

CS/ECE 4457

Computer Networks: Architecture and Protocols

Lecture 19

The IP protocol

DNS, Discovery protocols

Putting ALL the Pieces Together

Qizhe Cai



Goal of today's lecture

- **A brief introduction to Domain Name System**
- **Discovery protocols**
- **End-to-end: how everything fits together**

Domain Name System (DNS)

What is DNS?

- **User has name of entity she/he wants to access**
 - E.g., www.cnn.com
 - Content, host, etc.
- **However, Internet routes and forwards requests based on IP addresses**
 - Need to convert name (e.g., www.cnn.com) to an IP address
- **Domain Name System (DNS)**
 - **Provides the mapping from name to IP address**
 - User asks DNS: what is the IP address for www.cnn.com
 - DNS responds: 157.166.255.18

Correctness Requirements

- **Addresses can change underneath**
 - Move www.cnn.com to 4.125.91.21
 - Humans/Applications should be unaffected
- **Name could map to multiple IP addresses**
 - www.cnn.com to multiple replicas to the Web site
 - To enable “load balancing” or reduced latency
 - Replicas may see different load (eg, due to geographic location)
 - Some replicas may be closer to the user
- **Multiple names for the same address**
 - E.g., www.cnn.com and cnn.com should map to same IP addresses

Goals and Approach

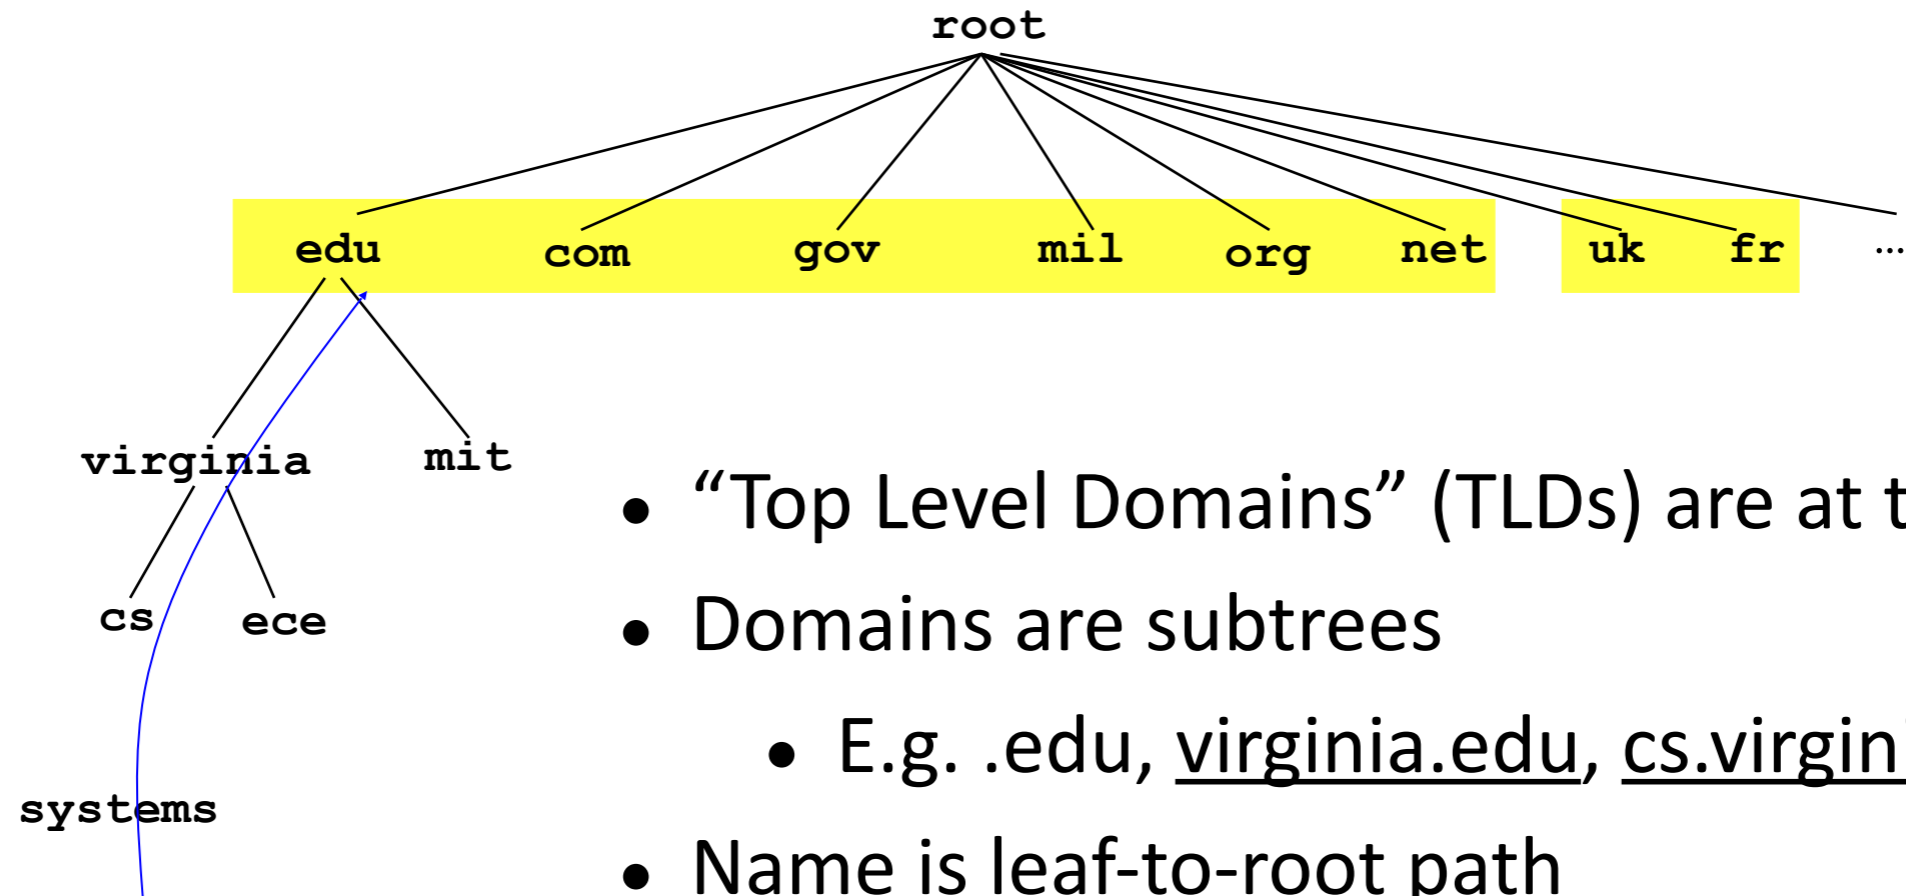
- **Goals**

- Correctness (from previous slide)
- Scaling (names, users, updates, etc.)
- Ease of management (uniqueness of names, etc.)
- Availability and consistency
- Fast lookups

