

Network Programming

Kai Shen

1

IP Address vs. Machine Name

People: multiple identifiers

- SSN – good for machine bookkeeping, indexing
- name – human friendly, easy to remember

Internet hosts:

- IP address – used for addressing, routing on the Internet
- “name”, e.g., cycle1.csug.rochester.edu – human friendly

Q: map between IP addresses and names?

- **DNS query:** find the IP address for a given name
- not a fundamental part of the Internet as a network
- ease of use for Internet applications
 - “ssh cycle1.csug.rochester.edu” vs. “ssh 128.151.69.85”
 - <http://www.google.com> vs. <http://173.194.204.103>

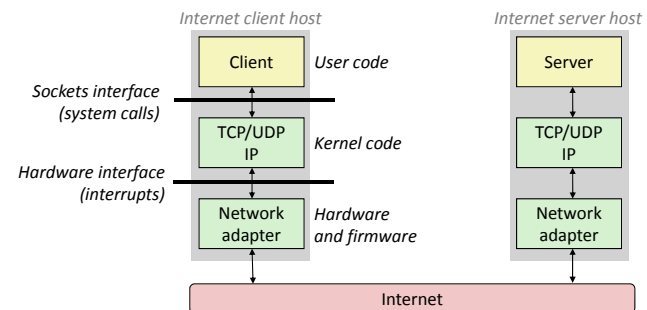
2

Domain Name System (DNS)

