

CIS 553: Networked Systems

Design Goals

January 27, 2021




Agenda

- Networking Basics ✓
- Layout of the Internet ✓
- Programming a switch ✓
- Design goals of the Internet ← NEXT
 - Taxonomy of networks
 - Design goals (Clark '88)
 - Layering and the end-to-end principle



Taxonomy of Communication Networks






Broadcast Networks



- Nodes share a common channel
- Everyone can hear everyone else

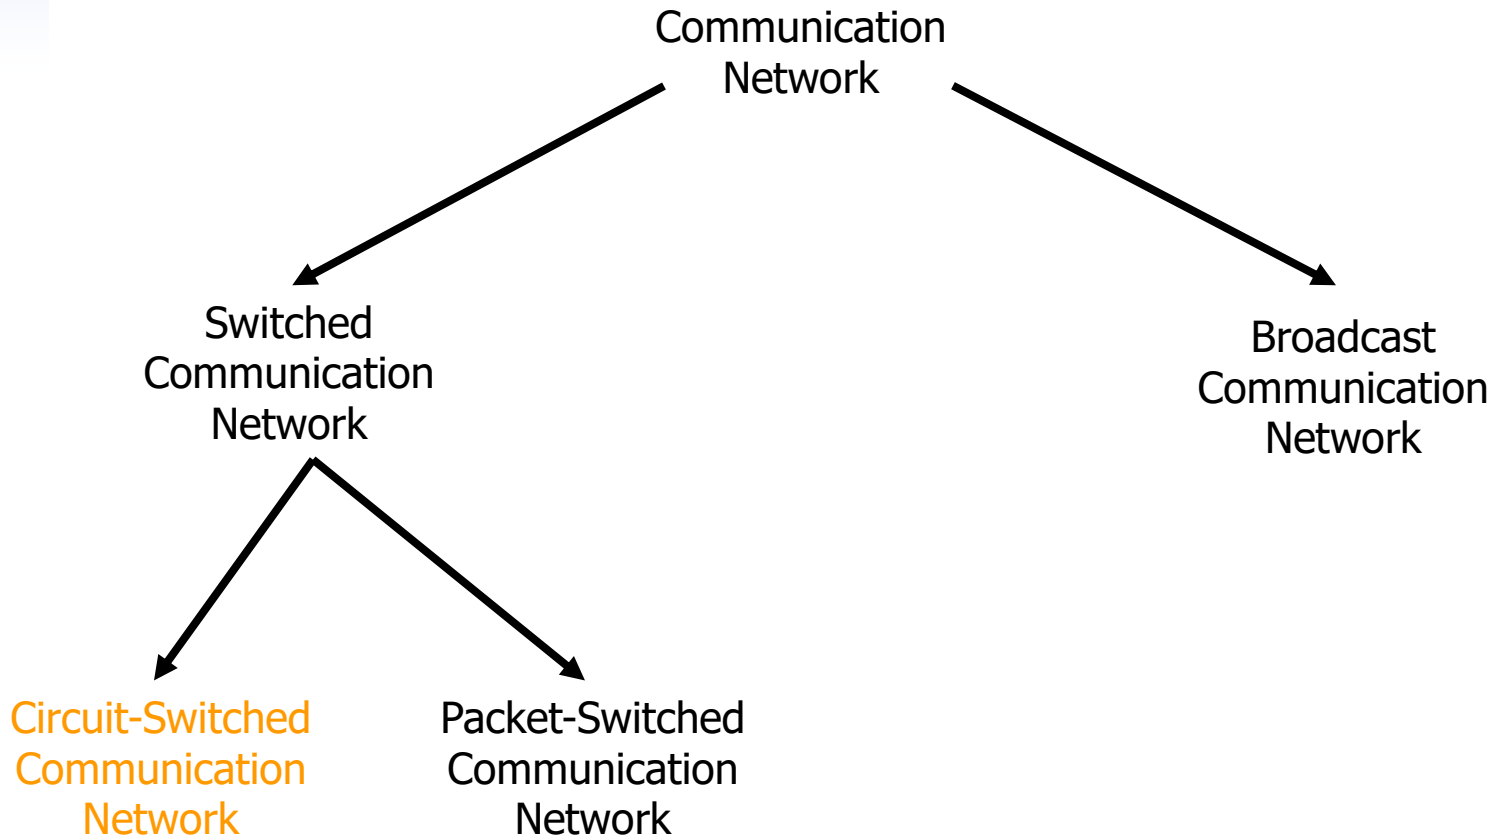
Switched Networks



- Information is transmitted to a small subset (usually one) of the nodes



Taxonomy of Communication Networks





Circuit Switching



- Three phases
 1. Circuit establishment (from source to destination)
 2. Data transfer
 3. Circuit termination

