



# IP Addressing and Forwarding

*Acknowledgments: Lecture slides are from Computer networks course thought by Jennifer Rexford at Princeton University. When slides are obtained from other sources, a reference will be noted on the bottom of that slide. A full list of references is provided on the last slide.*

# Goals of Today's Lecture



- IP addresses
  - Dotted-quad notation
  - IP prefixes for aggregation
- Address allocation
  - Classful addresses
  - Classless InterDomain Routing (CIDR)
  - Growth in the number of prefixes over time
- Packet forwarding
  - Forwarding tables
  - Longest-prefix match forwarding
  - Where forwarding tables come from



# IP Address (IPv4)

- A unique 32-bit number
- Identifies an interface (on a host, on a router, ...)
- Represented in dotted-quad notation

**12**



**34**



**158**



**5**



00001100

00100010


10011110

00000101



# Grouping Related Hosts

- The Internet is an “inter-network”
  - Used to connect *networks* together, not *hosts*
  - Needs a way to address a network (i.e., group of hosts)

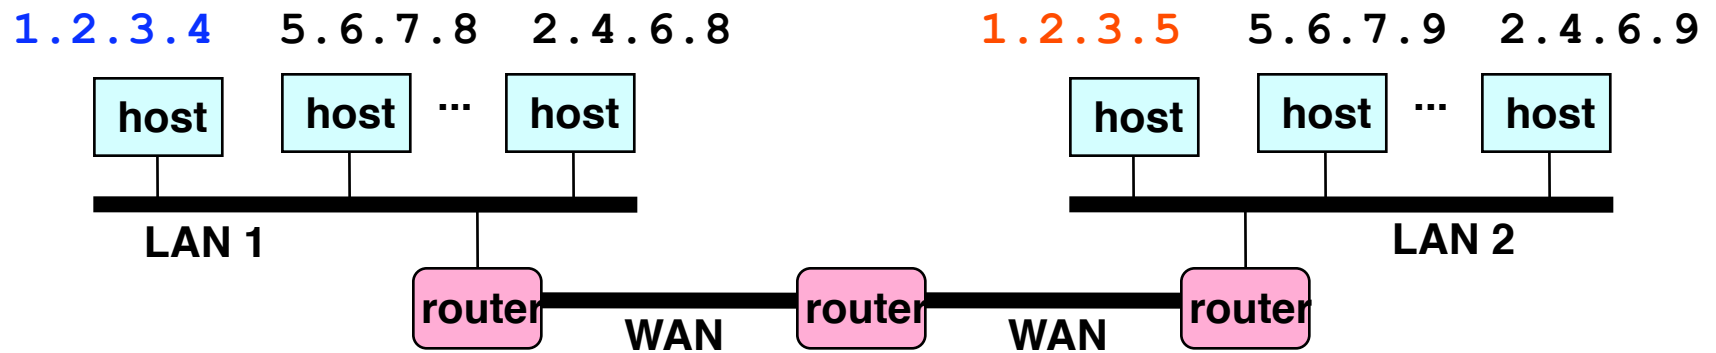


**LAN = Local Area Network**  
**WAN = Wide Area Network**



# Scalability Challenge

- Suppose hosts had arbitrary addresses
  - Then every router would need a lot of information
  - ...to know how to direct packets toward *every* host



1.2.3.4	←
1.2.3.5	→
⋮	

forwarding table



# Standard CS Trick

Have a scalability problem?


Introduce hierarchy...

# Hierarchical Addressing in U.S. Mail



- Addressing in the U.S. mail

- Zip code: 08540
- Street: Olden Street
- Building on street: 35
- Room in building: 306
- Name of occupant: Jennifer Rexford



- Forwarding the U.S. mail


- Deliver letter to the post office in the zip code
- Assign letter to mailman covering the street
- Drop letter into mailbox for the building/room
- Give letter to the appropriate person



# Hierarchical Addressing: IP Prefixes



- Divided into network & host portions (left and right)
- 12.34.158.0/24 is a 24-bit **prefix** with  $2^8$  addresses







