

15-441 Computer Networking

Lecture 7 – IP Addressing and Forwarding

Outline

- Methods for packet forwarding
- Traditional IP addressing
- CIDR IP addressing
- Forwarding example

Lecture 7: 9 - 18-01 2

Techniques for Forwarding Packets

- Source routing
 - Packet carries path
- Table of virtual circuits
 - Connection routed through network to setup state
 - Packets forwarded using connection state
- Table of global addresses (IP)
 - Routers keep next hop for destination
 - Packets carry destination address

Lecture 7: 9 - 18-01 3

Source Routing

- List entire path in packet
 - Driving directions (north 3 hops, east, etc..)
- Router processing
 - Examine first step in directions
 - Strip first step from packet
 - Forward to step just stripped off

Lecture 7: 9 - 18-01 4

Source Routing Example

The diagram shows a network topology with a Sender, three routers (R1, R2, R3), and a Receiver. The path taken by the packet is Sender → R1 → R2 → R3 → Receiver. The packet header at each router shows the remaining path: at R1, the header is 3,4,3; at R2, it is 4,3; and at R3, it is 3.

Lecture 7: 9 - 18-01 5

Source Routing

- Advantages
 - Switches can be very simple and fast
- Disadvantages
 - Variable (unbounded) header size
 - Sources must know or discover topology (e.g., failures)
- Typical use
 - Ad-hoc networks (DSR)
 - Machine room networks (Myrinet)

Lecture 7: 9 - 18-01 6

Global Addresses (IP)




- Each packet has destination address
- Each switch has forwarding table of destination → next hop
 - At v and x: destination → east
 - At w and y: destination → south
 - At z: destination → north
- Distributed routing algorithm for calculating forwarding tables

Lecture 7: 9-18-01

7

Global Address Example



Lecture 7: 9-18-01

8

Router Table Size



- One entry for every host on the Internet
 - 100M entries, doubling every year
- One entry for every LAN
 - Every host on LAN shares prefix
 - Still too many, doubling every year
- One entry for every organization
 - Every host in organization shares prefix
 - Requires careful address allocation

Lecture 7: 9-18-01

9

Global Addresses



- Advantages
 - Stateless – simple error recovery
- Disadvantages
 - Every switch knows about every destination
 - Potentially large tables
 - All packets to destination take same route

Lecture 7: 9-18-01

10

How do we set up Routing Tables?



- Graph theory to compute “shortest path”
 - Switches = nodes
 - Links = edges
 - Delay, hops = cost
- Need to adapt to changes in topology

Lecture 7: 9-18-01

11

Virtual Circuits/Tag Switching



- Connection setup phase
 - Use other means to route setup request
 - Each router allocates flow ID on local link
 - Creates mapping of inbound flow ID/port to outbound flow ID/port
- Each packet carries connection ID
 - Sent from source with 1st hop connection ID
- Router processing
 - Lookup flow ID – simple table lookup
 - Replace flow ID with outgoing flow ID
 - Forward to output port

Lecture 7: 9-18-01

12

Virtual Circuits Examples

Lecture 7: 9-18-01 13

Virtual Circuits

- Advantages
 - More efficient lookup (simple table lookup)
 - More flexible (different path for each flow)
 - Can reserve bandwidth at connection setup
 - Easier for hardware implementations
- Disadvantages
 - Still need to route connection setup request
 - More complex failure recovery – must recreate connection state
- Typical uses
 - ATM – combined with fix sized cells
 - MPLS – tag switching for IP networks

Lecture 7: 9-18-01 14

IP Datagrams on Virtual Circuits

- Challenge – when to setup connections
 - At bootup time – permanent virtual circuits (PVC)
 - Large number of circuits
 - For every packet transmission
 - Connection setup is expensive
 - For every connection
 - What is a connection?
 - How to route connectionless traffic?

Lecture 7: 9-18-01 15

IP Datagrams on Virtual Circuits

- Traffic pattern
 - Few long lived flows
 - Flow – set of data packets from source to destination
 - Large percentage of packet traffic
 - Improving forwarding performance by using virtual circuits for these flows
- Other traffic uses normal IP forwarding

Lecture 7: 9-18-01 16

Comparison

	Source Routing	Global Addresses	Virtual Circuits
Header Size	Worst	OK – Large address	Best
Router Table Size	None	Number of hosts (prefixes)	Number of circuits
Forward Overhead	Best	Prefix matching	Pretty Good
Setup Overhead	None	None	Connection Setup
Error Recovery	Tell all hosts	Tell all routers	Tell all routers and Tear down circuit and re-route

Lecture 7: 9-18-01 17

Forwarding vs. Routing

- **Forwarding:** the process of moving packets from input to output
 - The forwarding table
 - Information in the packet
- **Routing:** process by which the forwarding table is built and maintained
 - One or more routing protocols
 - Procedures (algorithms) to convert routing info to forwarding table.