- Guaranteed capacity until circuit termination

- If circuit not available: “Busy signal”
 - Even if no data is sent!




Circuit Switching Timeline





Taxonomy of Communication Networks



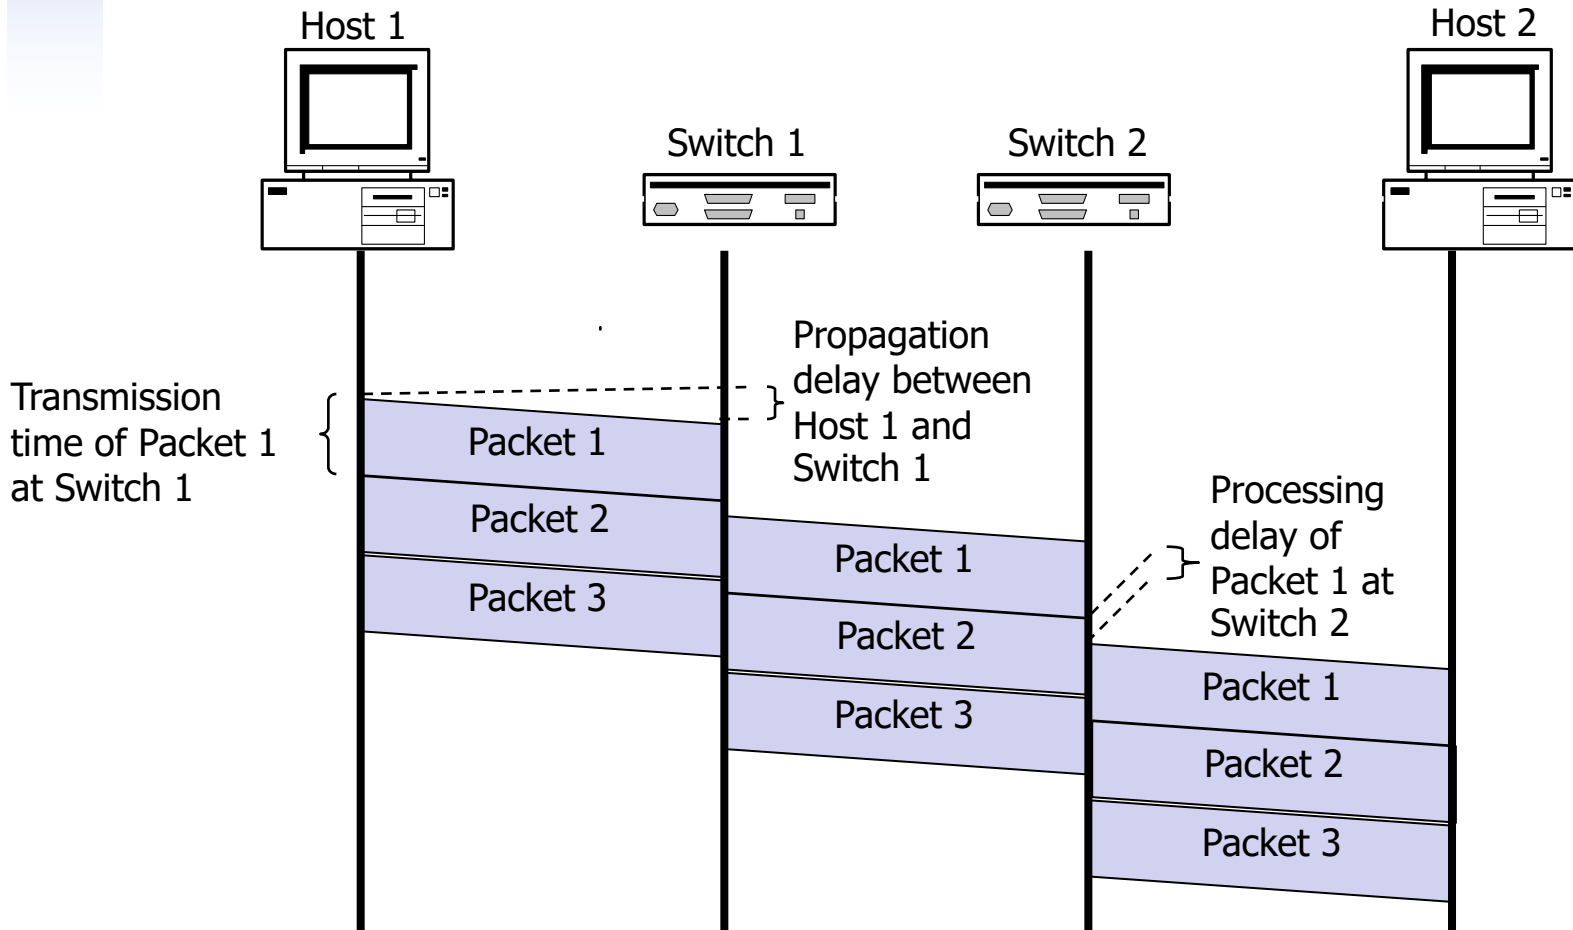


Datagram Packet Switching

- Each packet is independently switched
 - Each packet header contains destination address
 - Routing protocol is used to compute next hop
- No resources are pre-allocated (reserved) in advance



Packet Switching Timeline



What are the advantages/disadvantages of packet switching?



Packet switching vs Circuit switching

- Advantages:
 - No connection state required
 - Easy to recover from errors
 - Minimal network assumptions

- Disadvantages:
 - No guarantees
 - Slower when transmitting a lot of data




Agenda

- Design goals of the Internet ✓
 - Taxonomy of networks ✓
 - Design goals (Clark '88) ← NEXT
 - Layering and the end-to-end principle



Taxonomy of Communication Networks





The fundamental goal of the Internet

- An inter-network:
 - “technique for **multiplexed utilization** of **existing interconnected networks**”
- **Multiplexing** (**sharing**)
 - Shared use of a single communications channel
- **Existing networks** (**interconnection**)



The goals of the Internet

<http://ccr.sigcomm.org/archive/1995/jan95/ccr-9501-clark.pdf>

“technique for multiplexed utilization of existing interconnected networks”

- Second-level goals:
 - Resilience/survivability
 - Heterogeneity
 - Different types of services
 - Different types of networks
 - Distributed management
 - Cost effectiveness
 - Easy host attachment
 - Resource accountability

“This set of goals might seem to be nothing more than a checklist of all the desirable network features. It is important to understand that these goals are in order of importance, and an entirely different network architecture would result if the order were changed.”

These goals were prioritized for a military network.
Should priorities change as the network evolves?




Goal #2: Survivability

- If a device between you and your destination fails
- Goal: End-hosts should be able to continue communicating
 - Reconfigure to mask transient failures
 - Only stop working if there is a complete partition
- Key issue: how to deal with state
 - Ex: Information on location, authentication, what types of messages I'm allowed to send, how fast I'm allowed to send, etc.
 - **Replication vs Fate-sharing**



Replication

- The network keeps the state for you
- To protect from information loss, must be replicated
- Fault tolerance limited by the number of replicas
 - More failures will still wipe out the state
- Hard to engineer
 - End host and network need to deal with a lot of corner cases





Fate-sharing (chosen by the Internet)



- End hosts are responsible for their own state
- Lose information for a host if (and only if?) the host itself goes down, in which case it doesn't care
- Easy to engineer
 - Protects against any number of failures
 - Network devices are (mostly) stateless
 - More trust placed at end-hosts

Tradeoff: performance.
But mostly works.