# IP Address and a 24-bit Subnet Mask

Address

12

34


158

5



00001100	00100010	10011110	00000101
----------	----------	----------	----------

11111111	11111111	11111111	00000000
----------	----------	----------	----------



255

255

255

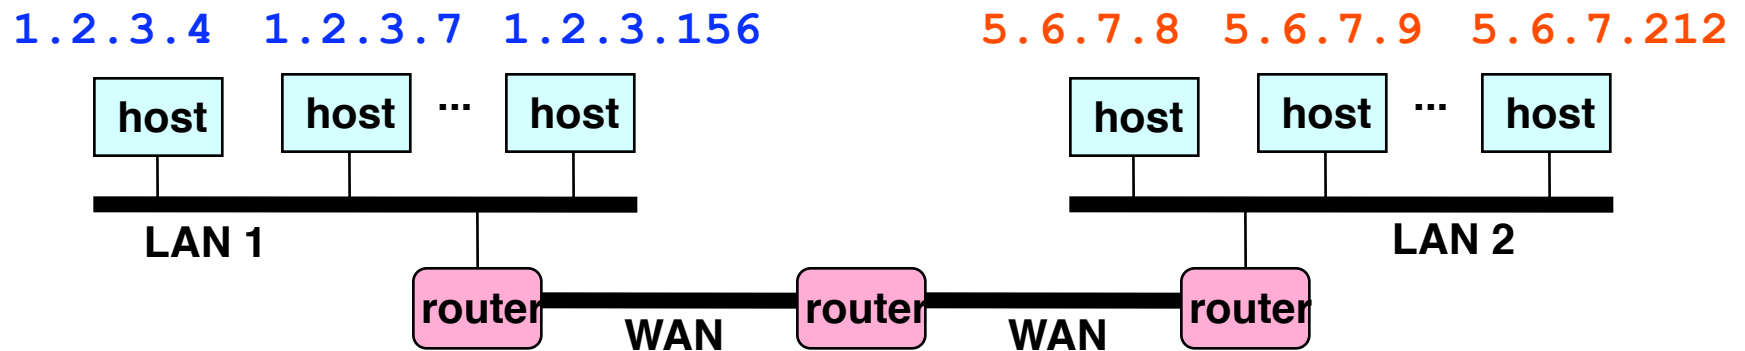
0

Mask



# Scalability Improved

- Number related hosts from a common subnet
  - 1.2.3.0/24 on the left LAN
  - 5.6.7.0/24 on the right LAN



1.2.3.0/24	←
5.6.7.0/24	→

forwarding table





# Address Allocation

# Classful Addressing



- In the olden days, only fixed allocation sizes
  - Class A: 0\*
    - Very large /8 blocks (e.g., MIT has 18.0.0.0/8)
  - Class B: 10\*
    - Large /16 blocks (e.g., Princeton has 128.112.0.0/16)
  - Class C: 110\*
    - Small /24 blocks (e.g., AT&T Labs has 192.20.225.0/24)
  - Class D: 1110\*
    - Multicast groups
  - Class E: 11110\*
    - Reserved for future use
- This is why folks use dotted-quad notation!

# Classless Inter-Domain Routing (CIDR)



Use two 32-bit numbers to represent a network.  
Network number = IP address + Mask


IP Address : 12.4.0.0      IP Mask: 255.254.0.0

Address

00001100	00000100	00000000	00000000
----------	----------	----------	----------

Mask

11111111	11111110	00000000	00000000
----------	----------	----------	----------




Written as 12.4.0.0/15



# CIDR: Hierarchical Address Allocation


- Prefixes are key to Internet scalability
  - Address allocated in contiguous chunks (prefixes)
  - Routing protocols and packet forwarding based on prefixes
  - Today, routing tables contain ~200,000 prefixes



