High-speed and Programmable Networks

ECE/CS598HPN

Instructor: Radhika Mittal

Evolution of Computer Networks

1876: Alexander Graham Bell invented telephone.

Evolution of Computer Networks

1876: Alexander Graham Bell invented telephone.

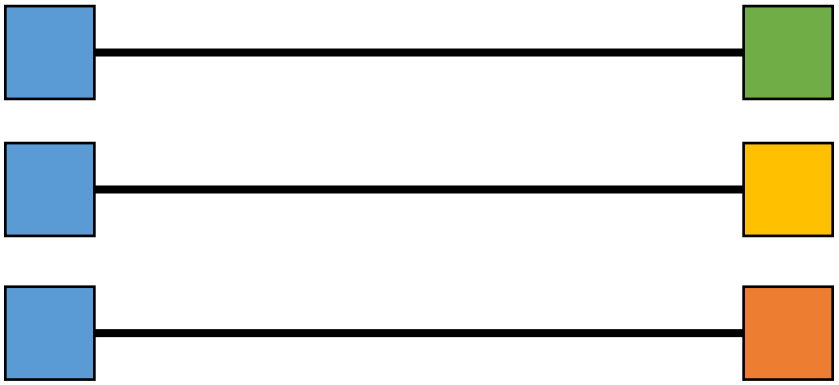

Such a design cannot scale!

Evolution of Computer Networks

Soon evolved to Public Switched Telephone Network.

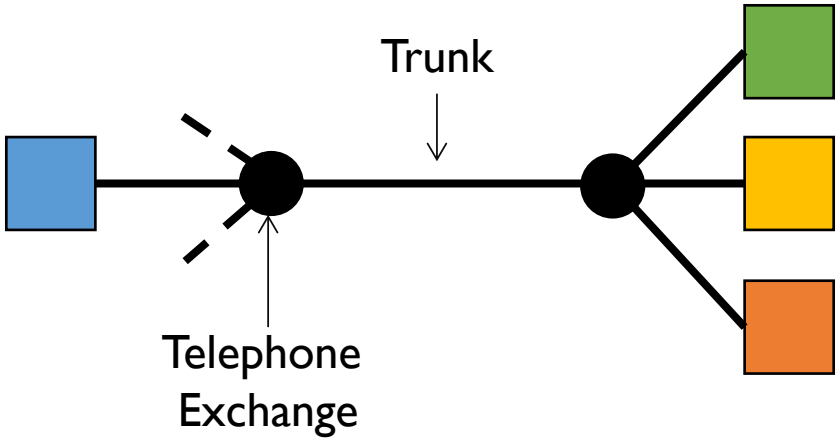

Evolution of Computer Networks

Manually operated!

Earliest circuit-switched network

Evolution of Computer Networks

Manually operated!

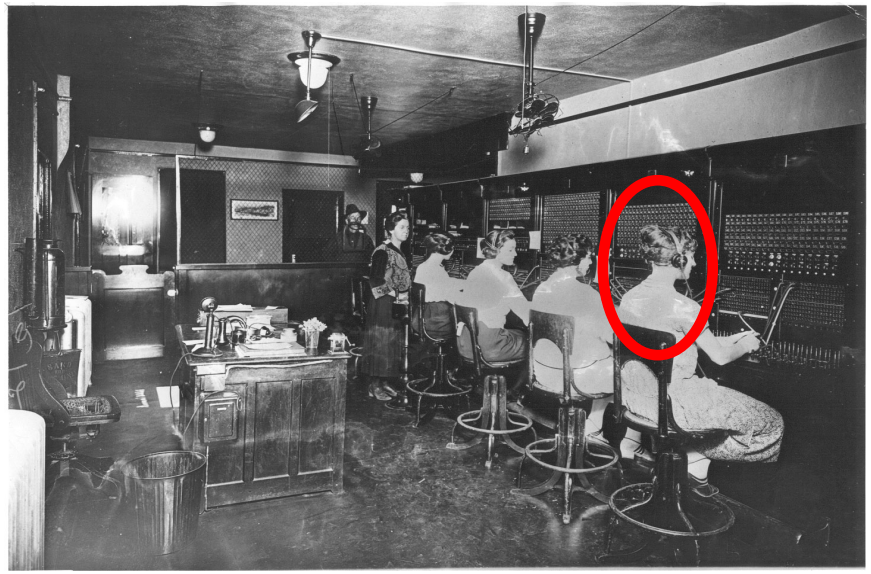

A.B. Strowger's competitor's wife

Earliest circuit-switched network

Evolution of Computer Networks

1889: AB Strowger invents first mechanical circuit switch.