Lecture 7: 9-18-01 18

Outline

- Methods for packet forwarding
- **Traditional IP addressing**
- CIDR IP addressing
- Forwarding example

Lecture 7: 9-18-01 19

How is IP Design Standardized?

- IETF
 - Voluntary organization
 - Meeting every 4 months
 - Working groups and email discussions
- "We reject kings, presidents, and voting; we believe in rough consensus and running code" (Dave Clark 1992)
 - Need 2 independent, interoperable implementations for standard
- IRTF
 - End2End
 - Reliable Multicast, etc..

Lecture 7: 9-18-01 20

Addressing in IP

- IP addresses are names of interfaces
- Domain Name System (DNS) names are names of hosts
- DNS binds host names to interfaces
- Routing binds interface names to paths

Lecture 7: 9-18-01 21

Addressing Considerations

- Fixed length or variable length?
- Issues:
 - Flexibility
 - Processing costs
 - Header size
- Engineering choice: IP uses fixed length addresses

Lecture 7: 9-18-01 22

Addressing Considerations

- Structured vs flat
- Issues
 - What information would routers need to route to Ethernet addresses?
 - Need structure for designing scalable binding from interface name to route!
 - How many levels? Fixed? Variable?

Lecture 7: 9-18-01 23

IP Addresses

- Fixed length: 32 bits
- Initial classful structure (1981)
- Total IP address size: 4 billion
 - Class A: 128 networks, 16M hosts
 - Class B: 16K networks, 64K hosts
 - Class C: 2M networks, 256 hosts

High Order Bits	Format	Class
0	7 bits of net, 24 bits of host	A
10	14 bits of net, 16 bits of host	B
110	21 bits of net, 8 bits of host	C

Lecture 7: 9-18-01 24

IP Address Classes (Some are Obsolete)

Class	Network ID	Host ID
Class A	0	31
Class B	10	30
Class C	110	29
Class D	1110	28
Class E	1111	27

Lecture 7: 9-18-01 25

Some Special IP Addresses

- 127.0.0.1: local host (a.k.a. the loopback address)
- Host bits all set to 0: network address
- Host bits all set to 1: broadcast address

Lecture 7: 9-18-01 26

Interaction with Link Layer

- How does one find the Ethernet address of a IP host?
- ARP
 - Broadcast search for IP address
 - E.g., "who-has 128.2.184.45 tell 128.2.206.138" sent to Ethernet broadcast (all FF address)
 - Destination responds (only to requester using unicast) with appropriate 48-bit Ethernet address
 - E.g., "reply 128.2.184.45 is-at 0:d0:bc:f2:18:58" sent to 0:c0:4f:d:ed:c6

Lecture 7: 9-18-01 27

Original IP Route Lookup

- Address classes
 - A: 0 | 7 bit network | 24 bit host (16M each)
 - B: 10 | 14 bit network | 16 bit host (64K)
 - C: 110 | 21 bit network | 8 bit host (255)
- Address would specify prefix for forwarding table
 - Simple lookup

Lecture 7: 9-18-01 28

Original IP Route Lookup – Example

- www.cmu.edu address 128.2.11.43
 - Class B address – class + network is 128.2
 - Lookup 128.2 in forwarding table
 - Prefix – part of address that really matters for routing
- Forwarding table contains
 - List of class+network entries
 - A few fixed prefix lengths (8/16/24)
- Large tables
 - 2 Million class C networks

Lecture 7: 9-18-01 29


Subnet Addressing RFC917 (1984)

- For class B & C networks
- Very few LANs have close to 64K hosts
 - For electrical/LAN limitations, performance or administrative reasons
- Need simple way to get multiple "networks"
 - Use bridging, multiple IP networks or split up single network address ranges (subnet)
 - Must reduce the total number of network addresses that are assigned
- CMU case study in RFC
 - Chose not to adopt – concern that it would not be widely supported ☹

Lecture 7: 9-18-01 30

Subnetting

- Variable length subnet masks
 - Could subnet a class B into several chunks



Lecture 7: 9 - 18-01

31

Subnetting Example


- Assume an organization was assigned address 150.100
- Assume < 100 hosts per subnet
- How many host bits do we need?
 - Seven
- What is the network mask?
 - 11111111 11111111 11111111 10000000
 - 255.255.255.128

Lecture 7: 9 - 18-01

32

Forwarding Example

- Assume a packet arrives with address 150.100.12.176
- Step 1: AND address with class + subnet mask



Lecture 7: 9 - 18-01

33

Outline

- Methods for packet forwarding
- Traditional IP addressing
- **CIDR IP addressing**
- Forwarding example

Lecture 7: 9 - 18-01

34


IP Address Problem (1991)

- Address space depletion
 - In danger of running out of classes A and B
- Why?
 - Class C too small for most domains
 - Very few class A – very careful about giving them out
 - Class B – greatest problem

Lecture 7: 9 - 18-01

35

IP Address Utilization ('98)