# Scalability: Address Aggregation



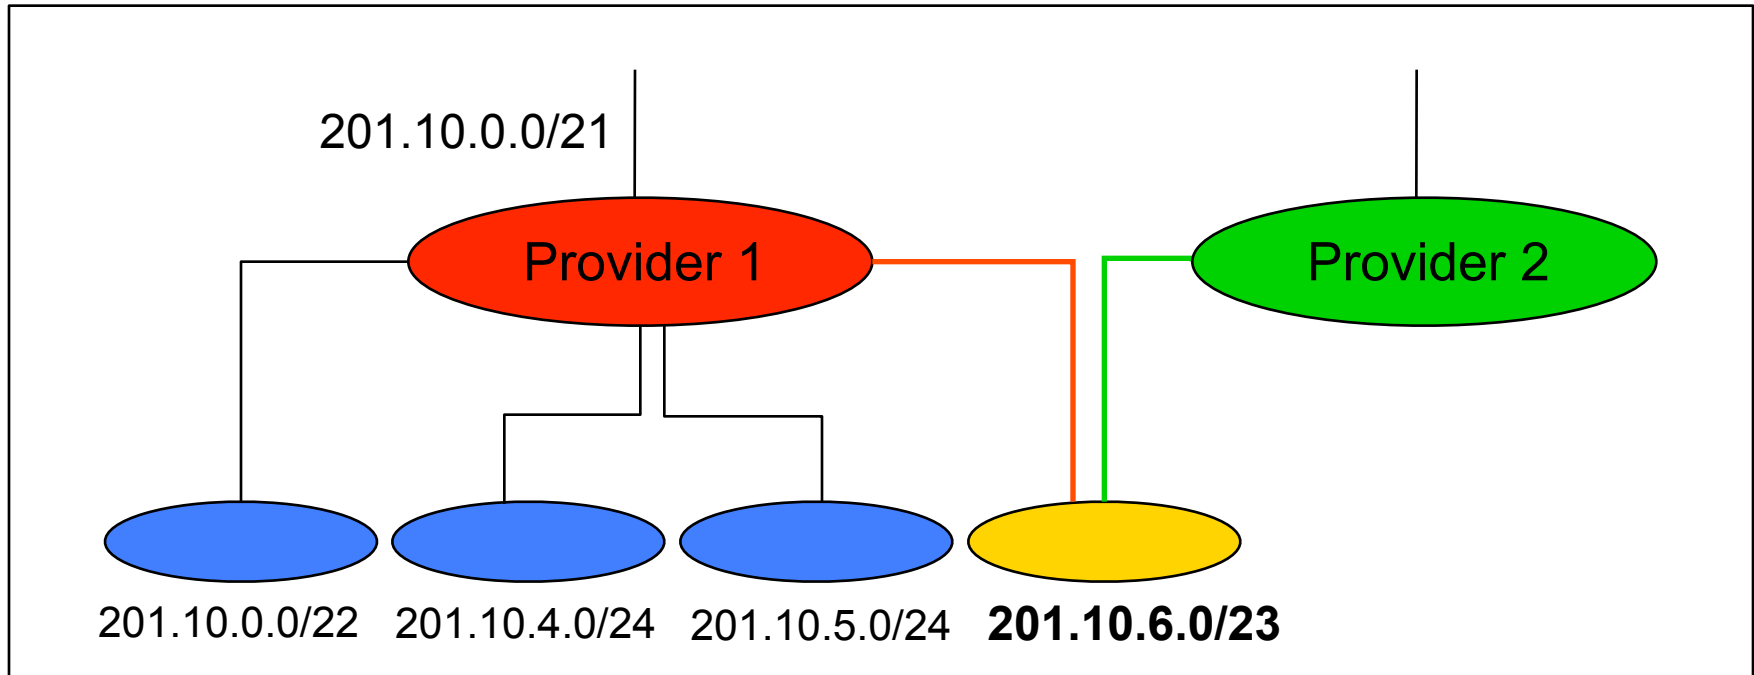
Provider is given 201.10.0.0/21



Routers in the rest of the Internet just need to know how to reach **201.10.0.0/21**. The provider can direct the IP packets to the appropriate **customer**.



