CS354: Machine Organization and Programming

Lecture 39
Friday the December 4th 2015

Section 2
Instructor: Leo Arulraj

Class Announcements

1. Special lecture on 12/11 and the review lecture on 12/14 will be handled by Jason.

2. Practice exam for Finals coming soon.

3. If you have not already started working on Programming Assignment 5, then do so now! Starting early helps you a lot.

Lecture Overview

1. Client Server Programming Model
2. Networks
3. The Global IP Internet
    1. IP Addresses
    2. Internet Domain Names
    3. Internet Connections

Client Server Programming Model

Application consists of one or more Client processes and one Server process.

Server manages some resource and it provides some service for its clients.

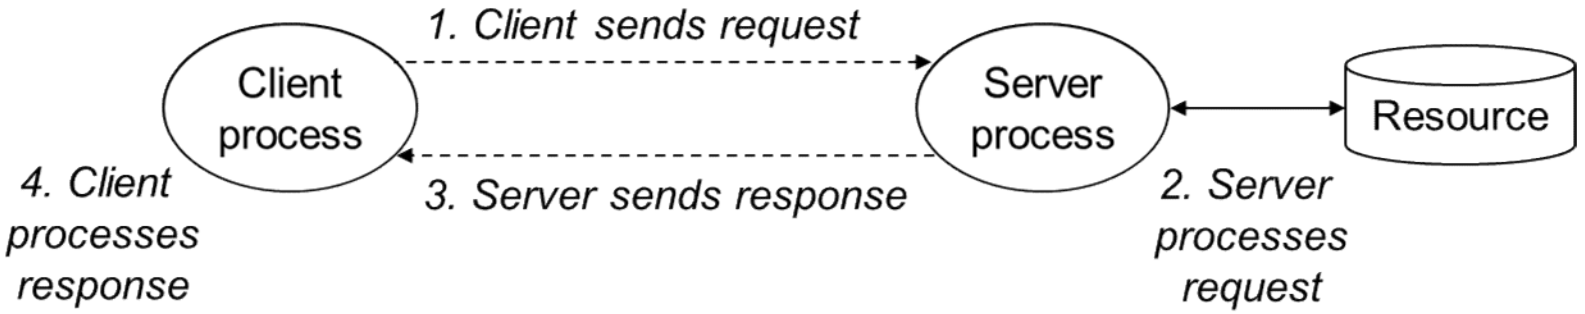

Networks

To the host, network is just another I/O device.