- **Approach: Three intertwined hierarchies**

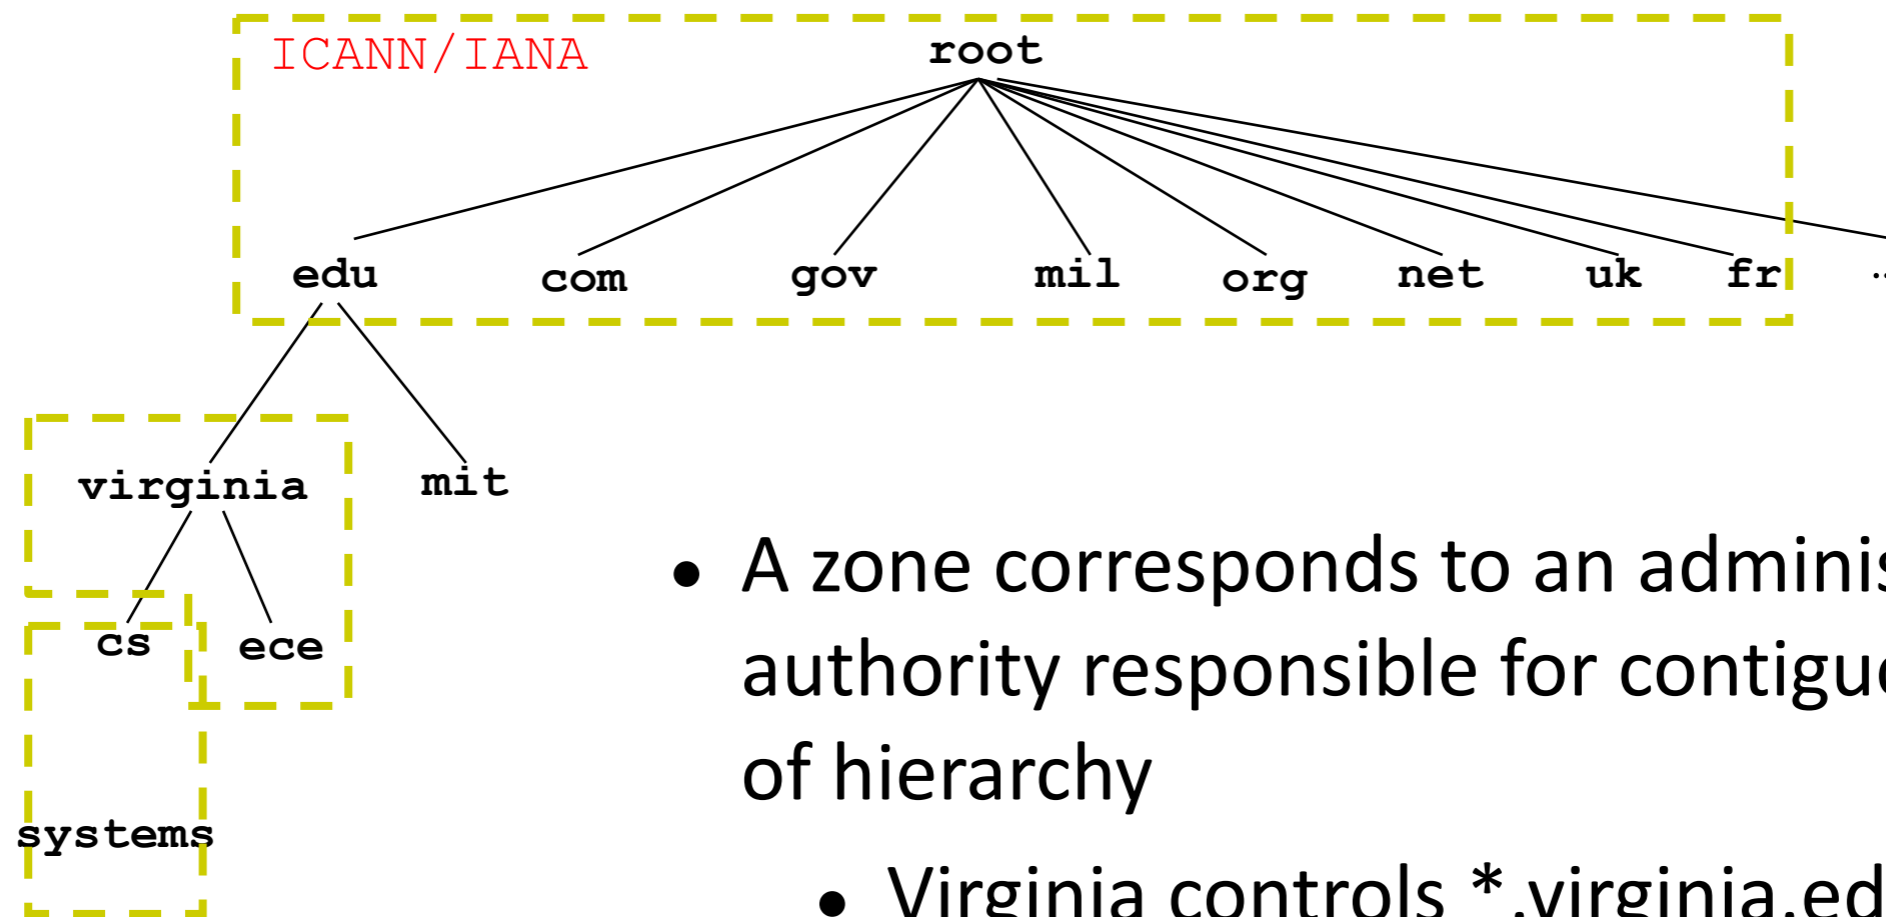
- Hierarchical Namespace: exploit structure in names
- Hierarchical Administration: hierarchy of authority over names
- Hierarchical Infrastructure: hierarchy of DNS servers

Hierarchical Namespace



- “Top Level Domains” (TLDs) are at the top
- Domains are subtrees
 - E.g. .edu, virginia.edu, cs.virginia.edu
- Name is leaf-to-root path
 - systems.cs.virginia.edu

Hierarchical Administration



- A zone corresponds to an administrative authority responsible for contiguous portion of hierarchy
 - Virginia controls *.virginia.edu
 - CS controls *.cs.virginia.edu
- Name collisions trivially avoided
 - Each domain can ensure this locally

Hierarchical Infrastructure

- Top of hierarchy: root
 - Location hardwired into other servers
- Next level: Top Level Domain (TLD) servers
 - .com, .edu, etc.
- Bottom level: Authoritative DNS servers
 - Actually do the mapping
 - Can be maintained locally or by a service provider

Who Knows What?


- Every server knows address of root name server
- Root servers know the address of all TLD servers
- Every node knows the address of all children
- An ***authoritative*** DNS server stores name-to-address mappings (“resource records”) for all DNS names in the domain that it has authority for
- Therefore, each server:
 - Stores only a subset of the total DNS database (scalable!)
 - Can discover server(s) for any portion of the hierarchy

Using DNS

- Two components
 - Local DNS servers
 - Resolver software on hosts
- Local DNS server (“default name server”)
 - Usually near the end hosts that use it
 - Local hosts configured with local server (e.g, /etc/resolv.conf) or learn server via DHCP (to be discussed soon)
- Client application
 - Obtain DNS name (e.g., from the URL)
 - Do **gethostbyname()** to trigger resolver code
 - Which then sends request to local DNS server

**Local DNS
server**


(mydns.virginia.edu)



DNS client

(me.cs.virginia.edu)