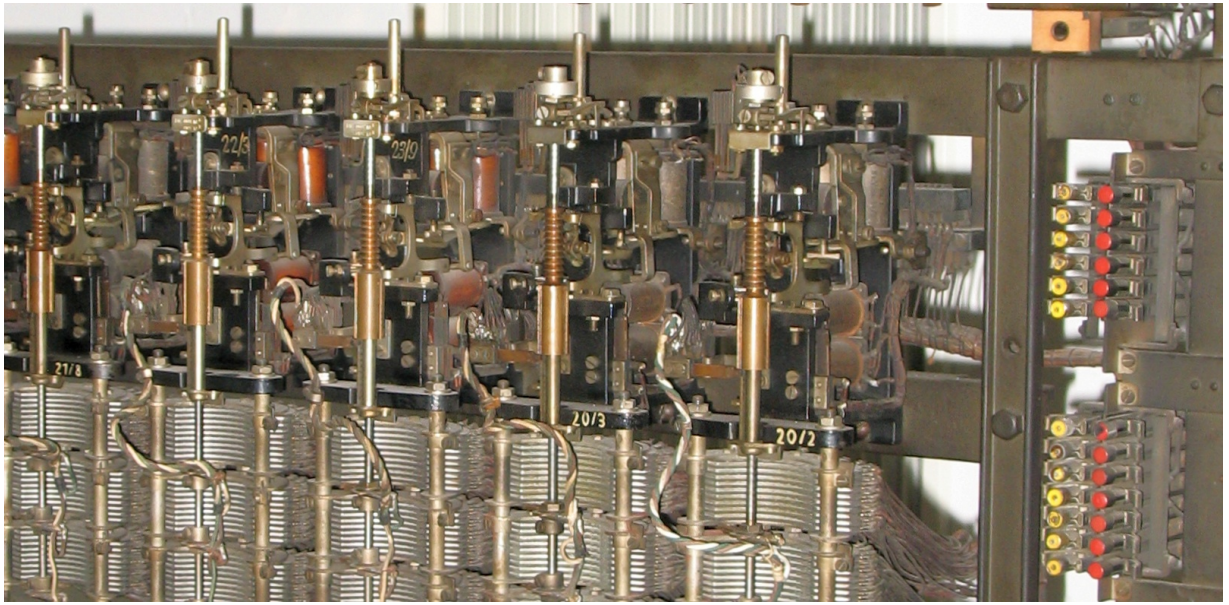

Earliest mechanical circuit-switched network!

Evolution of Computer Networks

Earliest circuit-switched network

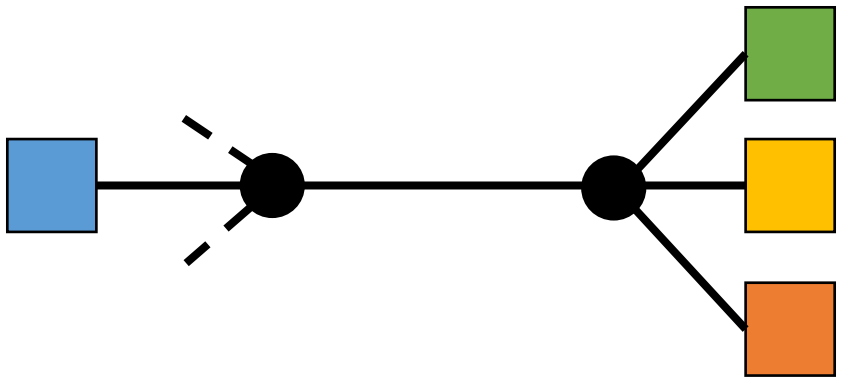

Evolution of Computer Networks

Circuit switching is wasteful!

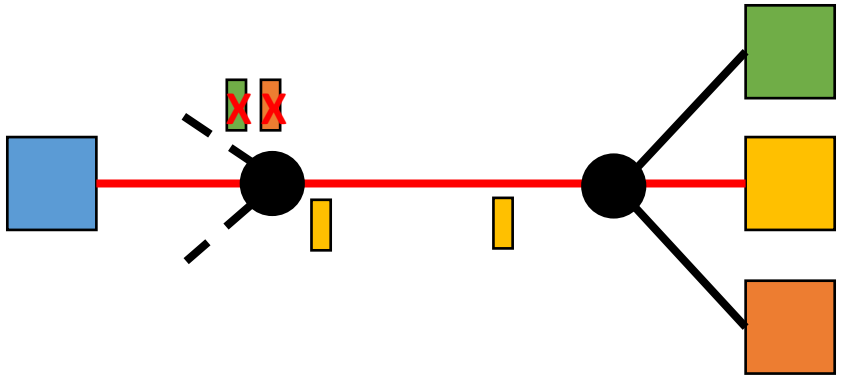

Evolution of Computer Networks

Packet switching is designed:
1959(Paul Baran), 1961(Leonard Kleinrock), 1965 (Donald Davies).