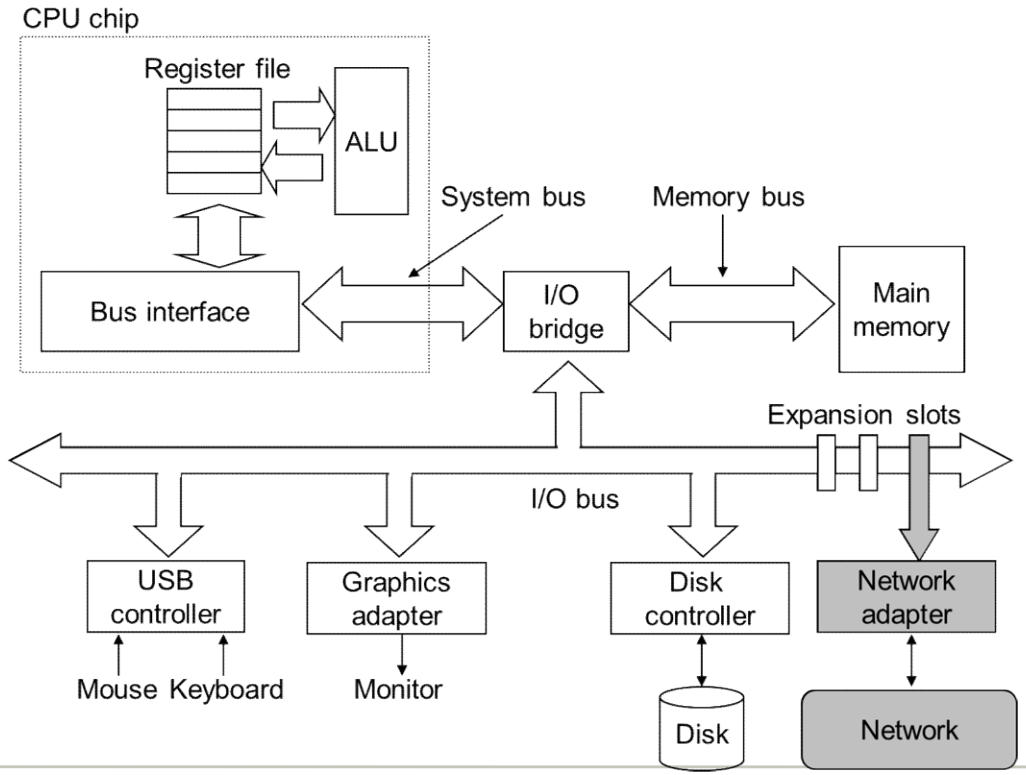

Network Hierarchy

Local Area Network (LAN): spans a building or a campus.

Most popular LAN technology is Ethernet which was developed at Xerox PARC and has proved to be remarkably resilient scaling from 3 Mb/s to 10 Gb/s.

Ethernet

Ethernet segments consists of some wires and a small box called a hub.

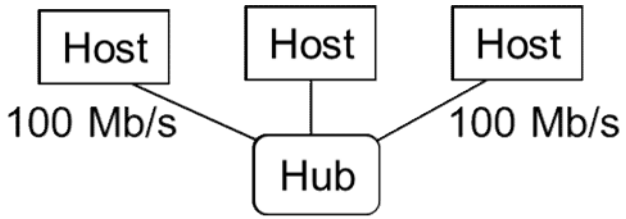

Each wire has same maximum bit bandwidth.

Hub copies every bit that it receives on each port to every other port. So, every host sees every bit.

Ethernet

Ethernet adapter consists of globally unique 48 bit address stored in non-volatile memory.

Each host can send a chunk of bits called an Ethernet frame to any other host on the segment.

Each Ethernet frame consists of header bits that identify source and destination of the frame, the frame length, followed by the payload of data bits.

Bridged Ethernets

Multiple Ethernet segments can be connected into larger LANs called bridged Ethernets, using a set of wires and small boxes called bridges.

