## Systemprogrammering 2006 Föreläsning 8 Network Programming

#### Topics

Client-server programming model

Networks

Internetworks

Global IP Internet

- IP addresses
- Domain names
- Connections

Socket interface

Writing clients and servers

## Hardware Org of a Network Host



## **A Client-Server Transaction**

#### Every network application is based on the client-server model:

- A server process and one or more client processes
- Server manages some resource.
- Server provides service by manipulating resource for clients.



Note: clients and servers are processes running on hosts (can be the same or different hosts).

F8 – 2 –

Systemprogrammering 2007

## **Computer Networks**

#### A network is a hierarchical system of boxes and wires organized by geographical proximity

- LAN (local area network) spans a building or campus.
   Ethernet is most prominent example.
- WAN (wide-area network) spans country or world.
  - Typically high-speed point-to-point phone lines.

## An *internetwork (internet)* is an interconnected set of networks.

The Gobal IP Internet (uppercase "I") is the most famous example of an internet (lowercase "i")

#### Let's see how we would build an internet from the ground up.

## Lowest Level: Ethernet Segment

Ethernet segment consists of a collection of *hosts* connected by wires (twisted pairs) to a *hub*.

Spans room or floor in a building.



#### Operation

- Each Ethernet adapter has a unique 48-bit Mac address.
- Hosts send bits to any other host in chunks called frames.
- Hub slavishly copies each bit from each port to every other port.
  - Every host sees every bit.
- Now hubs are usually replaced by switches which send data only to the right port

ports

F8 – 5 –

Systemprogrammering 2007

## **Conceptual View of LANs**

For simplicity, hubs, bridges, and wires are often shown as a collection of hosts attached to a single wire:



## Next Level: Bridged Ethernet Segment

Spans building or campus.

Bridges cleverly learn which hosts are reachable from which ports and then selectively copy frames from port to port.



### **Next Level: internets**

Multiple incompatible LANs can be physically connected by specialized computers called *routers*.

The connected networks are called an internet.



LAN 1 and LAN 2 might be completely different, totally incompatible LANs (e.g., Ethernet and ATM)

### The Notion of an internet Protocol

- How is it possible to send bits across incompatible LANs and WANs?
- Solution: *protocol software* running on each host and router smoothes out the differences between the different networks.

Implements an *internet protocol* (i.e., set of rules) that governs how hosts and routers should cooperate when they transfer data from network to network.

TCP/IP is the protocol for the global IP Internet.

Systemprogrammering 2007

## **Transferring Data Over an internet**



## What Does an internet Protocol Do?

#### 1. Provides a naming scheme

- An internet protocol defines a uniform format for host addresses.
- Each host (and router) is assigned at least one of these internet addresses that uniquely identifies it.

#### 2. Provides a delivery mechanism

- An internet protocol defines a standard transfer unit (packet)
- Packet consists of header and payload
  - Header: contains info such as packet size, source and destination addresses.
  - Payload: contains data bits sent from source host.

F8 – 10 –

Systemprogrammering 2007

## **Other Issues**

#### We are glossing over a number of important questions:

- What if different networks have different maximum frame sizes? (segmentation)
- How do routers know where to forward frames?
- How are routers informed when the network topology changes?
- What if packets get lost?

## These (and other) questions are addressed by the area of systems known as *computer networking*.

## **Global IP Internet**

#### Most famous example of an internet.

#### **Based on the TCP/IP protocol family**

- IP (Internet protocol) (v. 4, v. 6):
  - Provides basic naming scheme and unreliable delivery capability of packets (datagrams) from host-to-host.
- UDP (Unreliable Datagram Protocol)
  - Uses IP to provide unreliable datagram delivery from process-toprocess.
- TCP (Transmission Control Protocol)
  - Uses IP to provide reliable byte streams from process-to-process over connections.

## Accessed via a mix of Unix file VO and functions from the sockets interface.

F8 – 13 –

Systemprogrammering 2007

## A Programmer's View of the Internet

#### 1. Hosts are mapped to a set of 32-bit *IP addresses* (v. 4). 128.2.203.179

- 128.2.203.179
- IP v. 6 has 128-bit addresses, but is, so far, rarely used
- 2. The set of IP addresses is mapped to a set of identifiers called Internet *domain names*.
  - 128.2.203.179 is mapped to www.cs.cmu.edu
- 3. A process on one Internet host can communicate with a process on another Internet host over a *connection*.

