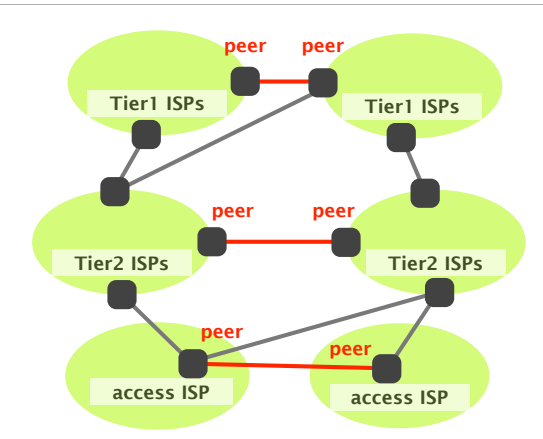

Interconnecting each network to its neighbors one-by-one is not cost effective

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

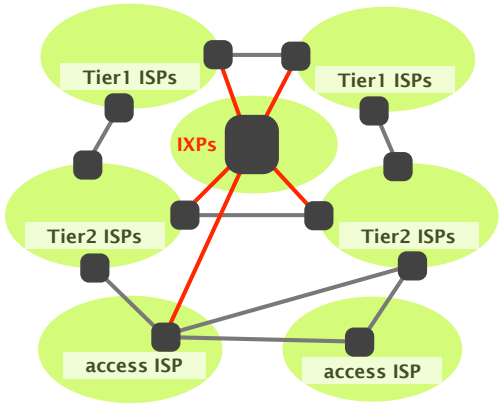

Communication Networks

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