root servers



.edu servers



**nyu.edu
servers**

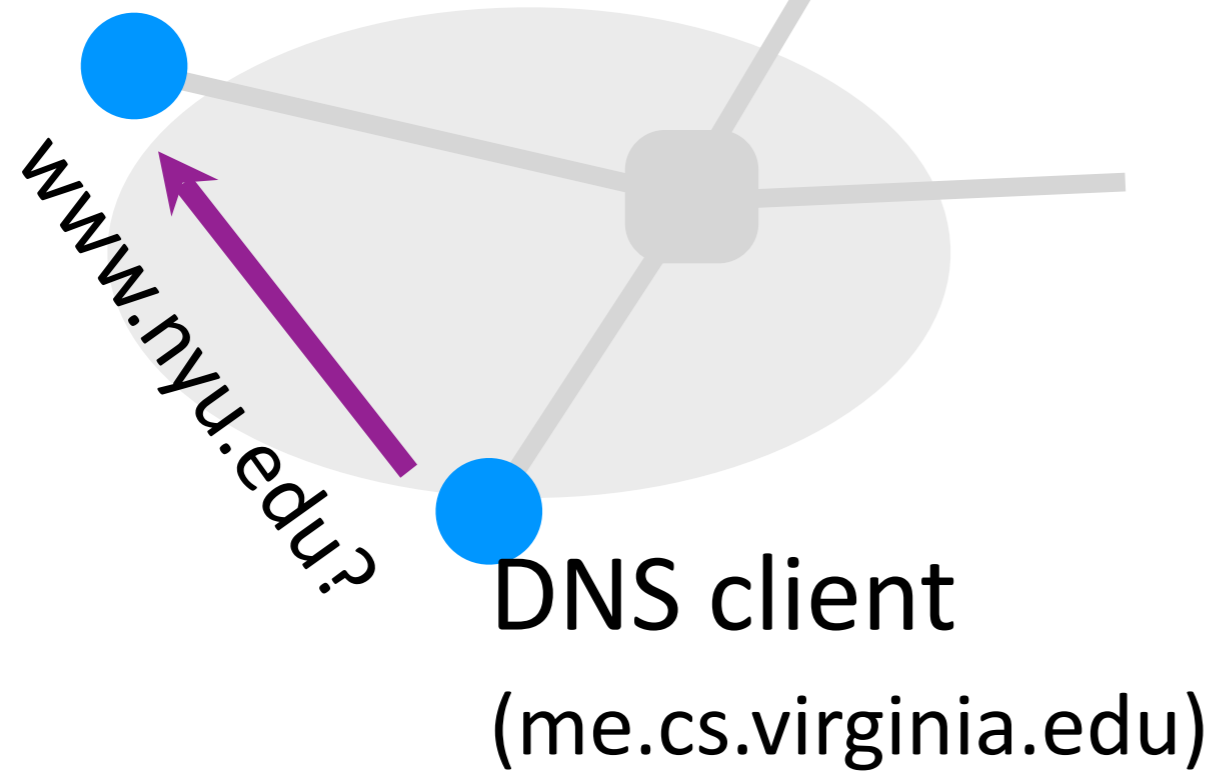
**Local DNS
server**

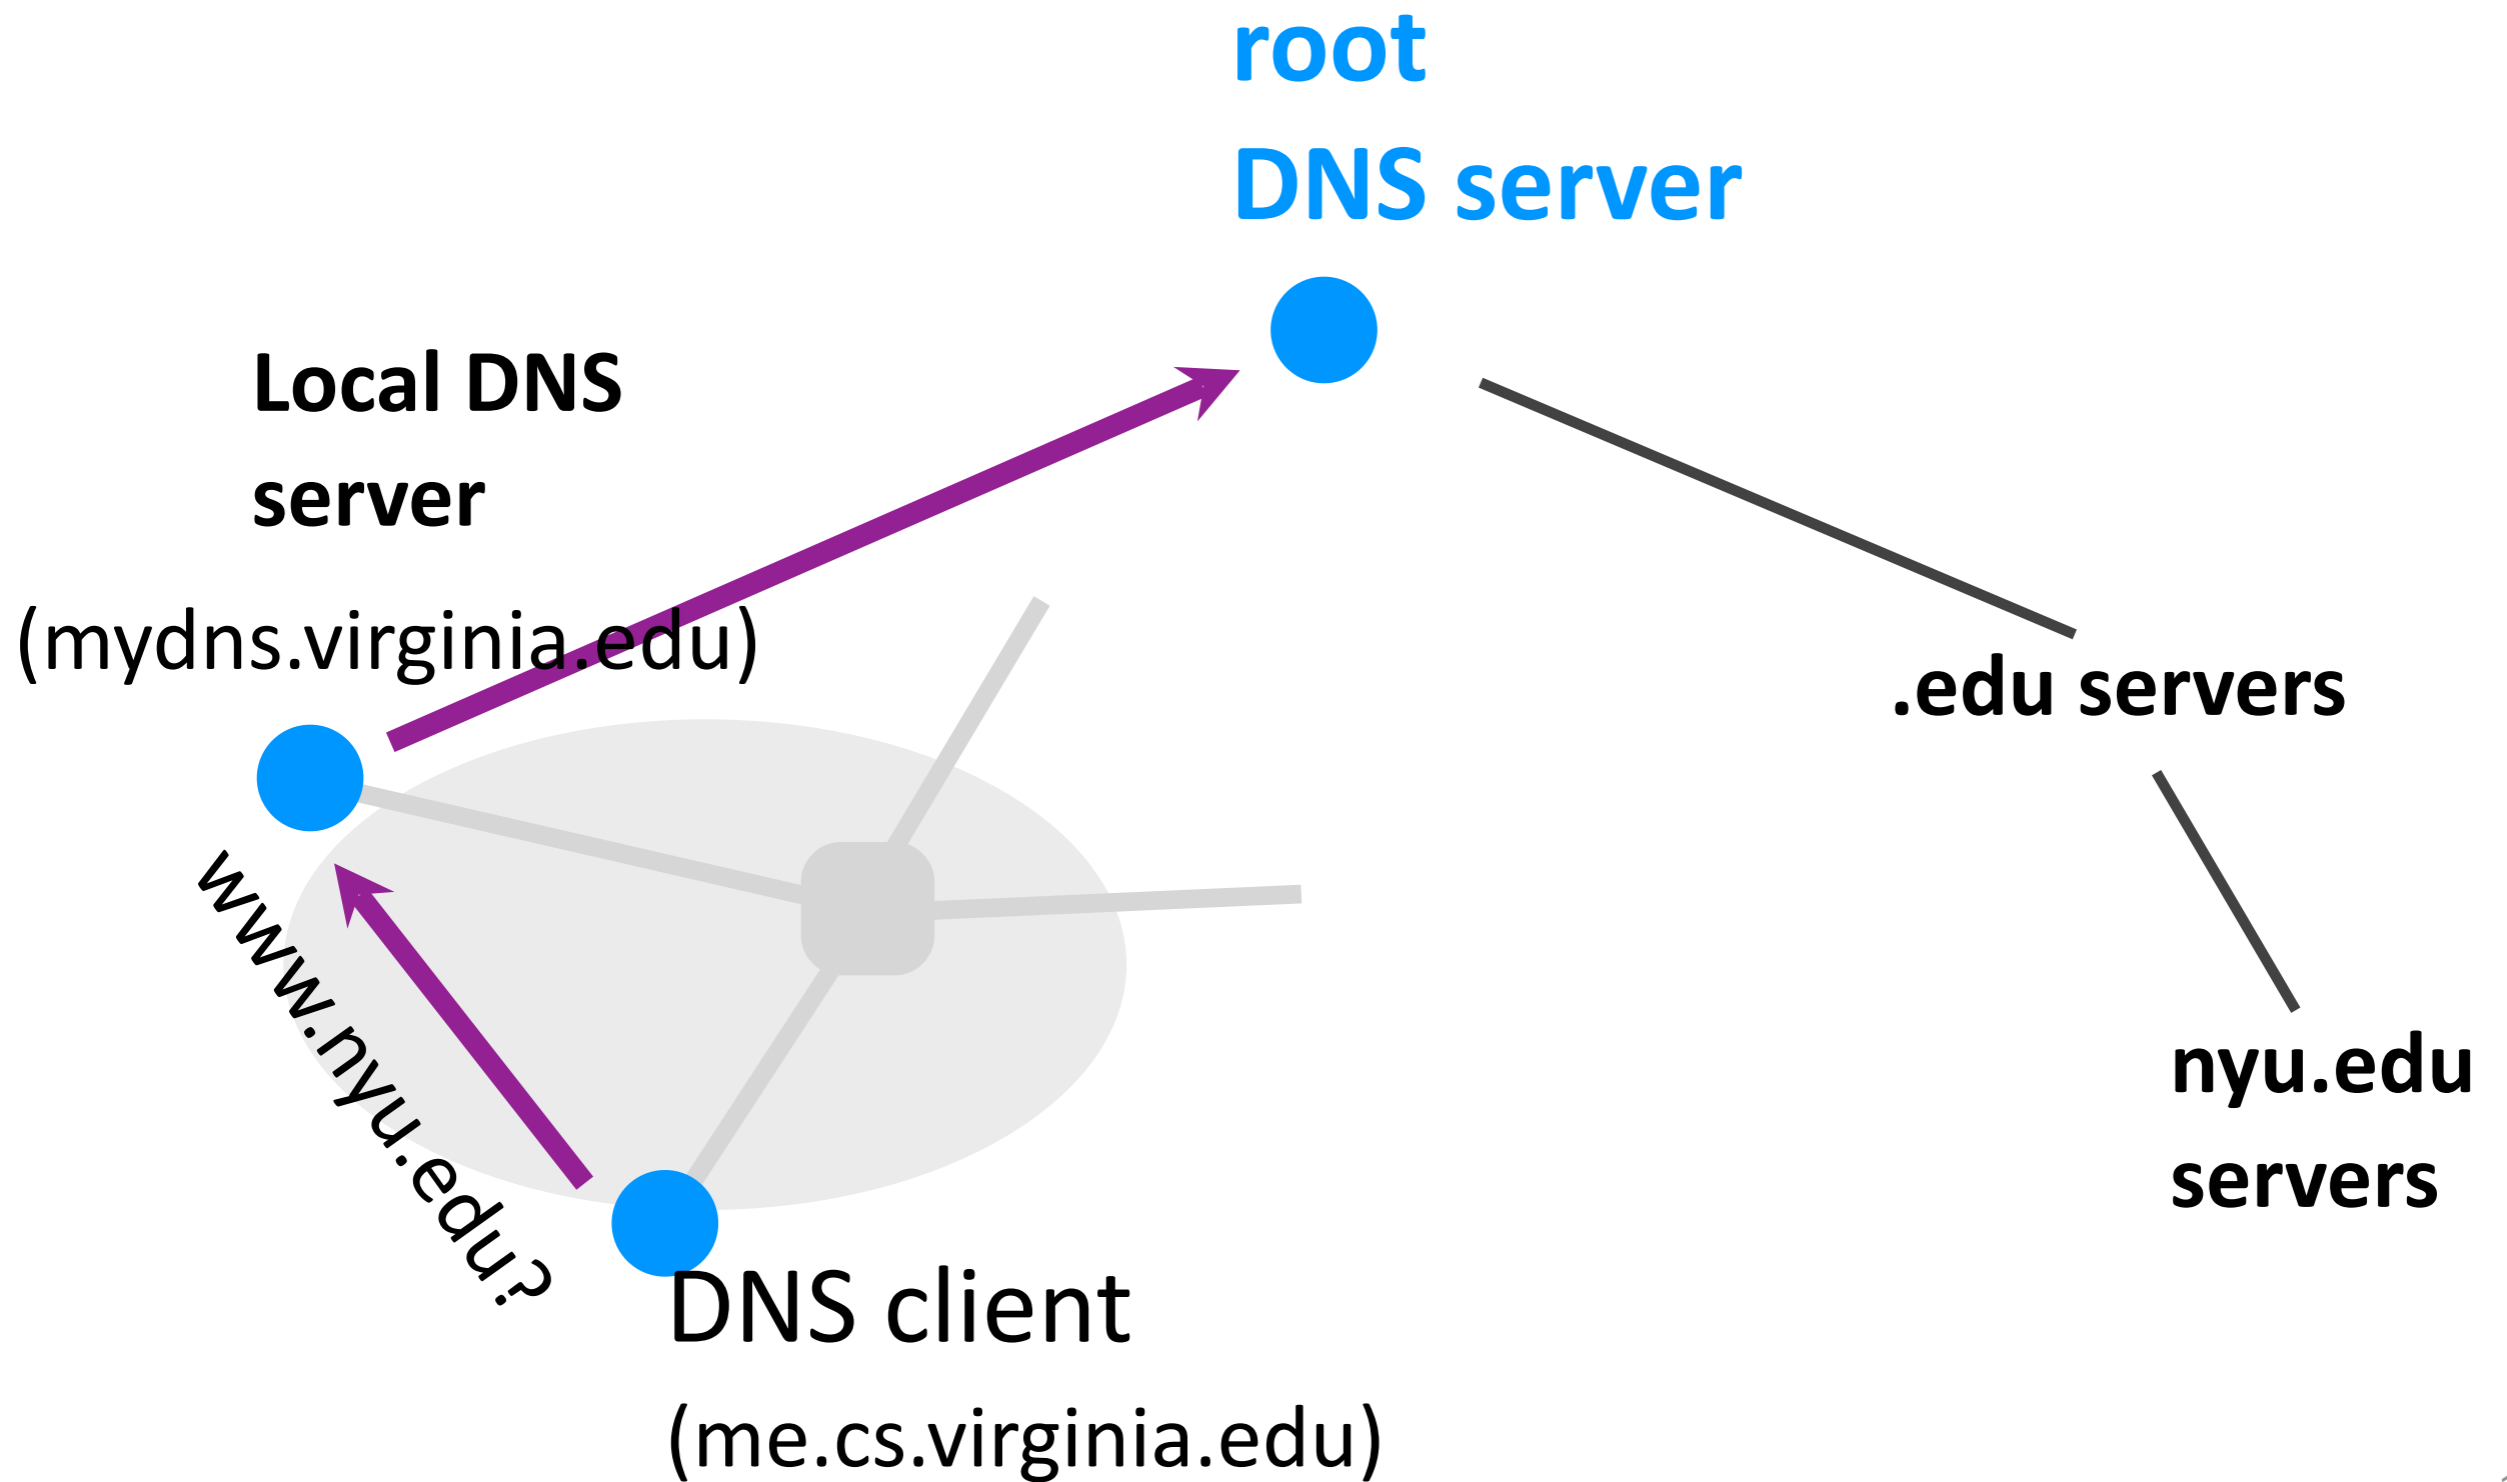
(mydns.virginia.edu)