## Hardware and Software Org of an Internet Application



## 1. IP Addresses

#### 32-bit IP addresses are stored in an IP address struct

- IP addresses are always stored in memory in network byte order (big-endian byte order)
- True in general for any integer transferred in a packet header from one machine to another.
  - E.g., the port number used to identify an Internet connection.

```
/* Internet address structure */
struct in_addr {
    unsigned int s_addr; /* network byte order (big-endian) */
};
```

#### Handy network byte-order conversion functions:

htonl: convert long int from host to network byte order.

htons: convert short int from host to network byte order.

ntohl: convert long int from network to host byte order.

ntohs: convert short int from network to host byte order. F8-16Systemprogrammering 2007

## **Dotted Decimal Notation**

By convention, each byte in a 32-bit IP address is represented by its decimal value and separated by a period

• IP address 0x8002C2F2 = 128.2.194.242

## Functions for converting between binary IP addresses and dotted decimal strings:

- inet\_aton: converts a dotted decimal string to an IP address in network byte order.
- inet\_ntoa: converts an IP address in network by order to its corresponding dotted decimal string.
- "n" denotes network representation. "a" denotes application

representation.

Systemprogrammering 2007

## Domain Naming System (DNS)

# The Internet maintains a mapping between IP addresses and domain names in a huge worldwide distributed database called *DNS*.

Conceptually, programmers can view the DNS database as a collection of millions of *host entry structures*:

```
/* DNS host entry structure */
struct hostent {
    char *h_name; /* official domain name of host */
    char *h_aliases; /* null-terminated array of domain names */
    int h_addrtype; /* host address type (AF_INET) */
    int h_length; /* length of an address, in bytes */
    char **h_addr_list; /* null-terminated array of in_addr structs */
};
```

#### Functions for retrieving host entries from DNS:

gethostbyname: query key is a DNS domain name.

gethostbyaddr: query key is an IP address.

Systemprogrammering 2007



## **Properties of DNS Host Entries**

Each host entry is an equivalence class of domain names and IP addresses.

Each host has a locally defined domain name localhost which

always maps to the loopback address 127.0.0.1

#### Different kinds of mappings are possible:

Simple case: 1-1 mapping between domain name and IP addr:

• kittyhawk.cmcl.cs.cmu.edu maps to 128.2.194.242

- Multiple domain names mapped to the same IP address:
  - eecs.mit.edu and cs.mit.edu both map to 18.62.1.6
- Multiple domain names mapped to multiple IP addresses:

• aol.com and www.aol.com map to multiple IP addrs.

Some valid domain names don't map to any IP address:

• for example: cmcl.cs.cmu.edu

Systemprogrammering 2007

### A Program That Queries DNS



#### Internet Connections Client socket address Server socket address 128.2.194.242:51213 208.216.181.15:80 Server Client Connection socket pair (port 80) (128.2.194.242:51213, 208.216.181.15:80) Client host address Server host address 128.2.194.242 208.216.181.15 Note: 80 is a well-known port Note: 51213 is an ephemeral port allocated associated with Web servers by the kernel Systemprogrammering 2007

## 3. Internet Connections

#### Clients and servers communicate by sending streams of bytes over connections:

Point-to-point, full-duplex (2-way communication), and reliable.

#### A socket is an endpoint of a connection

Socket address is an IPaddress:port pair

#### A port is a 16-bit integer that identifies a process:

- Ephemeral port: Assigned automatically on client when client makes a connection request
- Well-known port: Associated with some service provided by a server (e.g., port 80 is associated with Web servers)

#### A connection is uniquely identified by the socket addresses

#### of its endpoints (socket pair)

(cliaddr:cliport, servaddr:servport)

F8 – 22 –

Systemprogrammering 2007

## Clients

#### **Examples of client programs**

Web browsers, ftp, telnet, ssh

#### How does a client find the server?

- The IP address in the server socket address identifies the host (more precisely, an adapter on the host)
- The (well-known) port in the server socket address identifies the service, and thus implicitly identifies the server process that performs that service.
- Examples of well know ports
  - Port 7: Echo server
  - Port 23: Telnet server
  - Port 25: Mail server
  - Port 80: Web server

F8 – 23 –

### **Using Ports to Identify Services**



## **Server Examples**

#### Web server (port 80)

- Resource: files/compute cycles (CGI programs)
- Service: retrieves files and runs CGI programs on behalf of the client

#### FTP server (20, 21)

- Resource: files
- Service: stores and retrieve files

#### **Telnet server (23)**

- Resource: terminal
- Service: proxies a terminal on the server machine

#### Mail server (25)

- Resource: email "spool" file
- Service: stores mail messages in spool file

#### F8 – 27 –

Systemprogrammering 2007

See /etc/services for a

comprehensive list of the

machine.

services available on a Unix

## Servers

#### Servers are long-running processes (daemons).

- Created at boot-time (typically) by the init process (process 1)
- Run continuously until the machine is turned off.