Goal #3: Heterogeneous services/networks

- May require and provide different interfaces
 - Example: Voice vs Video
 - Example: Wi-Fi vs Backbone network
- Abstraction through **layering**
 - (We'll talk more about this later)
 - Can swap out individual layers, e.g., backbone for Wi-Fi
- “**Best effort**” service model
 - If you give me a message, I'll “try my best” to deliver it
 - Don't guarantee anything so that

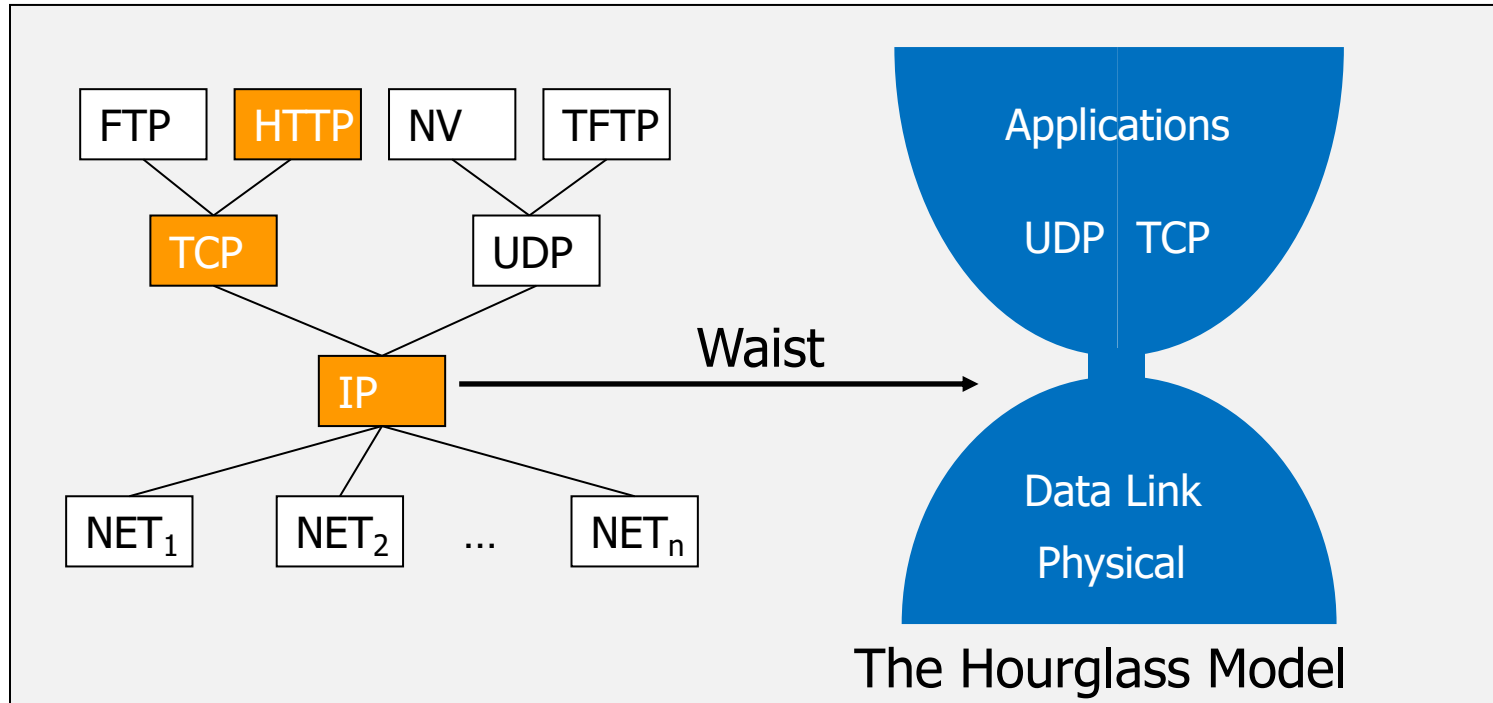
Tradeoff: No assumptions, no guarantees.
Also mostly works, but not in the way we expected.