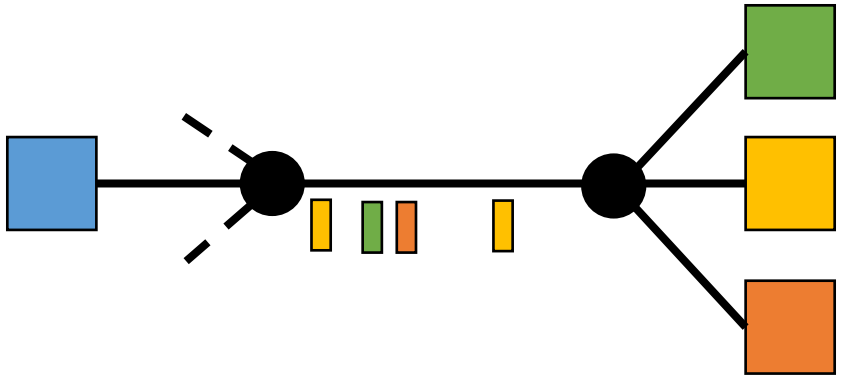

Evolution of Computer Networks

- Simultaneously, growing interest in connecting computers.
- Lawrence Roberts meets Davies' teammate at 1967 SOSP, and decides to use packet-switching for a network to connect computers.
- Roberts, Davies, Kleinrock, and Baran get together to design ARPANET.

Evolution of Computer Networks

1969: ARPANET is developed.

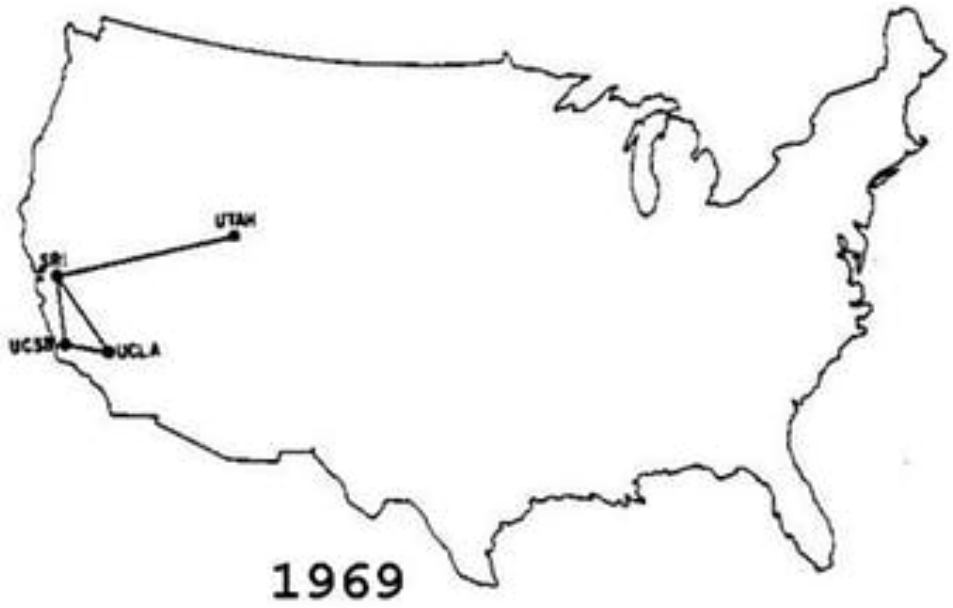

Evolution of Computer Networks

Early1970's: Vint Cerf develops NCP for transport and addressing.