Lecture 7: 9 - 18-01

36

IP Addressing Problems

- Class B sparsely populated
 - But people refuse to give it back
- Large forwarding tables
 - 2 Million possible class C groups
- Solution
 - Assign multiple class C addresses
 - Assign consecutive blocks
 - RFC1338 – Classless Inter-Domain Routing (CIDR)

Lecture 7: 9 - 18-01 37

Classless Inter-Domain Routing

- Do not use classes to determine network ID
- Assign any range of addresses to network
 - Use common part of address as network number
 - E.g., addresses 192.4.16 - 192.4.31 have the first 20 bits in common. Thus, we use these 20 bits as the network number
 - netmask is /20, /xx is valid for almost any xx
- Enables more efficient usage of address space (and router tables)

Lecture 7: 9 - 18-01 38

CIDR

- Supernetting
 - Assign adjacent net addresses to same org
 - Classless routing (CIDR)
- How does this help routing table?
 - Combine forwarding table entries whenever all nodes with same prefix share same hop

Lecture 7: 9 - 18-01 39

IP Addresses: How to Get One?

Network (network portion):

- Get allocated portion of ISP's address space:

ISP's block	11001000 00010111 00010000	00000000	200.23.16.0/20
Organization 0	11001000 00010111 00010000	00000000	200.23.16.0/23
Organization 1	11001000 00010111 00010010	00000000	200.23.18.0/23
Organization 2	11001000 00010111 00010100	00000000	200.23.20.0/23
...
Organization 7	11001000 00010111 00011110	00000000	200.23.30.0/23

Lecture 7: 9 - 18-01 40

IP addresses: How to Get One?

Q: How does an ISP get block of addresses?

A: ICANN: Internet Corporation for Assigned Names and Numbers

- Allocates addresses
- Manages DNS
- Assigns domain names, resolves disputes

Lecture 7: 9 - 18-01 41

IP addresses: How to Get One?


Hosts (host portion):

- Hard-coded by system admin in a file
- DHCP: Dynamic Host Configuration Protocol: dynamically get address: "plug-and-play"
 - Host broadcasts "DHCP discover" msg
 - DHCP server responds with "DHCP offer" msg
 - Host requests IP address: "DHCP request" msg
 - DHCP server sends address: "DHCP ack" msg

Lecture 7: 9 - 18-01 42

Hierarchical Addressing: Route Aggregation

- Hierarchical addressing allows efficient advertisement of routing information:




Lecture 7: 9-18-01

43

Hierarchical Addressing: More Specific Routes

- ISPs-R-Us has a more specific route to Organization 1



Lecture 7: 9-18-01

44


CIDR Example

- Network provide is allocated 8 class C chunks, 201.10.0.0 to 201.10.7.255
 - Allocation uses 3 bits of class C space
 - Remaining 21 bits are network number, written as 201.10.0.0/21
- Replaces 8 class C routing entries with 1 combined entry
 - Routing protocols carry prefix with destination network address
 - Longest prefix match for forwarding

Lecture 7: 9-18-01

45

CIDR Illustration




Lecture 7: 9-18-01

46

CIDR Addressing Shortcomings

- Multi-homing
- Customer selecting a new provider



Lecture 7: 9-18-01

47

Outline

- Methods for packet forwarding
- Traditional IP addressing
- CIDR IP addressing
- Forwarding example

Lecture 7: 9-18-01

48

Routing to the Network

- Packet to 10.1.1.3 arrives
- Path is R2 – R1 – H1 – H2

Lecture 7: 9-18-01 49

Routing Within the Subnet

- Packet to 10.1.1.3
- Matches 10.1.0.0/23

Routing table at R2

Destination	Next Hop	Interface
127.0.0.1	127.0.0.1	lo0
Default or 0/0	provider	10.1.16.1
10.1.8.0/24	10.1.8.1	10.1.8.1
10.1.2.0/23	10.1.2.1	10.1.2.1
10.1.0.0/23	10.1.2.2	10.1.2.1

Lecture 7: 9-18-01 50

Routing Within the Subnet

- Packet to 10.1.1.3
- Matches 10.1.1.1/31
- Longest prefix match

Routing table at R1

Destination	Next Hop	Interface
127.0.0.1	127.0.0.1	lo0
Default or 0/0	10.1.2.1	10.1.2.2
10.1.0.0/24	10.1.0.1	10.1.0.1
10.1.1.0/24	10.1.1.1	10.1.1.4
10.1.2.0/23	10.1.2.2	10.1.2.2
10.1.1.2/31	10.1.1.2	10.1.1.2

Lecture 7: 9-18-01 51

Routing Within the Subnet

- Packet to 10.1.1.3
- Direct route
- Longest prefix match

Routing table at H1

Destination	Next Hop	Interface
127.0.0.1	127.0.0.1	lo0
Default or 0/0	10.1.1.1	10.1.1.2
10.1.1.0/24	10.1.1.2	10.1.1.1
10.1.1.3/31	10.1.1.2	10.1.1.2

Lecture 7: 9-18-01 52