root servers

.edu servers


**nyu.edu
servers**





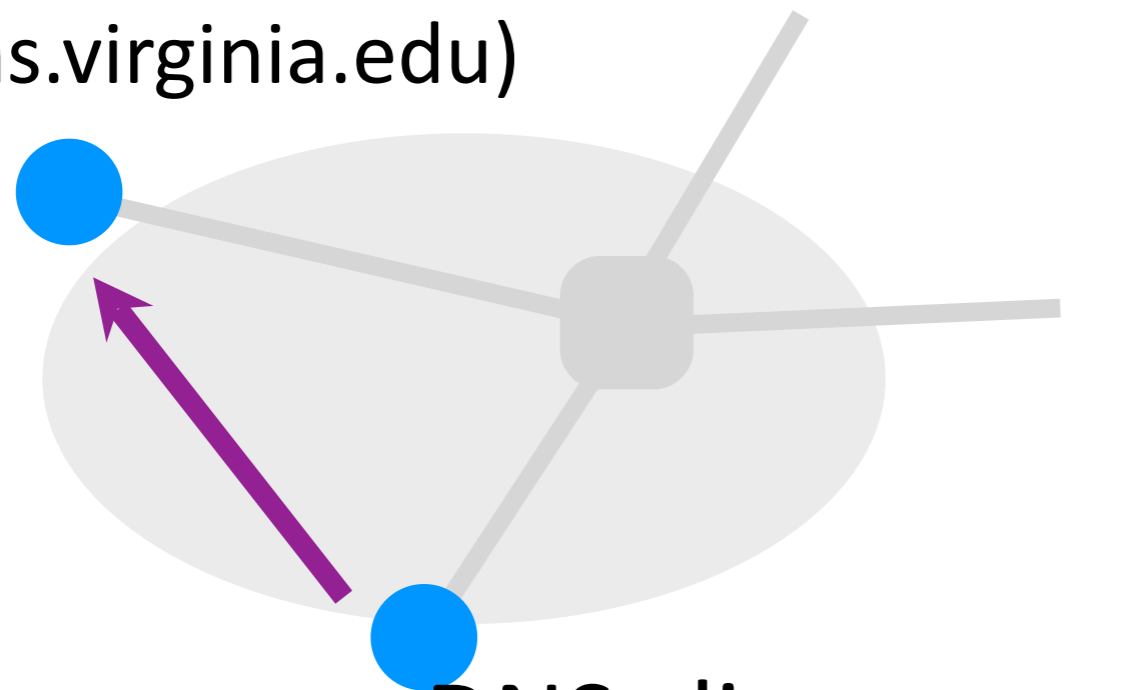


recursive DNS query



**Local DNS
server**

(mydns.virginia.edu)



DNS client

(me.cs.virginia.edu)