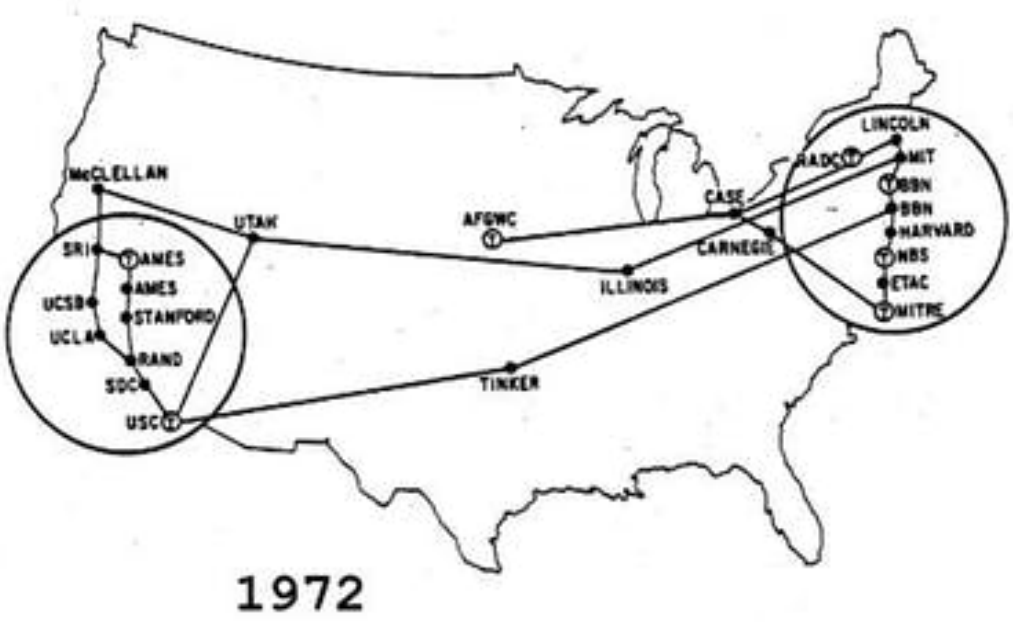

Evolution of Computer Networks

1973: European nodes added to ARPANET. The term Internet is born.

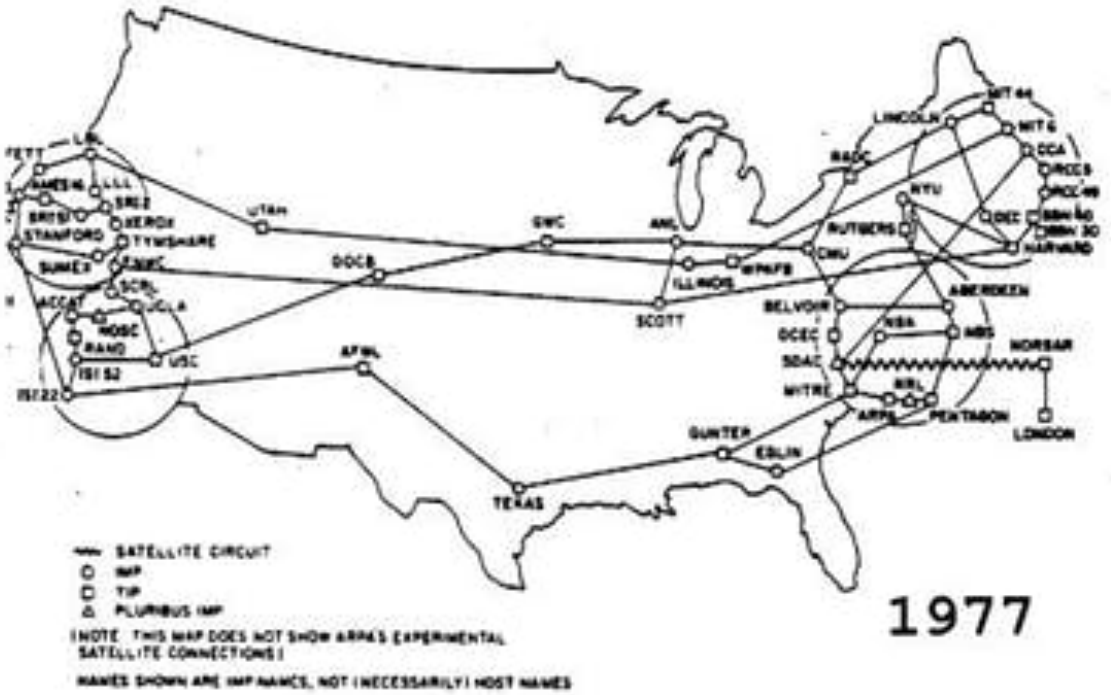

Evolution of Computer Networks

- mid-1970's: Vint Cerf and Bob Kahn develop TCP/IP, separating reliability from addressing.
- 1983: NCP becomes obsolete; all nodes switch to TCP/IP *(flag day).*
- Late 1970's: More scalable routing protocols was developed.
  - 1980: Link-state routing protocol was proposed.
- 1986: Series of congestion collapse; congestion control added to TCP.
- More interconnected networks emerge (Internet grows).
  - Early 1990's: BGP introduced for inter-domain routing.