# But, Aggregation Not Always Possible




***Multi-homed* customer with 201.10.6.0/23 has two providers. Other parts of the Internet need to know how to reach these destinations through *both* providers.**

# Scalability Through Hierarchy




- Hierarchical addressing
  - Critical for scalable system
  - Don't require everyone to know everyone else
  - Reduces amount of updating when something changes
- Non-uniform hierarchy
  - Useful for heterogeneous networks of different sizes
  - Initial class-based addressing was far too coarse
  - Classless InterDomain Routing (CIDR) helps
- Next few slides
  - History of the number of globally-visible prefixes
  - Plots of # of prefixes vs. time

# Pre-CIDR (1988-1994): Steep Growth




Growth faster than improvements in equipment capability

# CIDR Deployed (1994-1996): Much Flatter




Efforts to aggregate (even decreases after IETF meetings!)

# CIDR Growth (1996-1998): Roughly Linear




Good use of aggregation, and peer pressure in CIDR report

# Boom Period (1998-2001): Steep Growth



Internet boom and increased multi-homing

# Long-Term View (1989-2005): Post-Boom



# Obtaining a Block of Addresses



- Separation of control
  - Prefix: assigned *to* an institution
  - Addresses: assigned *by* the institution to their nodes
- Who assigns prefixes?
  - Internet Corporation for Assigned Names and Numbers
    - Allocates large address blocks to Regional Internet Registries
  - Regional Internet Registries (RIRs)
    - E.g., ARIN (American Registry for Internet Numbers)
    - Allocates address blocks within their regions
    - Allocated to Internet Service Providers and large institutions
  - Internet Service Providers (ISPs)
    - Allocate address blocks to their customers
    - Who may, in turn, allocate to their customers...



# Figuring Out Who Owns an Address



- Address registries
  - Public record of address allocations
  - Internet Service Providers (ISPs) should update when giving addresses to customers
  - However, records are notoriously out-of-date
- Ways to query
  - UNIX: “whois -h whois.arin.net 128.112.136.35”
  - <http://www.arin.net/whois/>
  - <http://www.geektools.com/whois.php>
  - ...



# Example Output for 213.233.168.1

inetnum: 213.233.168.0 - 213.233.175.255  
netname: SCHOOLNET-TEH-IR  
descr: Sharif University Of Technology  
country: IR  
person: Yahya Tabesh  
address: Computer Center, Sharif University of Technology  
address: Azadi Ave., Tehran, Iran.  
phone: +98 21 6005319  
fax-no: +98 21 6019568  
e-mail: [tabesh@sharif.ac.ir](mailto:tabesh@sharif.ac.ir)  
mnt-by: SHARIF-EDU-MNT

# Are 32-bit Addresses Enough?



- Not all that many unique addresses
  - $2^{32} = 4,294,967,296$  (just over four billion)
  - Plus, some are reserved for special purposes
  - And, addresses are allocated in larger blocks
- And, many devices need IP addresses
  - Computers, PDAs, routers, tanks, toasters, ...
- Long-term solution: a larger address space
  - IPv6 has 128-bit addresses ( $2^{128} = 3.403 \times 10^{38}$ )
- Short-term solutions: limping along with IPv4
  - Private addresses
  - Network address translation (NAT)
  - Dynamically-assigned addresses (DHCP)

# Hard Policy Questions



- How much address space per geographic region?
  - Equal amount per country?
  - Proportional to the population?
  - What about addresses already allocated?
- Address space portability?
  - Keep your address block when you change providers?
  - Pro: avoid having to renumber your equipment
  - Con: reduces the effectiveness of address aggregation
- Keeping the address registries up to date?
  - What about mergers and acquisitions?
  - Delegation of address blocks to customers?
  - As a result, the registries are horribly out of date