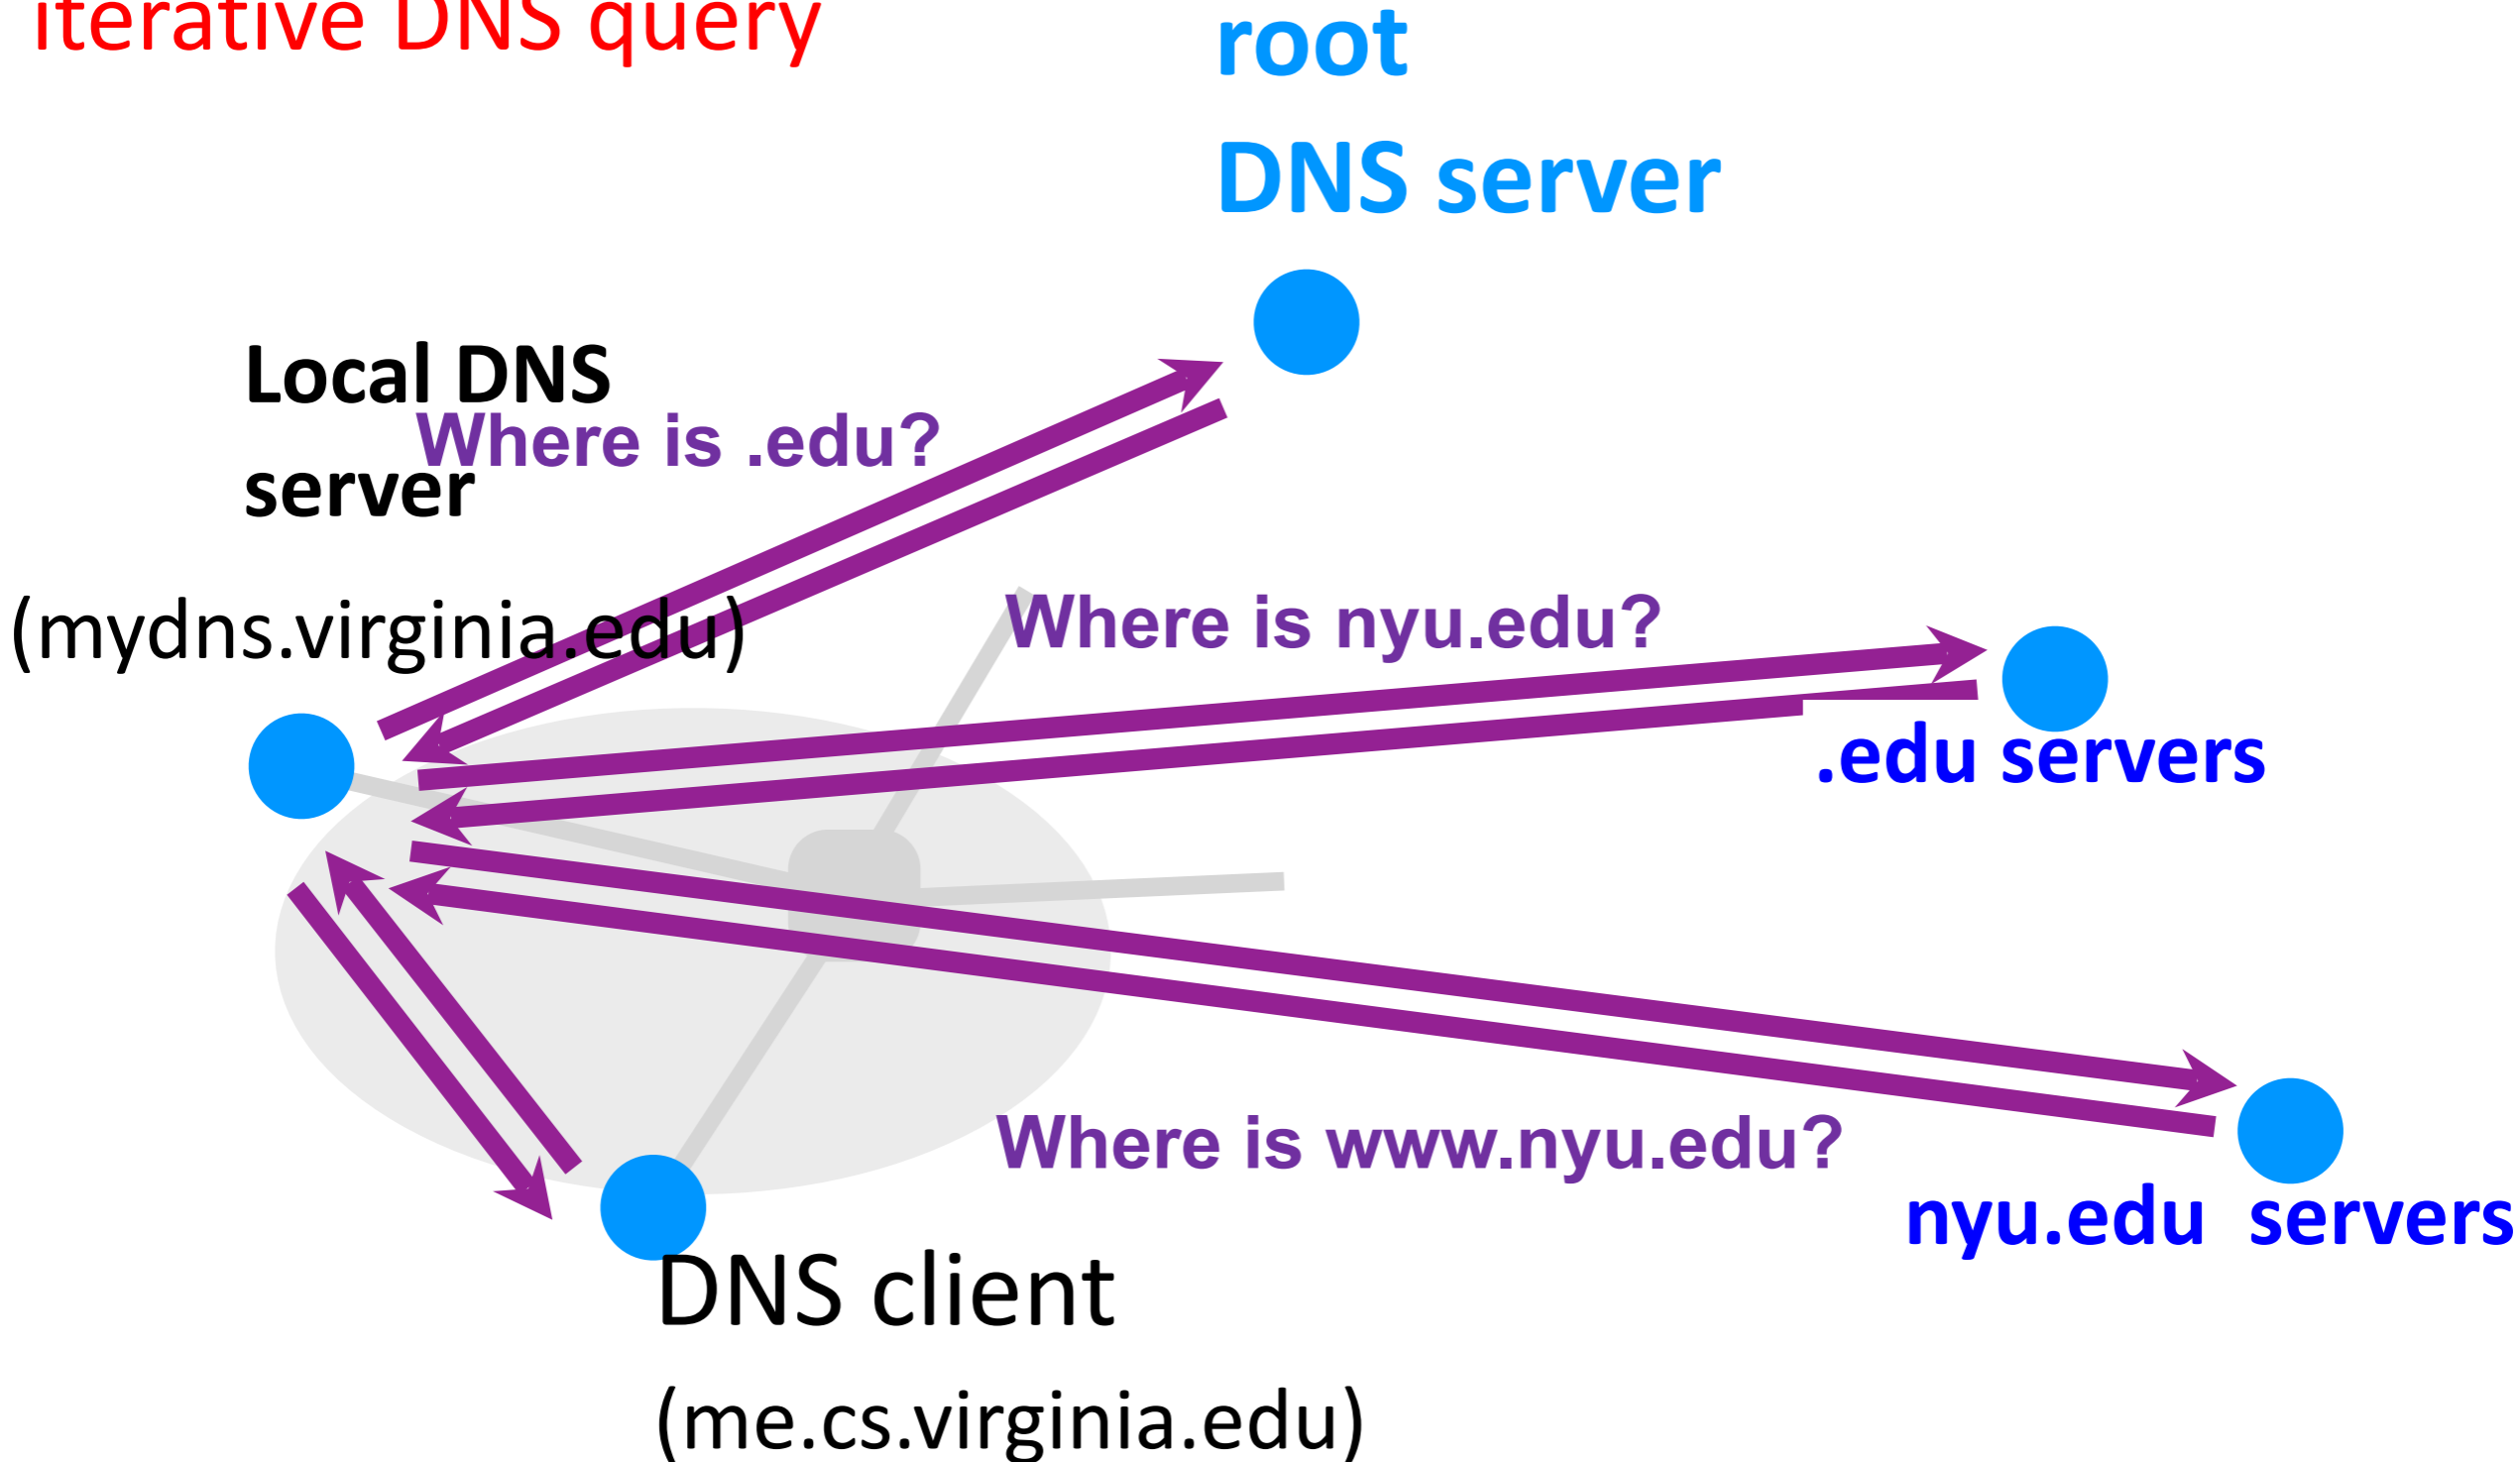
**root
DNS server**



.edu servers

nyu.edu servers

iterative DNS query



Discovery Protocols


Suppose Host A wants to communication with Host B

Discovery

- Suppose I am host A
- I want to communicate with B (say, www.google.com)
- I was “born” knowing **only** my name — my MAC address :-)
- Must discover some information before I can communicate with B
 - What is my IP address?
 - What is B’s IP address?
 - Using DNS
 - Is B within my LAN?
 - If yes, what is B’s MAC address?
 - If not, what is the address of my first-hop router to B?
 - ...