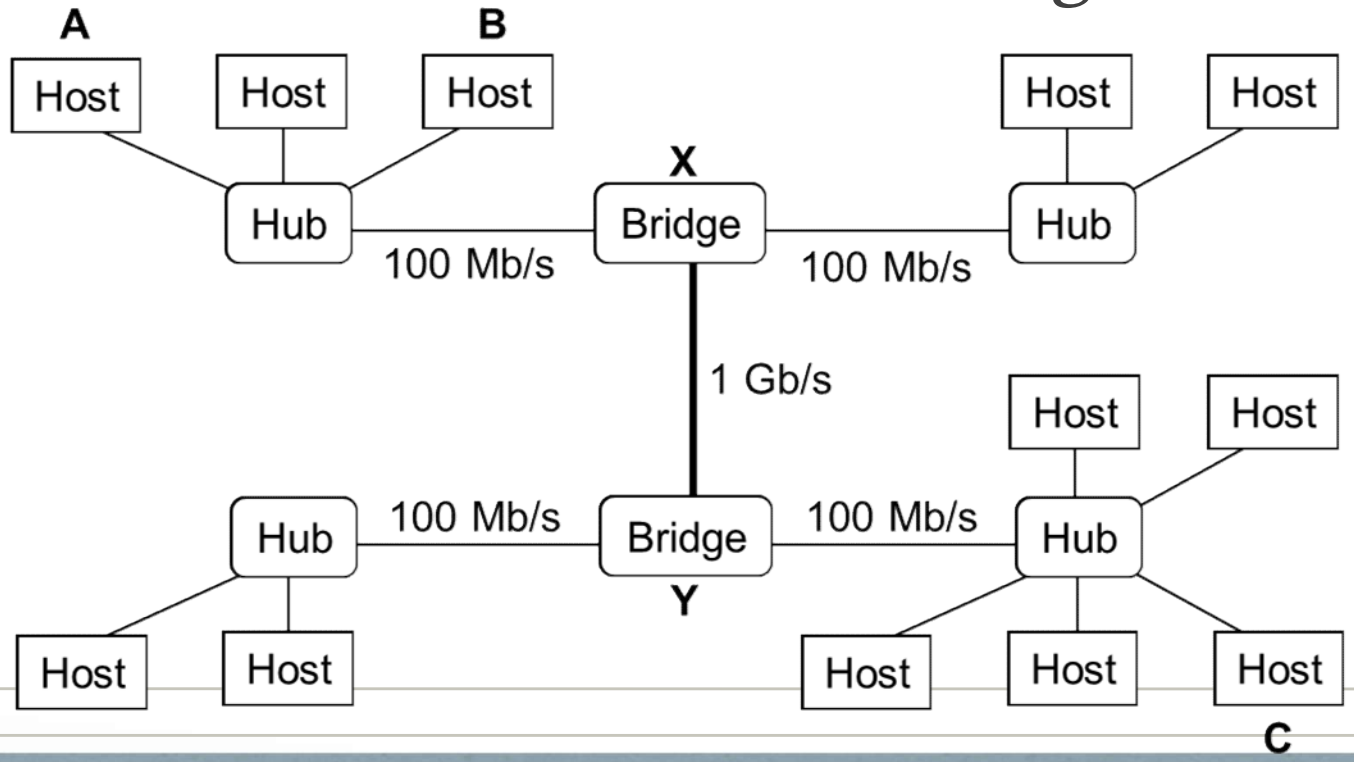

Bridged Ethernets

The maximum bit bandwidth of the wires connecting bridges and hubs can be different.

Bridges learn automatically over time which hosts are connected to which ports and selectively copy frames from one port to another only when necessary.

E.g. Bridge discards A->B frame but copies A->C frame only to the port connected to C's segment.

Hubs, Bridges, Switches and Routers

Network layer: packets (routers)
Link layer: frames (bridges and switches)
Physical layer: electrical signals (hubs)

Details of the OSI network model later.

Hubs:

- One large shared link

  Each bit is sent everywhere. So, aggregate throughput is limited.

- Cannot support multiple LAN technologies

  Does not buffer or interpret frames. So, can't interconnect between different rates or formats. E.g., 10 Mbps Ethernet and 100 Mbps Ethernet

Hubs, Bridges, Switches and Routers

Bridges:

Connects two or more LANs at the link layer

- Extracts destination address from the frame
- Looks up the destination in a table
- Forwards the frame to the appropriate LAN segment

Each segment can carry its own traffic

Hubs, Bridges, Switches and Routers

Switches (a.k.a. intelligent hubs):

Typically connects individual computer: A switch is essentially the same as a bridge … though typically used to connect hosts, not LANs.

Like bridges, support concurrent communication: Host A can talk to C, while B talks to D

Dedicated access: Host has direct connection to the switch… rather than a shared LAN connection

Full duplex: Each connection can send in both directions, Host sending to switch, and host receiving from switch.

Hubs, Bridges, Switches and Routers

Routers (a.k.a layer 3 switches):

- Routers are far more feature rich as compared to switches.
- Routers maintain routing table for data forwarding.
- Routers can be slow compared to switching because of the fact that routing table lookup time is considerably high if not implemented in hardware.
- Routers have lesser port densities as compared to switches.
- Routers are usually used as a forwarding network elements in Wide Area Networks.

internet

Multiple LANs and WANs are connected using Routers to form an internet.