Since then, for many years….

- No fundamental change in how we operate and use networks.
  - Distributed management of hardware switches.
  - Packet switching with store-and-forward design.
  - Endhost implements a TCP/IP stack in the kernel.

- Innovations in:
  - Transmission technology: wireless, cellular, more bandwidth.
  - Applications: HTTP, TLS, SSL, DNS.
  - Specific details: Congestion control algorithms, hierarchical addressing, etc.

But, changes emerged in the last decade...

This course tells the story of these changes.

Key enablers of the changes

- Increasing scale:
  - greater need to make networks *easier to manage*.
- More functionality:
  - greater need to make networks *more evolvable*.
- Commercialization:
  - greater emphasis on *performance*.

Key enablers of the changes

Emergence of large private networks.

In this course…

- What changes have been made to the networking infrastructure in the last decade?
- Why were the changes introduced?
- What do these changes enable?

In this course…

- Week 1: Review relevant concepts.
- Week 2: Historical perspective.
- Week 3-8: Switching infrastructure.
- Week 8-12: Endhost infrastructure.

In this course…

- Week 1: Review relevant concepts.
- Week 2: Historical perspective.
- Week 3-8: Switching infrastructure.
- Week 8-12: Endhost infrastructure.

Classical Papers

- End-to-end arguments in system design.
- The Design Philosophy of the DARPA Internet Protocols.
- Active networking.

In this course…

- Week 1: Review relevant concepts.
- Week 2: Historical perspective.
- Week 3-8: Switching infrastructure.
- Week 8-12: Endhost infrastructure.

Software-Defined Networking

- Philosophy
  - Individual switches focus on forwarding packets (data plane).
  - A centralized controller manages the switches (control plane).
- Enabling technology
  - OpenFlow, SDN controllers.
- Usecases
  - Google's software-defined WAN (B4), among others.

Limitation: switches can perform a limited set of actions, based on a fixed set of packet headers.

Programmable Data Plane

- Design and implementation of a software data plane.
- Programmable switching hardware
  - Reconfigurable match-action tables
- Language to program the hardware
  - P4
- Usecases
  - Networking functionality: telemetry, multicast, …
  - Others: caching, application-level load balancing, …
- Flexible packet scheduling.

In this course…

- Week 1: Review relevant concepts.
- Week 2: Historical perspective.
- Week 3-8: Switching infrastructure.
- Week 8-12: Endhost infrastructure.

Host Network Stack

- Standard kernel-based TCP stack is inefficient.
- User-space network stack (e.g. over DPDK).
- Accelerating kernel packet processing through eBPF.
- Offload network stack to hardware NIC (RDMA)

Smart NICs

- Software NIC to augment hardware.
- NICs equipped with FPGAs.
- NICs with multi-core SoC.

Systems built and used in industry

- Google's SNAP (unified host networking solution)
- Microsoft Azure's RDMA network
- Microsoft's AccelNet (network virtualization using FPGA-based smart NIC)

A recurring theme

Tussle between performance (high-speed) and programmability.

Logistics

Course Website

https://courses.engr.illinois.edu/ece598hpn/fa2023/

https://courses.engr.illinois.edu/cs598hpn/fa2023/

Office Hours

- Wednesdays 2:20pm-3:20pm, CSL 257.
- Meet by appointment: radhikam@illinois.edu

Reading assignments (30%)

- Each class: one full-length paper or two half-length papers.
- **Submit by 11:59pm the day before**:
  - 3-4 lines of summary.
  - 2 reasons why you would accept the paper.
  - 2 reasons why you would reject the paper.
  - One follow-up idea
    - Extension, weaker assumption, usecase.
- Submit via Google Forms (link on course website).
- 3 skips allowed (partial submission will be counted as a skip).
  - A submission that is late by more than 9hrs is counted as a skip.
  - Three late submissions (within 9hrs of deadline) counted as a skip.

Course Project (50%)

- Research style project in groups of up to two.
- Sept 8: I'll provide general pointers on project ideas, and setup meeting slots to discuss project proposals.
- Sept 29: Last date to discuss project proposal
- Oct 13: First progress report.
- Nov 3: Second progress report due.
- Last week of class: Final paper and presentation.

Warm-up assignments (10%)

- Three simple assignments to introduce different networking tools to you.
- Must be done individually.
- Team up with a partner to be each-other's TA.
- Submit a brief evaluation report for your partner.

Class Participation (10%)

- Actively engage in class discussions.
- A chance to present a paper of your choice towards the end of the course.
- A small quiz towards the end of the course score some *bonus* class participation points.

Questions?