#### Each server waits for requests to arrive on a well-known port

#### associated with a particular service.

- Port 7: echo server
- Port 23: telnet server
- Port 25: mail server
- Port 80: HTTP server

## A machine that runs a server process is also often referred to as a "server."

F8 – 26 –

Systemprogrammering 2007

## **Sockets Interface**

Created in the early 80's as part of the original Berkeley distribution of Unix that contained an early version of the Internet protocols.

Provides a user-level interface to the network.

Underlying basis for all Internet applications.

Based on client/server programming model.



## **Socket Address Structures**

#### Generic socket address:

- For address arguments to connect, bind, and accept.
- Necessary only because C did not have generic (void \*) pointers when the sockets interface was designed.

<pre>struct sockaddr {</pre>		
unsigned short char	sa_family; sa_data[14];	<pre>/* protocol family */ /* address data. */</pre>
};		

#### Internet-specific socket address:

Must cast (sockaddr_	_in *) t <b>o</b> (sockaddr	*) for connect, bind,
----------------------	-----------------------------	-----------------------

and accept.

```
struct sockaddr_in {
    unsigned short sin_family; /* address family (always AF_INET) */
    unsigned short sin_port; /* port num in network byte order */
    struct in_addr sin_addr; /* IP addr in network byte order */
    unsigned char sin_zero[8]; /* pad to sizeof(struct sockaddr) */
};
F8-31-
    Systemprogrammering 2007
```

## Sockets

#### What is a socket?

- To the kernel, a socket is an endpoint of communication.
- To an application, a socket is a file descriptor that lets the application read/write from/to the network.
  - Remember: All Unix I/O devices, including networks, are modeled as files.

Clients and servers communicate with each by reading from and writing to socket descriptors.

## The main distinction between regular file I/O and socket I/O is how the application "opens" the socket descriptors.

F8 - 30 -

F8 - 32 -

Systemprogrammering 2007

## **Echo Client Main Routine**



### Echo Client: open\_clientfd



F8 – 33 –

Systemprogrammering 2007

# Echo Client: open\_clientfd (gethostbyname)

#### The client then builds the server's Internet address.

<pre>int clientfd; struct hostent *hp; struct sockaddr_in serveraddr;</pre>	<pre>/* socket descriptor */ /* DNS host entry */ /* server's IP address */</pre>			
<pre>/* fill in the server's IP address and port */ if ((hp = gethostbyname(hostname)) == NULL)     return -2; /* check h_errno for cause of error */ bzero((char *) &amp;serveraddr, sizeof(serveraddr)); serveraddr.sin_family = AF_INET; bcopy((char *)hp-&gt;h_addr,</pre>				
serveraddr.sin_port = incons(por	,			

# Echo Client: open\_clientfd (socket)

#### socket creates a socket descriptor on the client.

- AF\_INET: indicates that the socket is associated with Internet protocols.
- SOCK\_STREAM: selects a reliable byte stream connection.

int clientfd; /\* socket descriptor \*/

if ((clientfd = socket(AF\_INET, SOCK\_STREAM, 0)) < 0)
 return -1; /\* check errno for cause of error \*/</pre>

... (more)

F8 – 34 –

Systemprogrammering 2007

# Echo Client: open\_clientfd (connect)

#### Finally the client creates a connection with the server.

- Client process suspends (blocks) until the connection is created.
- After resuming, the client is ready to begin exchanging messages with the server via Unix I/O calls on descriptor sockfd.

int clientfd;	<pre>/* socket descriptor */</pre>	
<pre>struct sockaddr_in serveraddr;</pre>	/* server address */	
typedef struct sockaddr SA;	/* generic sockaddr */	
<pre>/* Establish a connection with the server */ if (connect(clientfd, (SA *)&amp;serveraddr, sizeof(serveraddr)) &lt; 0) return -1;</pre>		
return clientfd;		

