

CS5503

Lecture # 16

2019-05-22

Administrica

OSSF

IPv4 - 32-bit addresses

IPv6 - 128-bit addresses

DNS system:

- distributed hierarchical database
- maps domain names to IP addresses

network OS

- networking applications

distributed OS

- resource allocation across entire database
- distributed lock service

Sockets: "high-level" abstraction

- once network connection is established, can use read/write system calls
- socket descriptor: specialized file descriptor

Enums (byte order)

#include <arpa/inet.h>

htonl } struct (16 bits)
htons }
htohl } long (32 bits)
htohs }

IP Address Manipulation

- text format

32-bit number

- host byte order (little endian on Intel)
- network byte order (big endian)

#include

- `sys/socket.h`

- `netinet/in.h`

- `arpa/inet.h`

typedef uint32_t in_addr_t

struct in_addr {

in_addr_t s_addr;

}

`inet_aton(char *cp, struct in_addr *in_p)`

string to 32-bit int
`strncpy(12.34.56.78) → 32-bit int`

`inet_ntoa`

number to dotted string

Network Programming API

- better-designed API (hypothetical):

namespace-oriented

e.g.

{
 Open ("/net/127.45.67.89/80")
 or /net/<ip>/tcp/<port>/
 or like URI (Uniform Resource Identifier)
 http://host:port/path(...)

- still need to translate that to wire format (protocol)
- text-based (URI): easier on eyes, harder on computers
 - good tradeoff these days
(programmers are expensive)
- but we're all about looking at the
 natty gatty here

- Given host name, get IP address:

library call

- looks up system config files

- /etc/hosts

/etc/host.conf

/etc/resolv.conf

- obsolete: struct hostent # gethostbyname (char * name)
pointer to 4 bytes (IP address)

getaddrinfo

- more complex API

- returns struct addrinfo #

- Once name lookup is complete, host name is not used further

- all connections via IP address

- higher-level protocol: http

http://bulk.com/ } => 12.34.56.78

http://thr.com/

=> "virtual" host field in http

- IP address manipulation

IP v4
address
object
(32-bits) {
- typedef uint32_t in_addr_t
- struct in_addr {
 in_addr_t s_addr,
};

inet_aton (char *cp, struct in_addr *ip)

- parse IP address string (dot notation)

inet_ntoa

- convert address to string rep

Network Programming API

o.g. X.25 - } ^{single} legacy
apps
not just TCP/IP
wrapper

- Socket: (unix) abstraction for all networks
 - also used by windows
 - AKA Berkeley Sockets
 - mostly TCP/UDP
 - unix-domain sockets (IPC)
 - X.25 (obsolete) packet switching
- create socket: `socket(2)` system call
 - returns socket descriptor
 - small integer, same namespace as file descriptors
 - once the socket is set up: may use `read/write/close` system calls
 - may also use `send/receive`
 - better control over message-oriented protocols
 - e.g. UDP
 - additional system calls to set up connections

⇒ acts like a file, but not completely transparent
↳ different "constructor"

Unix Socket

include

(`sys/types.h`)

(`sys/socket.h`)

Socket (int domain, int type, int protocol)

domain

AF_UNIX

(yet another IPC)

AF_INET

(internet protocol - IP)

type

SOCK_RAW

(raw IP packets)

SOCK_DGRAM

(UDP)

SOCK_STREAM

(TCP)

protocol

0 for AF_INET

* freshly-allocated socket object

→ not (yet) tied to connection

(no associated port)

Socket API (cont.)

- 2 operating modes

passive (server)

active (client)

⇒ difference between client & server:
who initiates the transaction

Server:

tcp

bind() - associate socket with port #

listen() - specify operating mode

accept()

- returns new socket descriptor
for each connection

udp

bind()

recvfrom()

↳ fd, buf, size,

flags, sockaddr, address

Socket API (cont.)

client

tcp (optional: bind) → associate local port #
connect() otherwise, port is assigned by system

udp

bind() (optional, receive reply)

sendto()

or { connect()
send() }

bind()

- associates socket & address (host:port)

int bind (int sd,

const struct sockaddr * addr,

socklen_t addrlen)

- sockaddr: essentially a void type

- protocol-specific address structs are cast
to sockaddr

local
address

Socket API (cont.)

#include <sys/types.h> } the usual
#include <sys/socket.h> } suspects

```
int accept ( int sock fd,  
            struct sockaddr * addr,  
            socklen_t * addr len )
```

← pointer
→ out parameters:

- filled in with remote address
- addr may be null
- don't care

- returns new sd for communication
original sock fd remains open
for additional connections

- listen(2) - ^{set} size of backlog

- may have multiple concurrent
connections (e.g. multiple web page
requests all go to port 80)

Addressing

text: host:port

text processing

- expensive
- but only have to do it once

address data structure

struct sockaddr

But: addressing is dependent on the protocol

- internet address: struct sockaddr_in

→ overlay: 1st field is family (AF_INET)

sin_family AF_INET

sin_port

sin_addr

} struct sin_addr
uint32 s_addr;
}

⇒ must construct IP Address object

& cast to sockaddr

Socket API (cont.)

internet address (TCP/IP or UDP/IP):

use struct sockaddr_in

sin_family = AF_INET (tag field)

sin_addr.s_addr = IP address 32 bits

sin_port = port 16 bits

host, port are in network byte order

⇒ use htons }
htonl }

- for server, use s_addr =
htonl(INADDR_ANY)

use
any network
interface on
machine
→ Computer may
have multiple
NICs

connect (int sd,
const struct sockaddr addr,
socklen_t addrlen)

Remote address

use port = 0 for randomly-assigned

unused port

- ephemeral ports

UDP - User Datagram Protocol

client:

Socket()

bind() (if response is needed)

sendto()
recvfrom()

or { connect()
send
recv

server:

Socket()

bind()

recvfrom()

sendto()

- connectionless
- packets may get lost, arrive out of order
- multicast

Applications

- TFTP/BOOTP
- NFS
- DNS
- real-time streaming

} either you don't care if some data is lost or you have your own handle error recovery scheme

UDP (cont.)

Client (connectionless)

- create socket
- sendto (sock
msg
len
flags
addr
addr len)

RTFM

⇒ line coding

- sockets & concurrency

- threads

- socket fd stays open across fork/exec

- unless you explicitly set close-on-exec

- like any file descriptor