# Packet Forwarding

# Hop-by-Hop Packet Forwarding



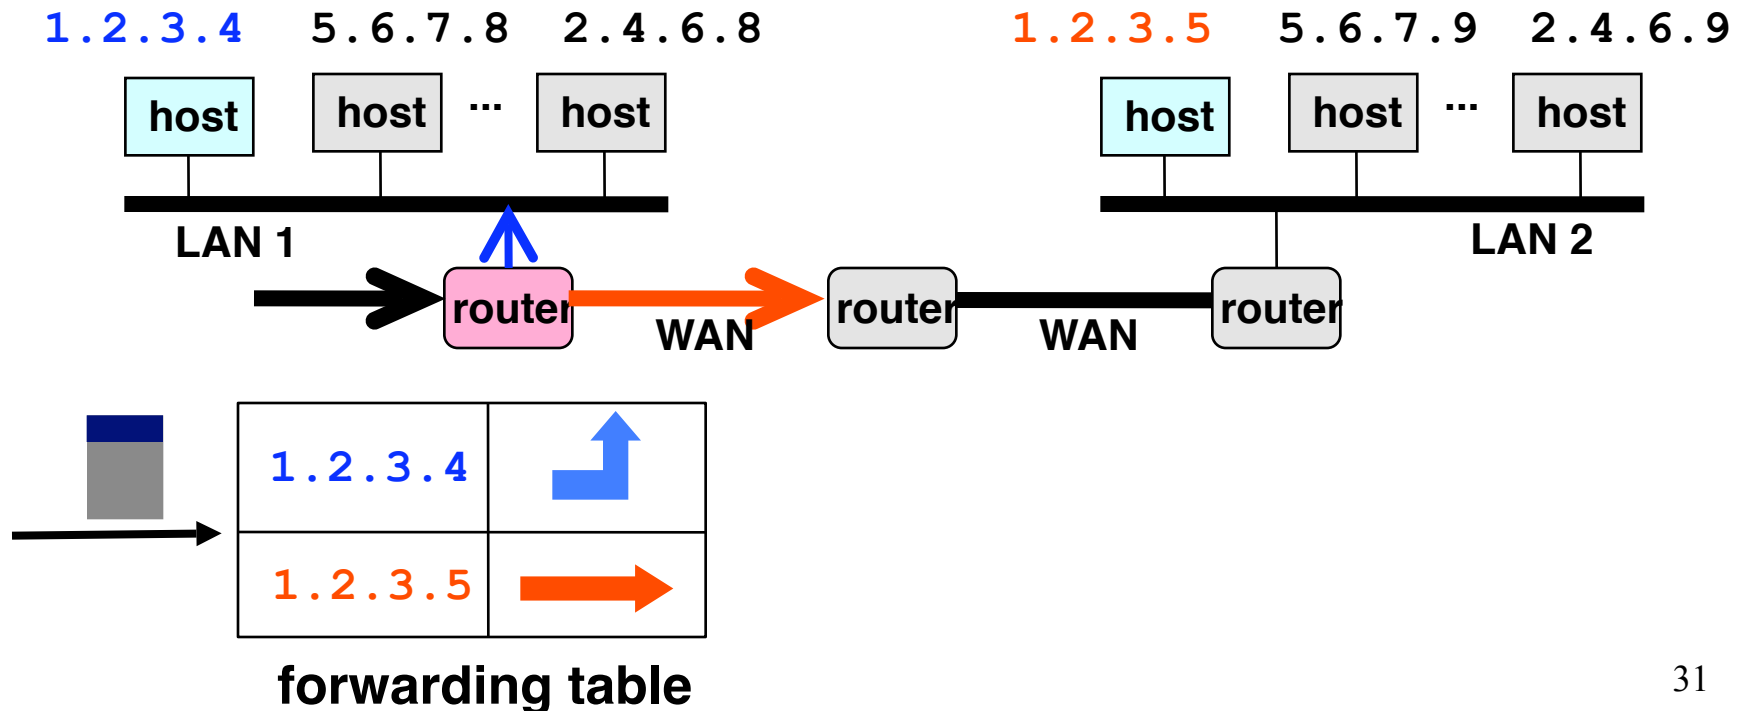
- Each router has a forwarding table
  - Maps destination addresses...
  - ... to outgoing interfaces
- Upon receiving a packet
  - Inspect the destination IP address in the header
  - Index into the table
  - Determine the outgoing interface
  - Forward the packet out that interface
- Then, the next router in the path repeats
  - And the packet travels along the path to the destination