DHCP and ARP

- Link layer discovery protocols
 - DHCP — Dynamic Host Configuration Protocol
 - ARP — Address Resolution Protocol
 - Configured to a single LAN
 - Rely on broadcast capability



DHCP and ARP

- Link layer discovery protocols
- Serve two functions
 1. Discovery of local end-hosts
 - For communication between hosts on the same LAN
 2. Bootstrap communication with remote hosts
 - What's my IP address?
 - Who/where is my local DNS server?
 - Who/where is my first hop router?

DHCP

- Dynamic Host Configuration Protocol
 - Defined in RFC 2131
- A host uses DHCP to discover
 - Its own IP address
 - Subnet masks — allows to test whether an IP address is local or not
 - IP address(es) for its local DNS name server(s)
 - IP address(es) for its first-hop “default” router(s)

DHCP: operation

1. One or more local DHCP servers maintain required information
 - IP address pool, netmask, DNS servers, etc.
 - Application that listens on UDP port 67

DHCP: operation

1. One or more local DHCP servers maintain required information
2. Client broadcasts a DHCP discovery message
 - L2 broadcast, to MAC address FF:FF:FF:FF:FF:FF

DHCP: operation

1. One or more local DHCP servers maintain required information
2. Client broadcasts a DHCP discovery message
3. One or more DHCP servers respond with a DHCP “offer” message
 - Proposed IP address for client, lease time
 - Other parameters