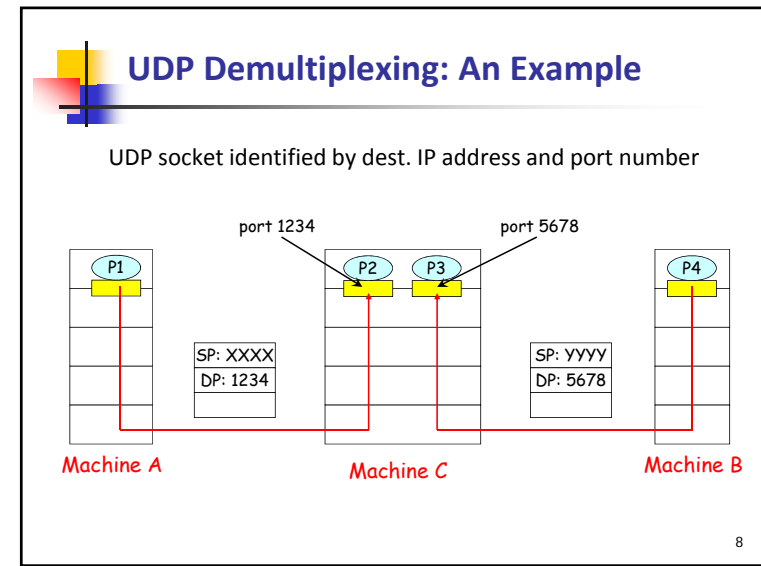
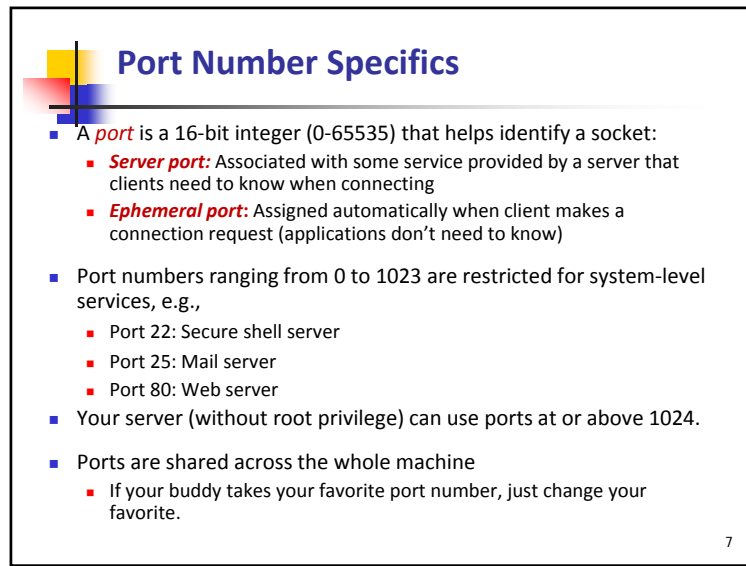
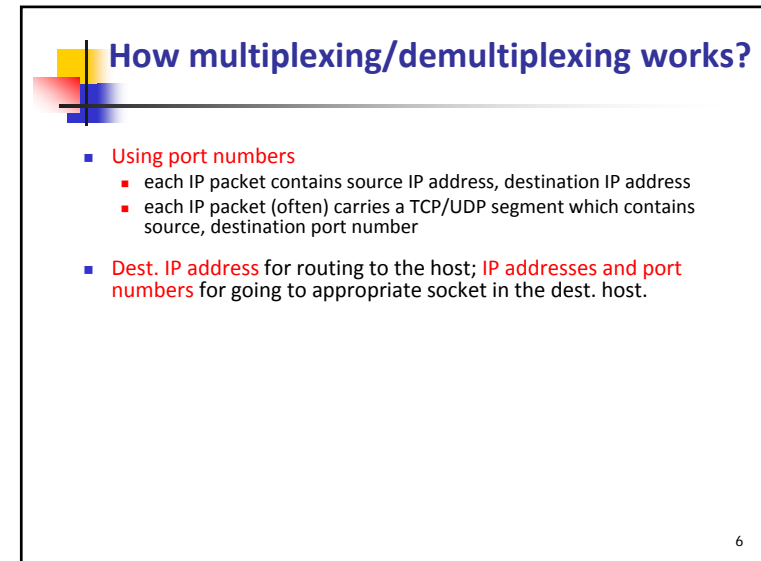
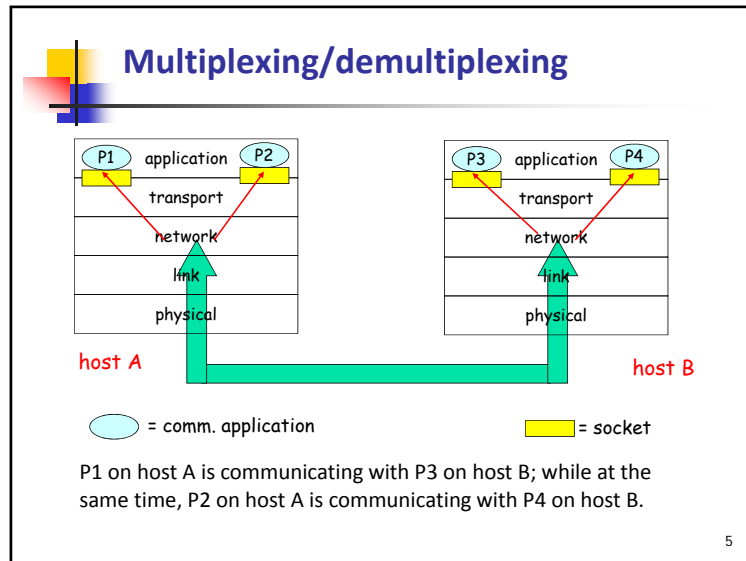
- Distributed database
 - implemented with collaboration of many name servers distributed all over the Internet in a hierarchical manner:
 - local name server manages the mappings for local hosts
 - a handful of global root name servers connect local name servers
 - extensive caching (local name server as well as local host)
- Try “nslookup www.google.com”
 - You see multiple answers. Why is that?

3

Hardware and Software Organization of an Internet Application



4




TCP Demultiplexing (Connection-oriented)

- Multiple TCP connections may simultaneously share a single port:
 - **Example:** web server (on default port 80) has multiple connections, one for each client
- TCP socket identified by 4-tuple:
 - (source IP address, source port number, dest IP address, dest port number)
 - recv host uses all four values to direct segment to appropriate socket

9

TCP Demultiplexing: An Example



10

TCP/UDP Sockets

Connection-oriented byte streams (TCP):