The Internet Protocol Suite

Eventually, different networks had to agree on *something*
That something was IP*

*now IP/TCP/HTTP



Led to something called the "Curse of the Narrow Waist"



Goal #4: Distributed Management

- Today's Internet is decentralized
 - 10s of thousands of ISPs across almost every country
- Autonomous systems exchange routes with each other
- Each autonomous system can manage their Internal networks, connect to whomever they want

Allows for organic growth, scalable management.
Tradeoff: No one party has visibility/control.



No Owner, No Responsible Party

“Some of the most significant problems with the Internet today relate to lack of sufficient tools for distributed management, especially in the area of routing.”

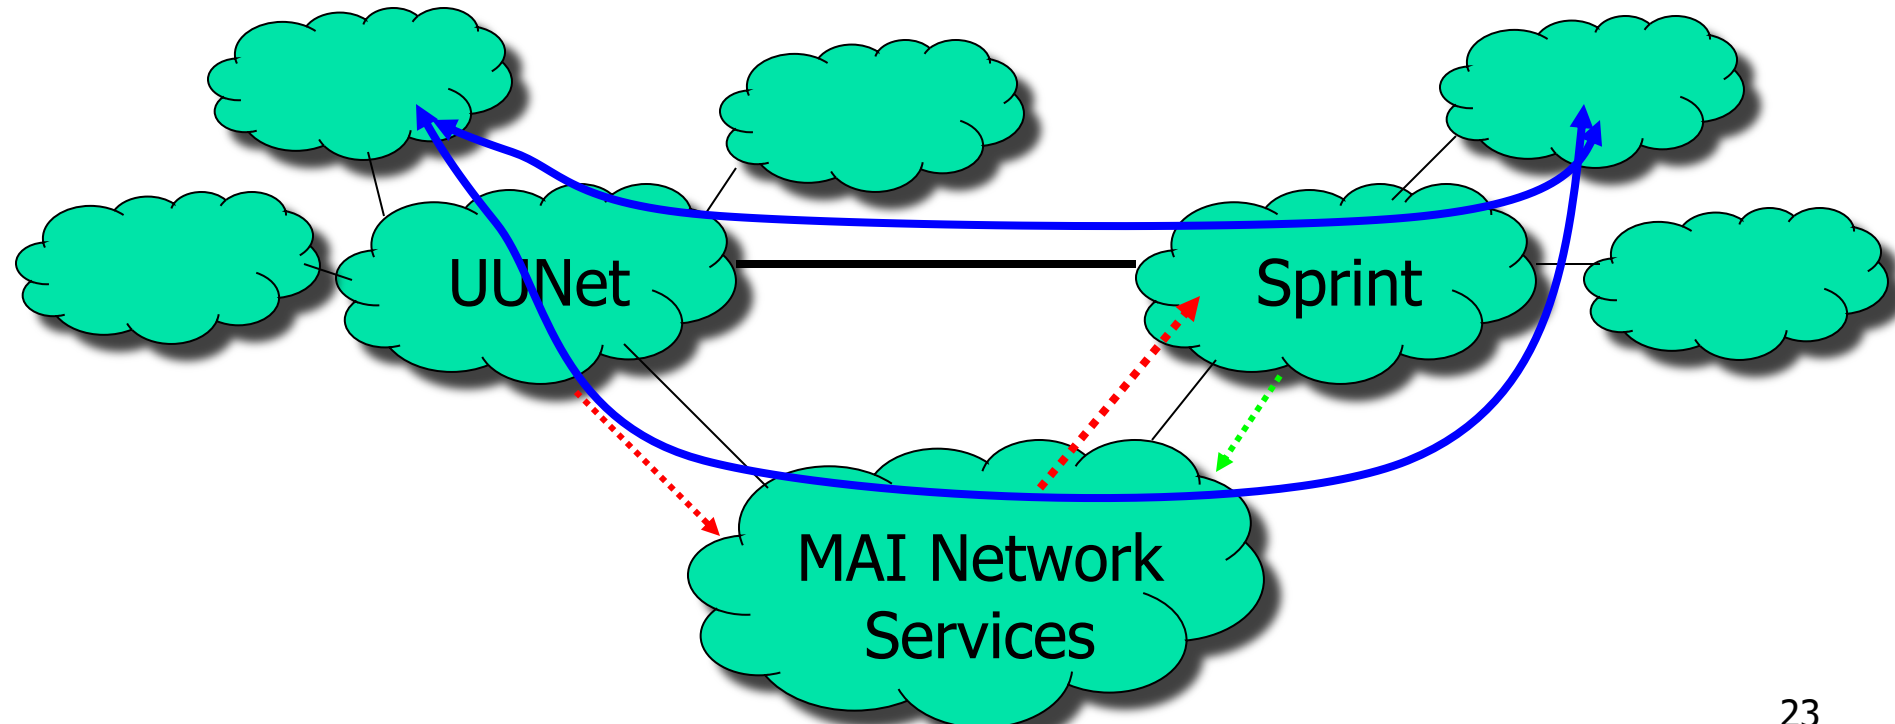
- Hard to figure out who/what’s causing a problem
- Worse yet, local actions have global effects...