# Separate Table Entries Per Address

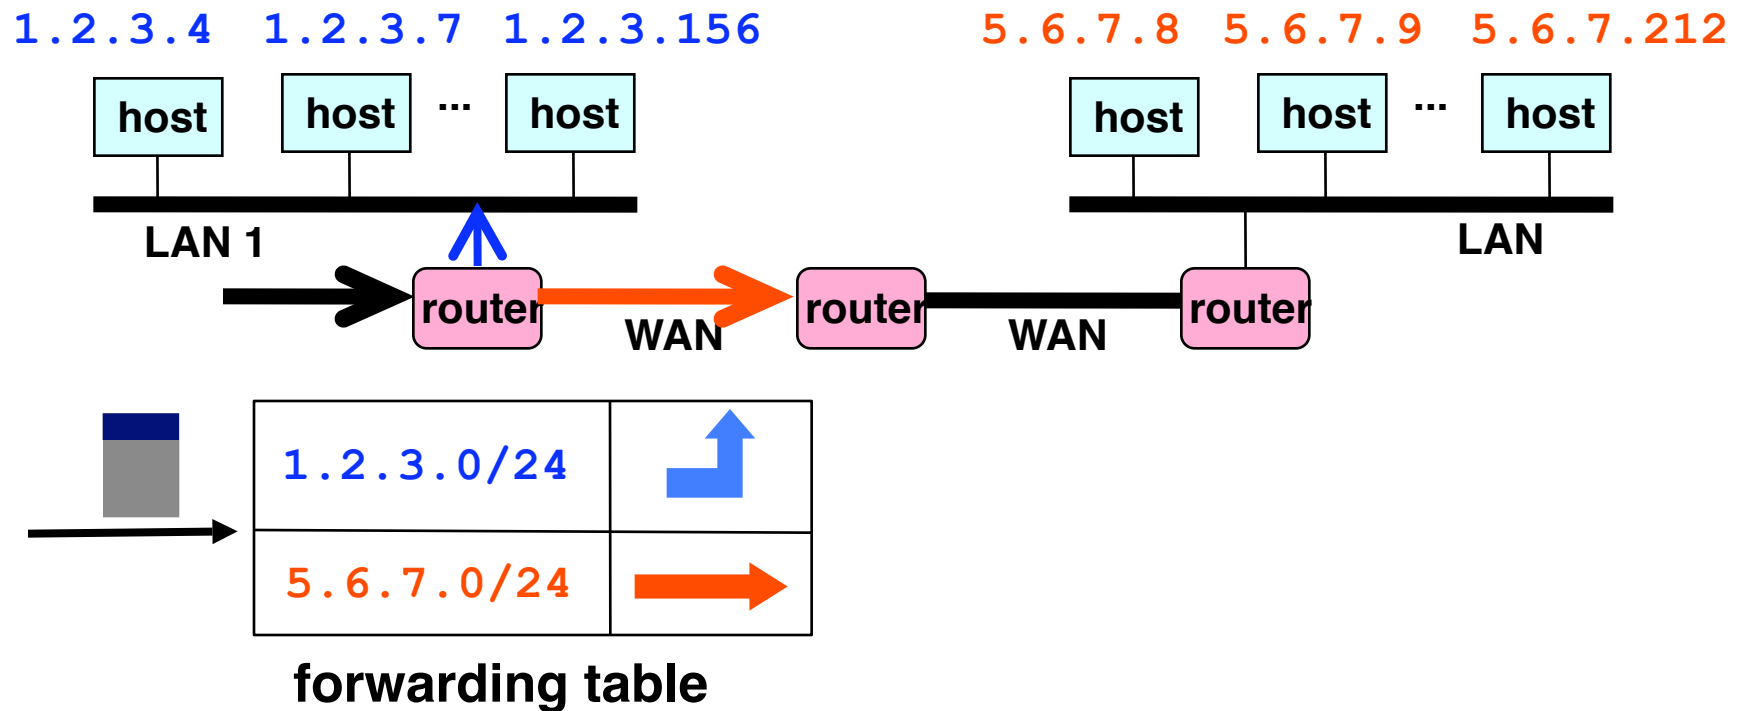
- If a router had a forwarding entry per IP address
  - Match *destination address* of incoming packet
  - ... to the *forwarding-table entry*
  - ... to determine the *outgoing interface*





# Separate Entry Per 24-bit Prefix

- If the router had an entry per 24-bit prefix
  - Look only at the top 24 bits of the destination address
  - Index into the table to determine the next-hop interface