Systemprogrammering 2007

F8 – 36 –

### Echo Server: Main Routine

<pre>int main(int argc, char **argv) {     int listenfd, connfd, port, clientlen;     struct sockaddr in clientaddr;</pre>	
struct hostent *hp:	
char *haddrp;	
<pre>port = atoi(argv[1]); /* the server listens on a port ;</pre>	passed
<pre>listenfd = open_listenfd(port);</pre>	
while (1) {	
clientlen = sizeof(clientaddr);	
<pre>connfd = Accept(listenfd, (SA *)&amp;clientaddr, &amp;clie</pre>	ntlen);
hp = Gethostbyaddr((const char *)&clientaddr.sin_a	ddr.s_addr,
sizeof(clientaddr.sin_addr.s_addr)	, AF_INET);
haddrp = inet_ntoa(clientaddr.sin_addr);	
<pre>printf("server connected to %s (%s)\n", hp-&gt;h_name echo(connfd):</pre>	, haddrp);
Close(connfd):	
}	

1 F8 – 37 –

Systemprogrammering 2007

Systemprogrammering 2007

## Echo Server: open\_listenfd (cont)

. . .

```
/* Listenfd will be an endpoint for all requests to port
     on any IP address for this host */
  bzero((char *) &serveraddr, sizeof(serveraddr));
  serveraddr.sin_family = AF_INET;
  serveraddr.sin_addr.s_addr = htonl(INADDR_ANY);
  serveraddr.sin_port = htons((unsigned short)port);
  if (bind(listenfd, (SA *)&serveraddr, sizeof(serveraddr)) < 0)</pre>
      return -1;
  /* Make it a listening socket ready to accept
     connection requests */
  if (listen(listenfd, LISTENQ) < 0)</pre>
```

```
return -1;
```

return listenfd;

## Echo Server: open\_listenfd

int open\_listenfd(int port) int listenfd, optval=1; struct sockaddr\_in serveraddr; /\* Create a socket descriptor \*/

if ((listenfd = socket(AF INET, SOCK STREAM, 0)) < 0) return -1;

/\* Eliminates "Address already in use" error from bind. \*/ if (setsockopt(listenfd, SOL\_SOCKET, SO\_REUSEADDR, (const void \*)&optval , sizeof(int)) < 0)</pre> return -1;

... (more)

F8 - 38 -

£

Systemprogrammering 2007

### Echo Server: open\_listenfd (socket)

#### socket creates a socket descriptor on the server.

- AF\_INET: indicates that the socket is associated with Internet protocols.
- SOCK\_STREAM: selects a reliable byte stream connection.

int listenfd; /\* listening socket descriptor \*/

- /\* Create a socket descriptor \*/
- if ((listenfd = socket(AF\_INET, SOCK\_STREAM, 0)) < 0)</pre> return -1;

F8 - 40 -

# Echo Server: open\_listenfd (setsockopt)

#### The socket can be given some attributes.

## Handy trick that allows us to rerun the server immediately after we kill it.

- Otherwise we would have to wait about 15 secs.
- Eliminates "Address already in use" error from bind().

## Strongly suggest you do this for all your servers to simplify debugging.

```
F8 – 41 –
```

. . .

Systemprogrammering 2007

# Echo Server: open\_listenfd (bind)

## bind associates the socket with the socket address we just created.

<pre>int listenfd; struct sockaddr_in serveraddr;</pre>	/* listening socket */ /* server's socket addr */
•••	
<pre>/* listenid will be an endpo     on any IP address for thi</pre>	int for all requests to port s host */
if (bind(listenfd, (SA *)&se return -1;	<pre>rveraddr, sizeof(serveraddr)) &lt; 0)</pre>

# Echo Server: open\_listenfd (initialize socket address)

## Next, we initialize the socket with the server's Internet address (IP address and port)

struct sockaddr\_in serveraddr; /\* server's socket addr \*/
...
/\* listenfd will be an endpoint for all requests to port
 on any IP address for this host \*/
bzero((char \*) &serveraddr, sizeof(serveraddr));
serveraddr.sin\_family = AF\_INET;
serveraddr.sin\_addr.s\_addr = htonl(INADDR\_ANY);
serveraddr.sin\_port = htons((unsigned short)port);

#### IP addr and port stored in network (big-endian) byte order

- hton1() converts longs from host byte order to network byte order.
- htons() convers shorts from host byte order to network byte order.

F8 – 42 –

Systemprogrammering 2007

# Echo Server: open\_listenfd (listen)

## listen indicates that this socket will accept connection (connect) requests from clients.

int listenfd; /\* listening socket \*/
...
/\* Make it a listening socket ready to accept connection requests \*/
if (listen(listenfd, LISTENQ) < 0)
 return -1;
 return listenfd;</pre>

## We're finally ready to enter the main server loop that accepts and processes client connection requests.

Systemprogrammering 2007

F8 – 44 –

## Echo Server: Main Loop

The server loops endlessly, waiting for connection requests, then reading input from the client, and echoing the input back to the client.