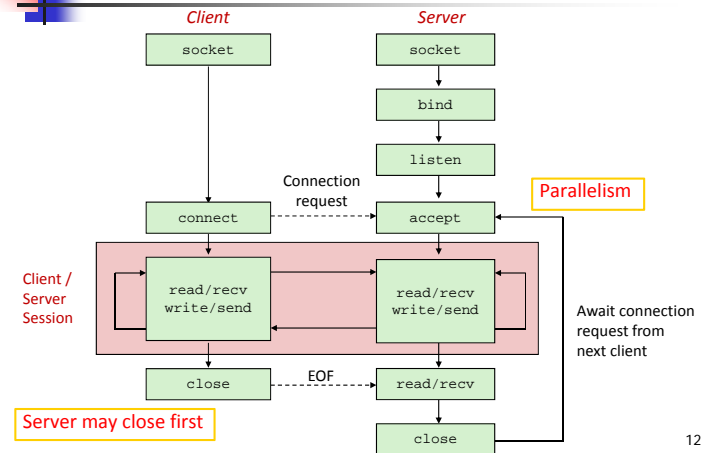
- Client/server need to establish connection before communications
- Communications flow in streams of ordered bytes
- But no clear sense of messages

Connectionless datagram (UDP):

- No connection is needed before communications;
- Clear boundaries between groups of bytes (messages)
- But no ordering between messages.

11

Overview of TCP Sockets



12

Socket Programming Notes

- **gethostbyname()**
 - Query name server for name to IP translation.
- Note
 - "gdb tells me **gethostbyname** fails with segmentation fault"!!!
 - Not reentrant. Why?
 - Only provide one IP address.

13

Socket Programming Notes

- **connfd = accept(listenfd, ...)**
 - Server waiting for a connection request.
- Notes
 - Returns when the connection with a new client is established and ready for data transfers
 - Data communications are done through the connection socket, not the listening socket
 - One listening socket for the whole server, one connection socket per active connection with a client

14

Socket Programming Notes

- **read/recv, write/send**
 - Data transfer calls.
- Note
 - TCP communications are stream-oriented, with no clear sense of message boundaries
 - when one sends 2000 bytes, the other side may receive twice (with 1460 bytes the first time, and 540 bytes the next time)
- How do you know when you are done with receiving?
 - Your protocol semantics should tell, three typical ways:
 - Fixed-size messages defined by protocol
 - Fixed-size header which contains a field that indicates the size of body
 - Always ends with a special pattern (e.g. an empty line "\n\n")

15

Socket Programming Notes

- How to write your receiver code?
- Fixed-size messages defined by protocol
 - **read/recv** for a certain size?
 - Must retry in case you don't get all the first time
 - Likely fine with one try if the size is small
- Fixed-size header which contains a field that indicates the size of body
 - Similar to above
- Always ends with a special pattern (e.g. an empty line "\n\n")
 - Be careful that the ending pattern may straddle over two messages

16

Socket Programming Notes

- Why I can't **bind** while restarting my server?
 - Your server port is kept by the OS for a certain amount of time in case some lingering packets belonging to the previous socket connection suddenly appears
 - Can be disabled by setting certain socket option

17

Socket Programming Notes

- Important!
 - Check the return code of every socket call (in fact, every system call)
 - Do error processing, warning messages for all cases
 - See UNIX manual pages for error codes and possible meaning

18

Testing Servers Using telnet

- The **telnet** program is invaluable for testing servers that transmit ASCII strings over Internet connections
 - Your simple FOOBAR server
 - Web servers
 - Mail servers
- Usage:
 - `unix> telnet <host> <portnumber>`
 - Creates a connection with a server running on `<host>` and listening on port `<portnumber>`

19

Testing Servers With telnet

```
linux> telnet www.google.com 80
Trying 173.194.73.99...
Connected to www.google.com.
Escape character is '^]'.
GET / HTTP/1.0

HTTP/1.0 200 OK
Date: Thu, 12 Apr 2012 17:28:59 GMT
... ..
```

```
linux> telnet <gmail mail server> 25
220 gmail.com
HELO cs.rochester.edu
250 Hello cs.rochester.edu, pleased to meet you
MAIL FROM: <bob@cs.rochester.edu>
250 bob@cs.rochester.edu... Sender ok
RCPT TO: <alice@gmail.com>
250 alice@gmail.com... Recipient ok
... ..
```

20



Learn Network Programming

- Learn to use UNIX manual pages
- Complete versions of an echo client and server in the text
 - Updated versions linked to course website
- Find information on the web
- Reference book
 - Network programming in C: Stevens, "Unix Network Programming, Volume I", 2nd edition, 1998.

21



Disclaimer

Parts of the lecture slides contain original work of James Kurose, Larry Peterson, and Keith Ross. The slides are intended for the sole purpose of instruction of computer organization course at the University of Rochester. All copyrighted materials belong to their original owner(s).

22