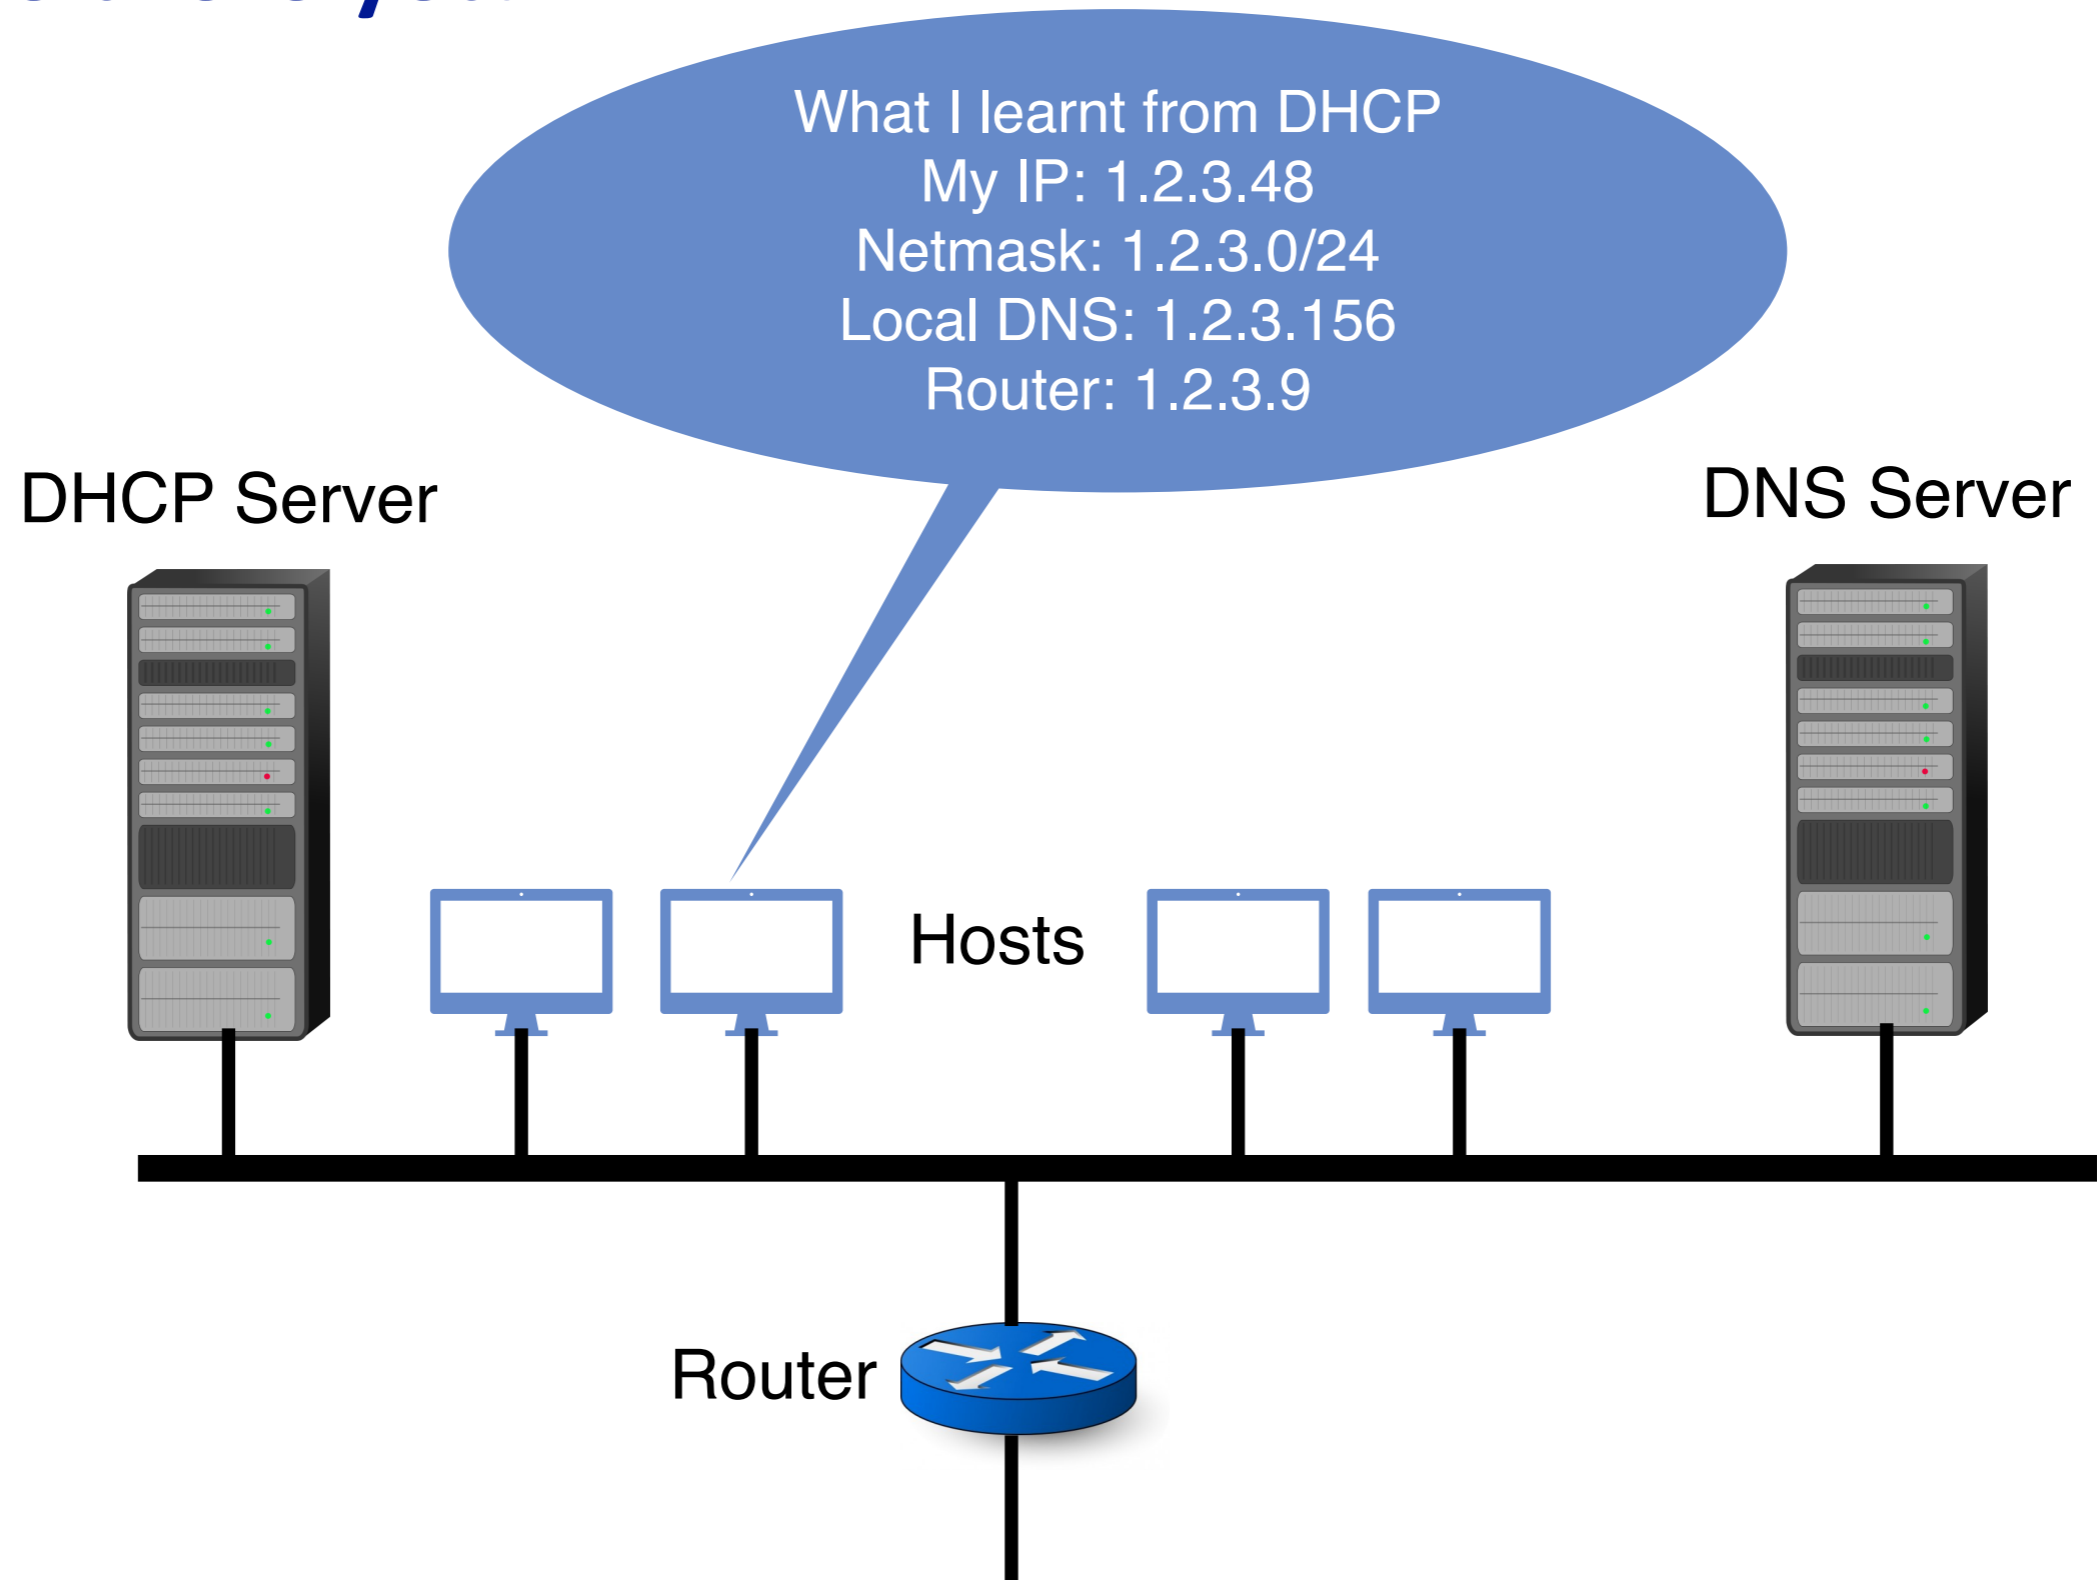
DHCP: operation

1. One or more local DHCP servers maintain required information
2. Client broadcasts a DHCP discovery message
3. One or more DHCP servers respond with a DHCP “offer” message
4. Client broadcasts a DHCP request message
 - Specifies which offer it wants
 - Echoes accepted parameters
 - Other DHCP servers learn they were not chosen

DHCP: operation

1. One or more local DHCP servers maintain required information
2. Client broadcasts a DHCP discovery message
3. One or more DHCP servers respond with a DHCP “offer” message
4. Client broadcasts a DHCP request message
5. Selected DHCP server responds with an ACK

Are we there yet?



ARP: Address Resolution Protocol

- Every host maintains an ARP table
 - List of (IP address — MAC address) pairs
 - For IP addresses within the same LAN
- Consult the table when sending a packet
 - Map destination IP address to destination MAC address
- But: what if IP address not in the table?
 - Either its not local (detected using DHCP)
 - If its local:
 - Sender broadcasts: “Who has IP address 1.2.3.156?”
 - Caches the answer in ARP table

Key Ideas in Both ARP and DHCP

- Broadcasting: can use broadcast to make contact
 - Scalable because of limited size
- Caching: remember the past for a while
 - Store the information you learn to reduce overhead


Taking Stock: Discovery

Layer	Examples	Structure	Configuration	Resolution Service
App Layer	<u>www.cs.virginia.edu</u>	Organizational hierarchy	~ manual	↕ DNS
Network Layer	123.45.6.78	Topological hierarchy	DHCP	
Link Layer	45-CC-4E-12-F0-97	Vendor(flat)	Hard-coded	↕ ARP

Putting all the pieces together


What is a computer network?

A set of network elements connected together, that implement a set of protocols for the purpose of sharing resources at the end hosts