F8 – 45 –

Systemprogrammering 2007

## Echo Server: accept Illustrated



1. Server blocks in accept, waiting for connection request on listening descriptor

2. Client makes connection request by calling and blocking in connect.



Systemprogrammering 2007

### Echo Server: accept

#### accept() blocks waiting for a connection request.

```
int listenfd; /* listening descriptor */
int connfd; /* connected descriptor */
struct sockaddr in clientaddr;
int clientlen;
```

clientlen = sizeof(clientaddr); connfd = Accept(listenfd, (SA \*)&clientaddr, &clientlen);

#### accept returns a connected descriptor (connfd) with the

#### same properties as the listening descriptor (listenfd)

but with a new (ephemeral) port.

- Returns when the connection between client and server is created and ready for I/O transfers.
- All I/O with the client will be done via the connected socket.

#### accept also fills in client's IP address.

F8 – 46 –

Systemprogrammering 2007

### **Connected vs. Listening Descriptors**

#### Listening descriptor

- End point for client connection requests.
- Created once and exists for lifetime of the server.

#### **Connected descriptor**

- End point of the connection between client and server.
- A new descriptor is created each time the server accepts a connection request from a client.
- Exists only as long as it takes to service client.

#### Why the distinction?

- Allows for concurrent servers that can communicate over many client connections simultaneously.
  - E.g., Each time we receive a new request, we fork a child to handle the request (more about this in later lectures).

## Echo Server: Identifying the Client

## The server can determine the domain name and IP address of the client.

F8 – 49 –

Systemprogrammering 2007

### Testing Servers Using telnet

## The telnet program is invaluable for testing servers that transmit ASCII strings over Internet connections

- Our simple echo server
- Web servers
- Mail servers

#### Usage:

- unix> telnet <host> <portnumber>
- Creates a connection with a server running on <host> and listening on port <portnumber>.

#### Systemprogrammering 2007

### Echo Server: echo

## The server uses RIO to read and echo text lines until EOF (end-of-file) is encountered.

- EOF notification caused by client calling close(clientfd).
- IMPORTANT: EOF is a condition, not a particular data byte.

void echo(int connfd)	
{	
<pre>size_t n;</pre>	
<pre>char buf[MAXLINE];</pre>	
rio_t rio;	
Rio_readinitb(&rio, connfd);	
<pre>while((n = Rio_readlineb(&amp;rio, buf, MAXLINE)) != 0) {</pre>	
<pre>printf("server received %d bytes\n", n);</pre>	
<pre>Rio_writen(connfd, buf, n);</pre>	
}	
}	

F8 – 50 –

Systemprogrammering 2007

### Testing the Echo Server With telnet

```
bass> echoserver 5000
server established connection with KITTYHAWK.CMCL (128.2.194.242)
server received 5 bytes: 123
server established connection with KITTYHAWK.CMCL (128.2.194.242)
server received 8 bytes: 456789
kittyhawk> telnet bass 5000
Trving 128.2.222.85...
Connected to BASS.CMCL.CS.CMU.EDU.
Escape character is '^]'.
123
123
Connection closed by foreign host.
kittyhawk> telnet bass 5000
Trying 128.2.222.85...
Connected to BASS.CMCL.CS.CMU.EDU.
Escape character is '^]'.
456789
456789
Connection closed by foreign host.
kittyhawk>
```

```
F8 – 52 –
```

Systemprogrammering 2007

### **Running the Echo Client and Server**

bass> echoserver 5000

server established connection with KITTYHAWK.CMCL (128.2.194.242)
server received 4 bytes: 123
server established connection with KITTYHAWK.CMCL (128.2.194.242)
server received 7 bytes: 456789
...

kittyhawk> echoclient bass 5000 Please enter msg: 123 Echo from server: 123

kittyhawk> echoclient bass 5000 Please enter msg: 456789 Echo from server: 456789 kittyhawk>

F8 – 53 –

Systemprogrammering 2007

### **For More Information**

- W. Richard Stevens, "Unix Network Programming: Networking APIs: Sockets and XTI", Volume 1, Second Edition, Prentice Hall, 1998.
  - THE network programming bible.

## Complete versions of the echo client and server are developed in the text.

- Available from /info/sysprog06/examples
- You should compile and run them for yourselves to see how they work.
- Feel free to borrow any of this code.

F8 – 54 –

Systemprogrammering 2007