Local Actions, Global Consequences

"...a glitch at a small ISP... triggered a **major outage in Internet access** across the country. The problem started when MAI Network Services...passed **bad router information** from one of its customers onto Sprint."

-- *news.com*, April 25, 1997





The goals that started
to fall by the wayside



Goal #5: Cost Effectiveness

- Many inefficiencies in the Internet
 - For small packets, headers introduce high overhead
 - End-to-end retransmission of lost packets leads to wasted bandwidth
- Arguably a good tradeoff as network speeds have grown rapidly
- Exception is wireless, where they try to correct for some of this



Goal #6: Ease of Attachment

- Anything with an IP stack can connect (hourglass model)
- Some amount of burden on end systems and programmers

- But once you get a working network stack...
 - A huge success!!



The “last” Goal #7: Resource Accountability

- Not prioritized at all
- Datagram networks make accounting tricky
 - The phone network has had an easier time figuring out billing
 - Payments/billing on the Internet is much less precise
- Convoluted business relationships
- Hot potato routing



What's Missing?

- Security
- Availability
- Accountability (the other kind)
- Support for disconnected/intermittent operation
- Mobility
- ...



Agenda

- Design goals of the Internet ✓
 - Taxonomy of networks ✓
 - Design goals (Clark '88) ✓
 - Layering and the end-to-end principle






Protocols

- An agreement between parties on how to communicate
- Defines the **syntax** of communication
- And **semantics**
 - “First a hello, then a request...”
- We will study many protocols
 - Exist at many levels (e.g., hardware and software)
 - Exist for many purposes (e.g., data transfer and routing control)
 - Defined by standards bodies like IETF, IEEE, ITU



Layering

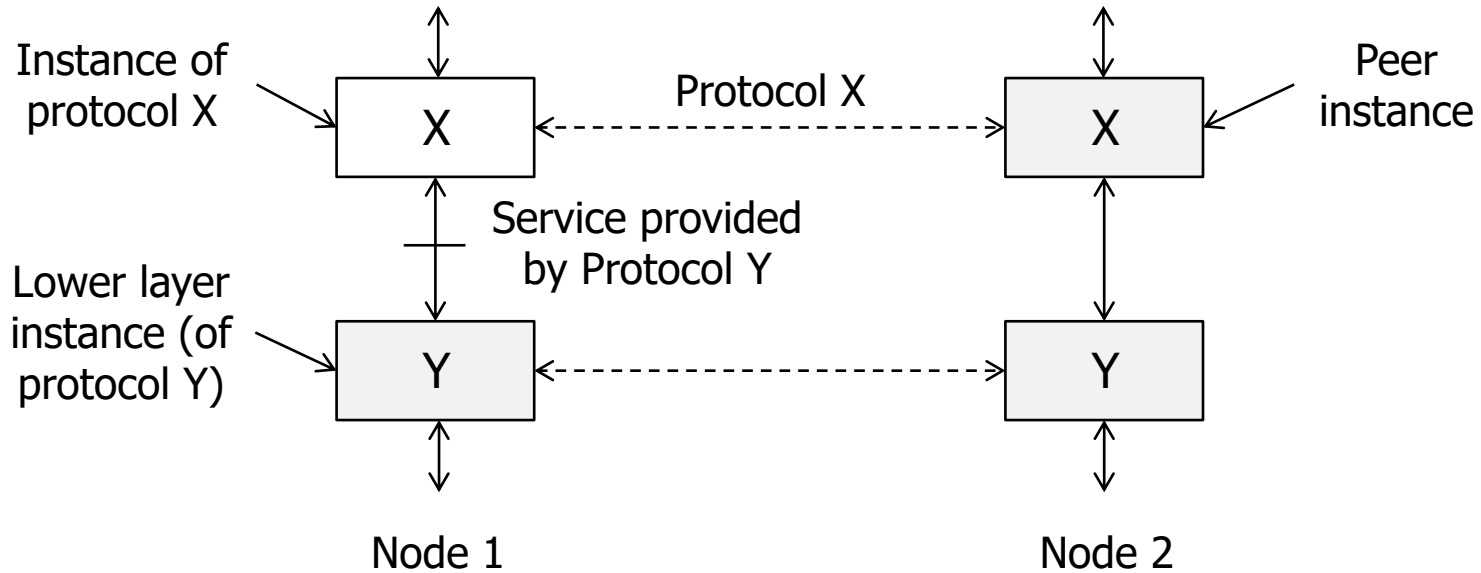
- The ability to mix and match protocols
- Set of protocols in use is called a **protocol stack**
- Ex:





Protocols and Layers


- Protocols are horizontal, layers are vertical



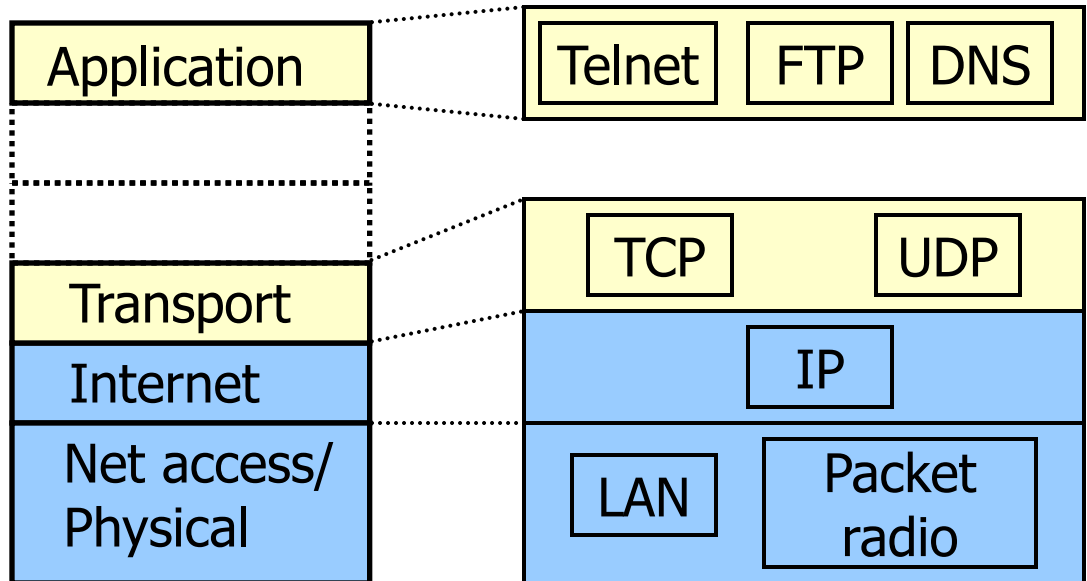


OSI vs. Internet

- OSI: conceptually define services, interfaces, protocols
- Internet: provide a successful implementation



OSI (formal)



Internet (informal)




Encapsulation

- **Encapsulation** is the mechanism used to effect protocol layering
 - Lower layer wraps higher layer content, adding its own information to make a new message for delivery
 - Like sending a letter in an envelope; postal service doesn't look inside




Encapsulation (Cont.)



- Message “on the wire” begins to look like an onion
 - Lower layers are outermost



Encapsulation (Cont.)





Encapsulation (Cont.)

- Normally draw message like this:
 - Each layer adds its own header



First bits on the
wire

Last
bits

- More involved in practice
 - Trailers as well as headers, encrypt/compress contents
 - Segmentation (divide long message) and reassembly