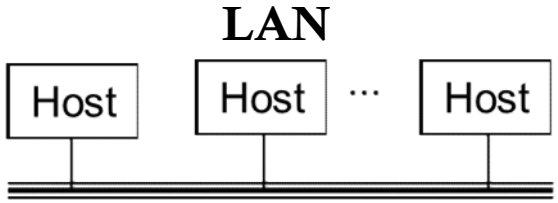

An internet

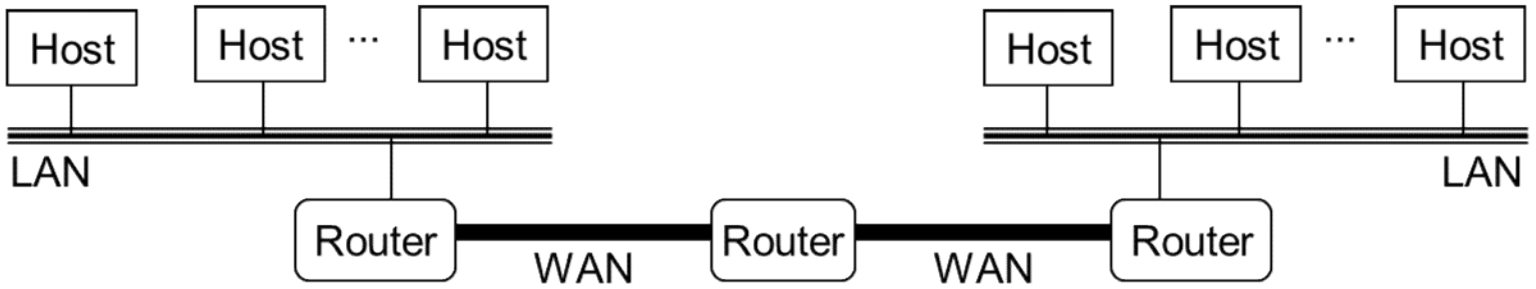

internet

A Wide Area Network (WAN) is connected by Routers using high speed point-to-point phone connections. WANs span large geographical areas.

Routers are used to build internets from arbitrary collection of LANs and WANs that use radically different and incompatible technologies.

A *protocol software* running on each host and router smoothens out the differences between different networks.

internet

Protocol software provides two basic capabilities:

Naming Scheme: provides a uniform format for host addresses.

Delivery Mechanism: defines a uniform way to bundle up data bits into discrete chunks of packets.

internet

Packets consist of:

Header: contains packet size, addresses of source and destination hosts.

Payload: contains the data bits sent from the source host.