# Separate Entry Classful Address




- If the router had an entry per classful prefix
  - Mixture of Class A, B, and C addresses
  - Depends on the first couple of bits of the destination
- Identify the mask automatically from the address
  - First bit of 0: class A address (/8)
  - First two bits of 10: class B address (/16)
  - First three bits of 110: class C address (/24)
- Then, look in the forwarding table for the match
  - E.g., 129.2.3.4 maps to 129.2.3.0/24
  - Then, look up the entry for 129.2.3.0/24
  - ... to identify the outgoing interface

# CIDR Makes Packet Forwarding Harder




- There's no such thing as a free lunch
  - CIDR allows efficient use of the limited address space
  - But, CIDR makes packet forwarding much harder
- Forwarding table may have many matches
  - E.g., table entries for 201.10.0.0/21 and 201.10.6.0/23
  - The IP address 201.10.6.17 would match *both*!





# Longest Prefix Match Forwarding

- Forwarding tables in IP routers
  - Maps each IP prefix to next-hop link(s)
- Destination-based forwarding
  - Packet has a destination address
  - Router identifies longest-matching prefix
  - Cute algorithmic problem: very fast lookups




# Simplest Algorithm is Too Slow



- Scan the forwarding table one entry at a time
  - See if the destination matches the entry
  - If so, check the size of the mask for the prefix
  - Keep track of the entry with longest-matching prefix
- Overhead is linear in size of the forwarding table
  - Today, that means 200,000 entries!
  - And, the router may have just a few nanoseconds
  - ... before the next packet is arriving
- Need greater efficiency to keep up with *line rate*
  - Better algorithms
  - Hardware implementations

# Patricia Tree

- Store the prefixes as a tree
  - One bit for each level of the tree
  - Some nodes correspond to valid prefixes
  - ... which have next-hop interfaces in a table
- When a packet arrives
  - Traverse the tree based on the destination address
  - Stop upon reaching the longest matching prefix





# Even Faster Lookups

- Patricia tree is faster than linear scan
  - Proportional to number of bits in the address
- Patricia tree can be made faster
  - Can make a k-ary tree
    - E.g., 4-ary tree with four children (00, 01, 10, and 11)
  - Faster lookup, though requires more space
- Can use special hardware
- Huge innovations in the mid-to-late 1990s
  - After CIDR was introduced (in 1994)
  - ... and longest-prefix match was a major bottleneck



# Where do Forwarding Tables Come From?

- Routers have forwarding tables
  - Map prefix to outgoing link(s)
- Entries can be statically configured
  - E.g., “map 12.34.158.0/24 to Serial0/0.1”
- But, this doesn’t adapt
  - To failures
  - To new equipment
  - To the need to balance load
  - ...
- That is where other technologies come in...
  - Routing protocols, DHCP, and ARP (later in course)

# How Do End Hosts Forward Packets?



- End host with single network interface
  - PC with an Ethernet link
  - Laptop with a wireless link
- Don't need to run a routing protocol
  - Packets to the host itself (e.g., 1.2.3.4/32)
    - Delivered locally
  - Packets to other hosts on the LAN (e.g., 1.2.3.0/24)
    - Sent out the interface
  - Packets to external hosts (e.g., 0.0.0.0/0)
    - Sent out interface to local gateway
- How this information is learned
  - Static setting of address, subnet mask, and gateway
  - Dynamic Host Configuration Protocol (DHCP)






# What About Reaching the End Hosts?



- How does the last router reach the destination?



- Each interface has a persistent, global identifier
  - MAC (Media Access Control) address
  - Burned in to the adaptors Read-Only Memory (ROM)
  - Flat address structure (i.e., no hierarchy)
- Constructing an address resolution table
  - Mapping MAC address to/from IP address
  - Address Resolution Protocol (ARP)



# Conclusions

- IP address
  - A 32-bit number
  - Allocated in prefixes
  - Non-uniform hierarchy for scalability and flexibility
- Packet forwarding
  - Based on IP prefixes
  - Longest-prefix-match forwarding
- Next lecture
  - Transmission Control Protocol (TCP)
- We'll cover some topics later
  - Routing protocols, DHCP, and ARP