internet protocol

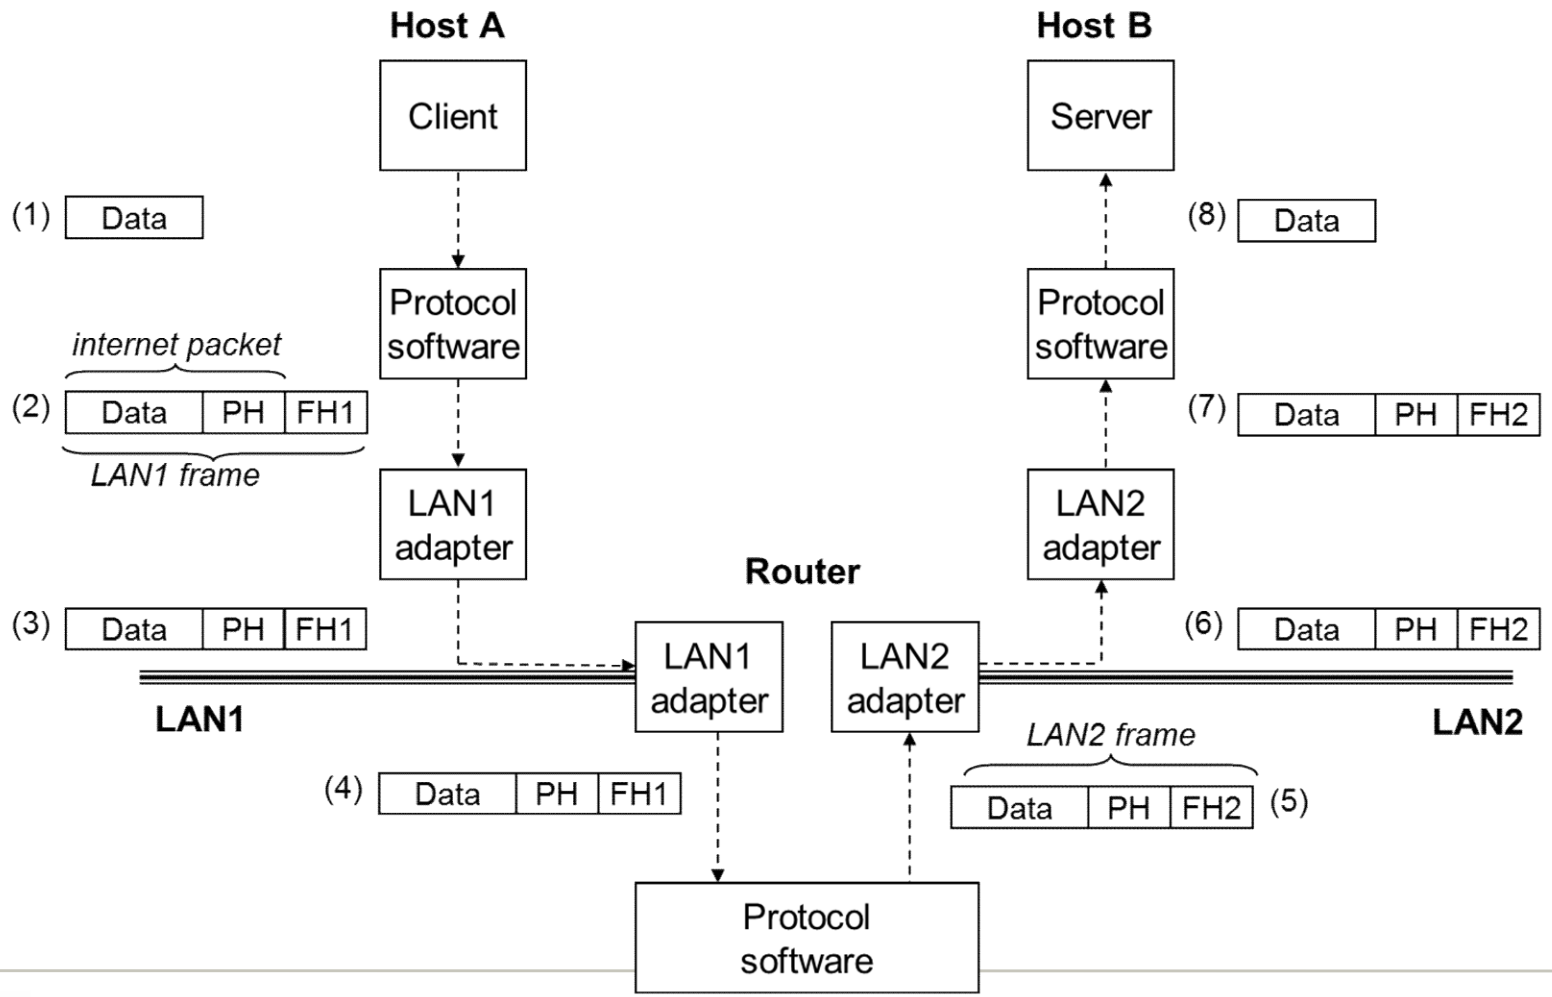

an internet vs The Internet

internet is the general concept of connecting hosts across LANs and WANs.

The Internet is the specific implementation – the global IP Internet.

The Global IP Internet

Most famous and successful implementation of an internet.

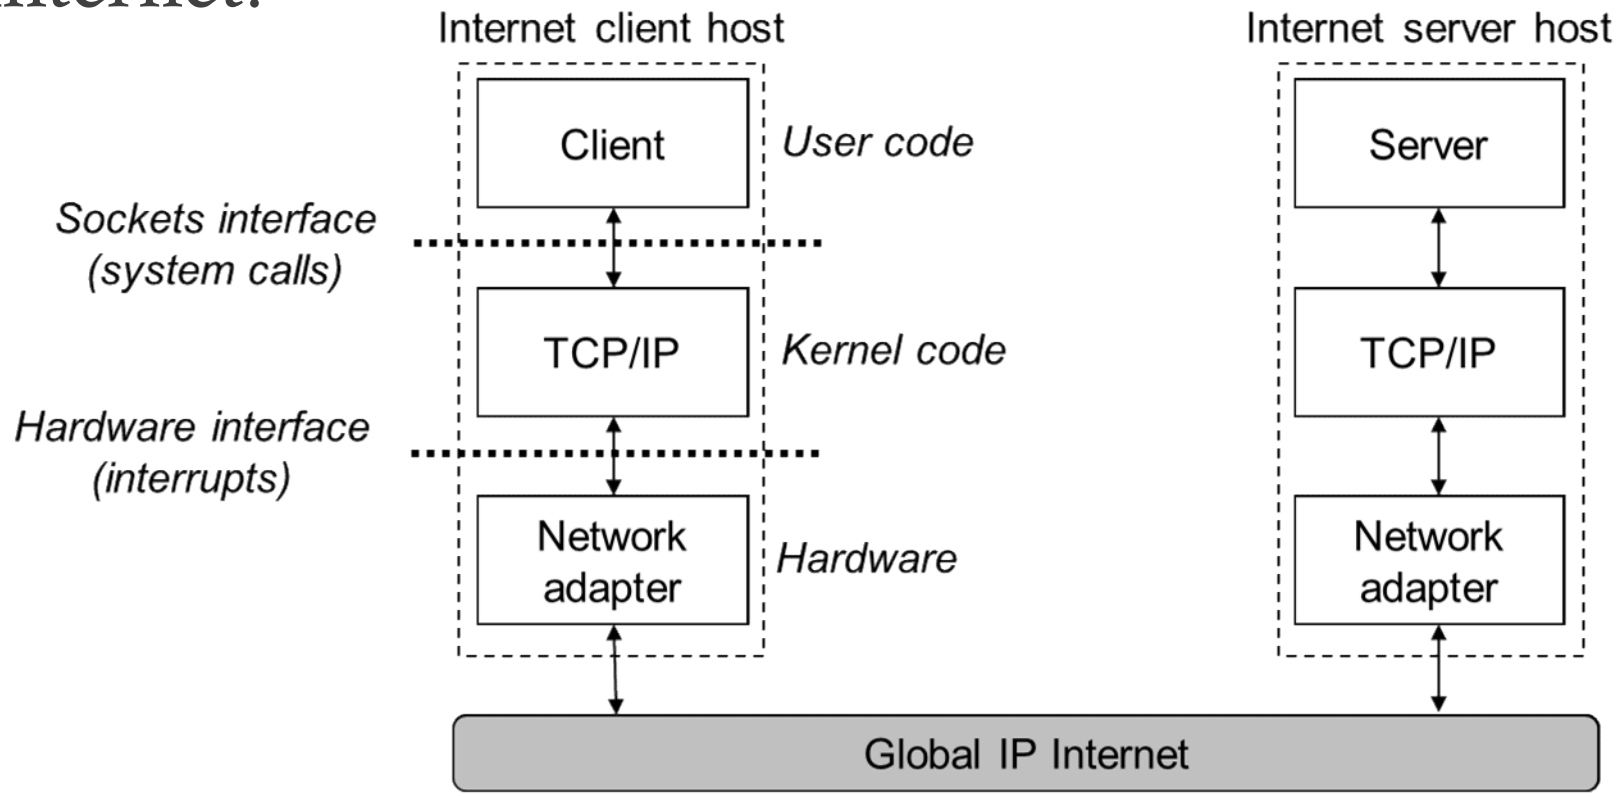

The Global IP Internet

Each Internet host runs software that runs TCP/IP protocol (Transmission Control Protocol/Internet Protocol).

TCP/IP is a family of protocols consisting of: IP, UDP, TCP

TCP/IP

1. IP: provides basic naming scheme and delivery mechanism that can send packets/datagrams from one host to another. Is Unreliable.
2. UDP (Unreliable Datagram Protocol): extends IP to provide full duplex (bidirectional) connections between processes.
3. TCP (Transmission Control Protocol): a complex protocol that builds on IP to provide reliable full duplex connections between processes.

Programmer's perspective of the Global IP Internet

The set of hosts is mapped to a set of 32-bit IP addresses.

The set of IP addresses is mapped to a set of identifiers called Internet Domain Names.

A process on one Internet host can communicate with a process on any other Internet host over a connection.

IP Addresses

IP addresses typically presented to humans using a *dotted decimal notation*. E.g. 128.2.194.242

TCP/IP defines a uniform network byte order (big endian byte order) for any integer data item such as an IP address.

Convert between network and host byte order using htonl(), htons(), ntohl(), ntohs().

IP Addresses

linux> hostname –I

inet_aton

inet_ntoa

Internet address structure:
struct in_addr{ unsigned int s_addr;};

IPv4 vs IPv6

- The original Internet Protocol, with its 32-bit addresses, is known as *Internet Protocol Version 4* (IPv4)
- 1996: Internet Engineering Task Force (IETF) introduced *Internet Protocol Version 6* (IPv6) with 128-bit addresses
  - Intended as the successor to IPv4
- As of 2015, vast majority of Internet traffic still carried by IPv4
  - Only 4% of users access Google services using IPv6.
- We will focus on IPv4, but will show you how to write networking code that is protocol-independent.

Internet Domain Names

Since large integers are difficult to remember, a human friendly *domain name* is used along with a mechanism that maps domain names to IP addresses.

A domain name is a sequence of words separated by periods. E.g. cs.wisc.edu

The set of domain names form a hierarchy and each domain name encodes its position in the hierarchy.