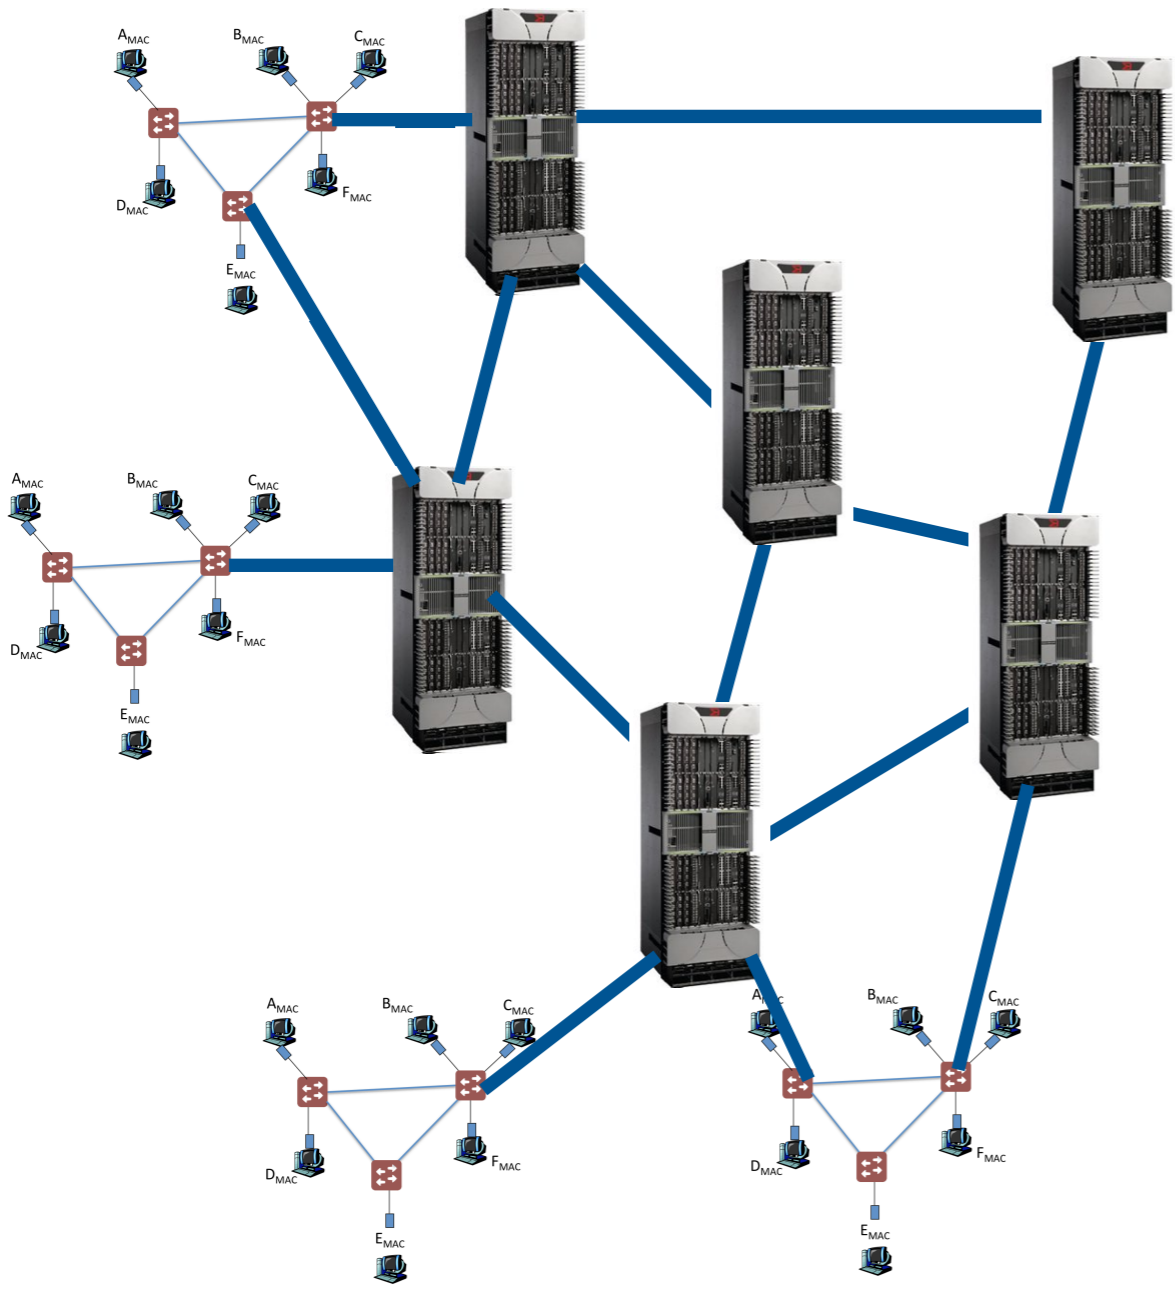


What does Internet actually look like?

- **The smallest component:**
 - A Network Interface Card (NIC), or a machine, or a server
 - Has a **Link Layer MAC name/address**
- **Multiple NICs connected in a Local Area Network (LAN) via**
 - Broadcast Ethernet,
 - Or, Switched Ethernet
- **Switches in LAN**
 - Connected to larger routers

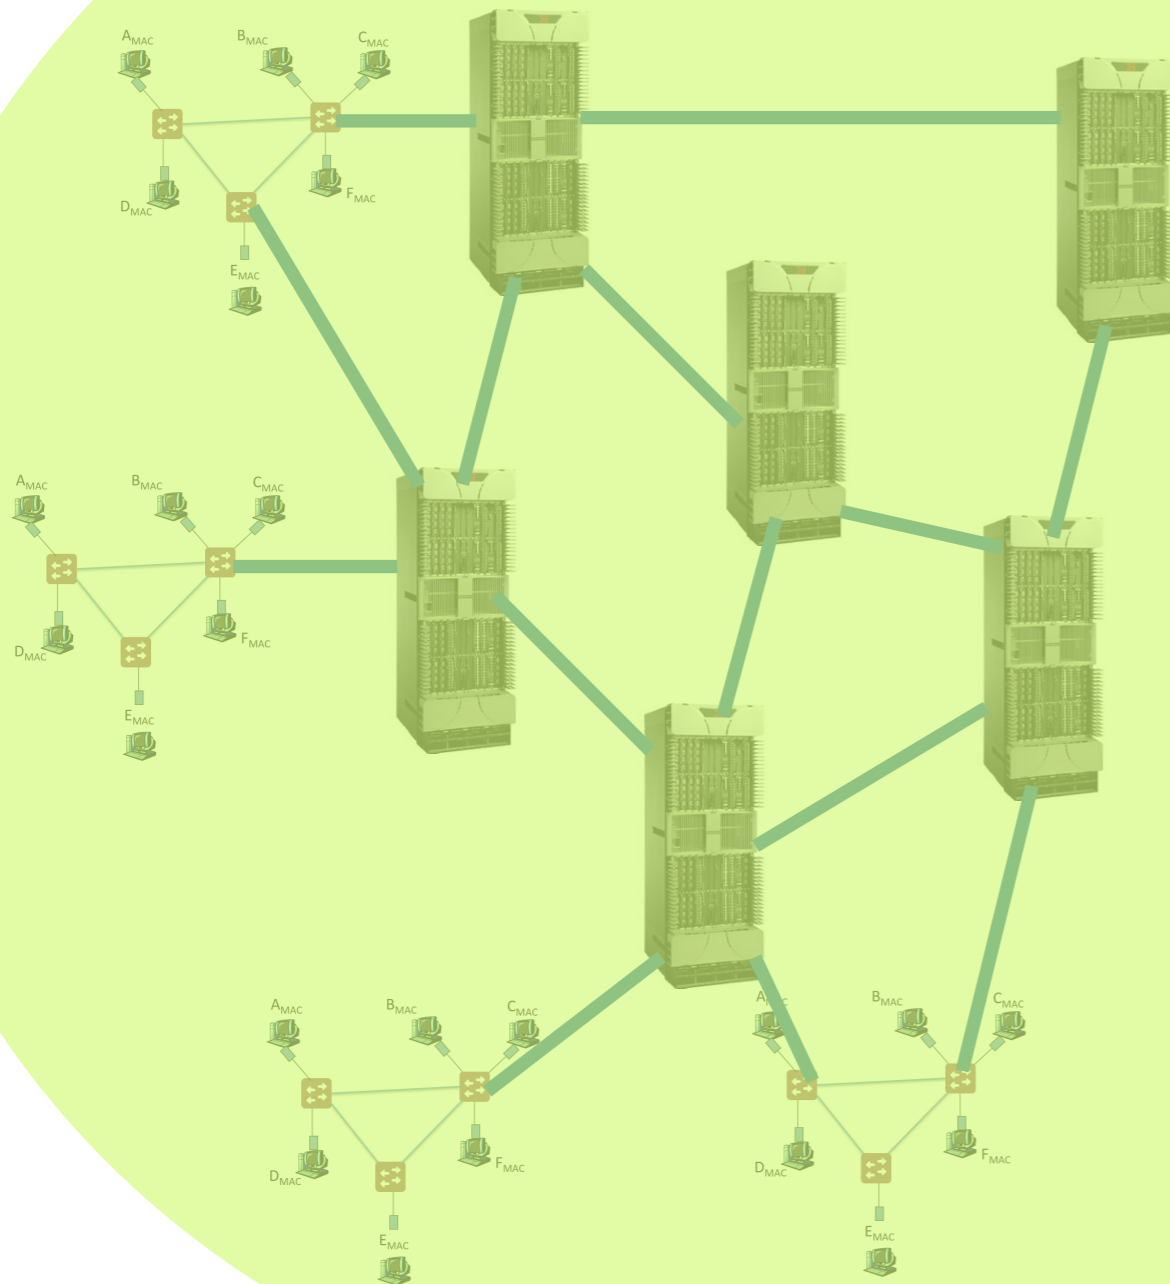


What does Internet actually look like?




What does Internet actually look like?

“Autonomous System (AS)” or “Domain”
Region of a network under a single administrative entity




What does Internet actually look like?

Multiple “Autonomous Systems (AS)” or “Domains” connect together using Border Routers



This entire infrastructure is a part of the INTERNET :-)




What is the other part of the Internet?

Protocols!


What protocols have we learnt on LAN?

- **Addresses**
 - Link Layer MAC names/addresses: come with the hardware
- **CSMA/CD Protocol:**
 - For transmitting frames on broadcast Ethernet
- **Spanning Tree Protocol:**
 - For transmitting frames on switched Ethernet



What have we learnt beyond LAN?

- **Link-state and Distance-vector Protocols:**
 - For finding routes (and a next-hop) to an IP address within an ISP
- **Border Gateway Protocol:**
 - For finding routes to an IP address range
- **Forwarding at routers**
 - Store **routing tables** (map **destination prefixes** to outgoing port)
 - Longest prefix match for destination address lookup



How does the Internet work?

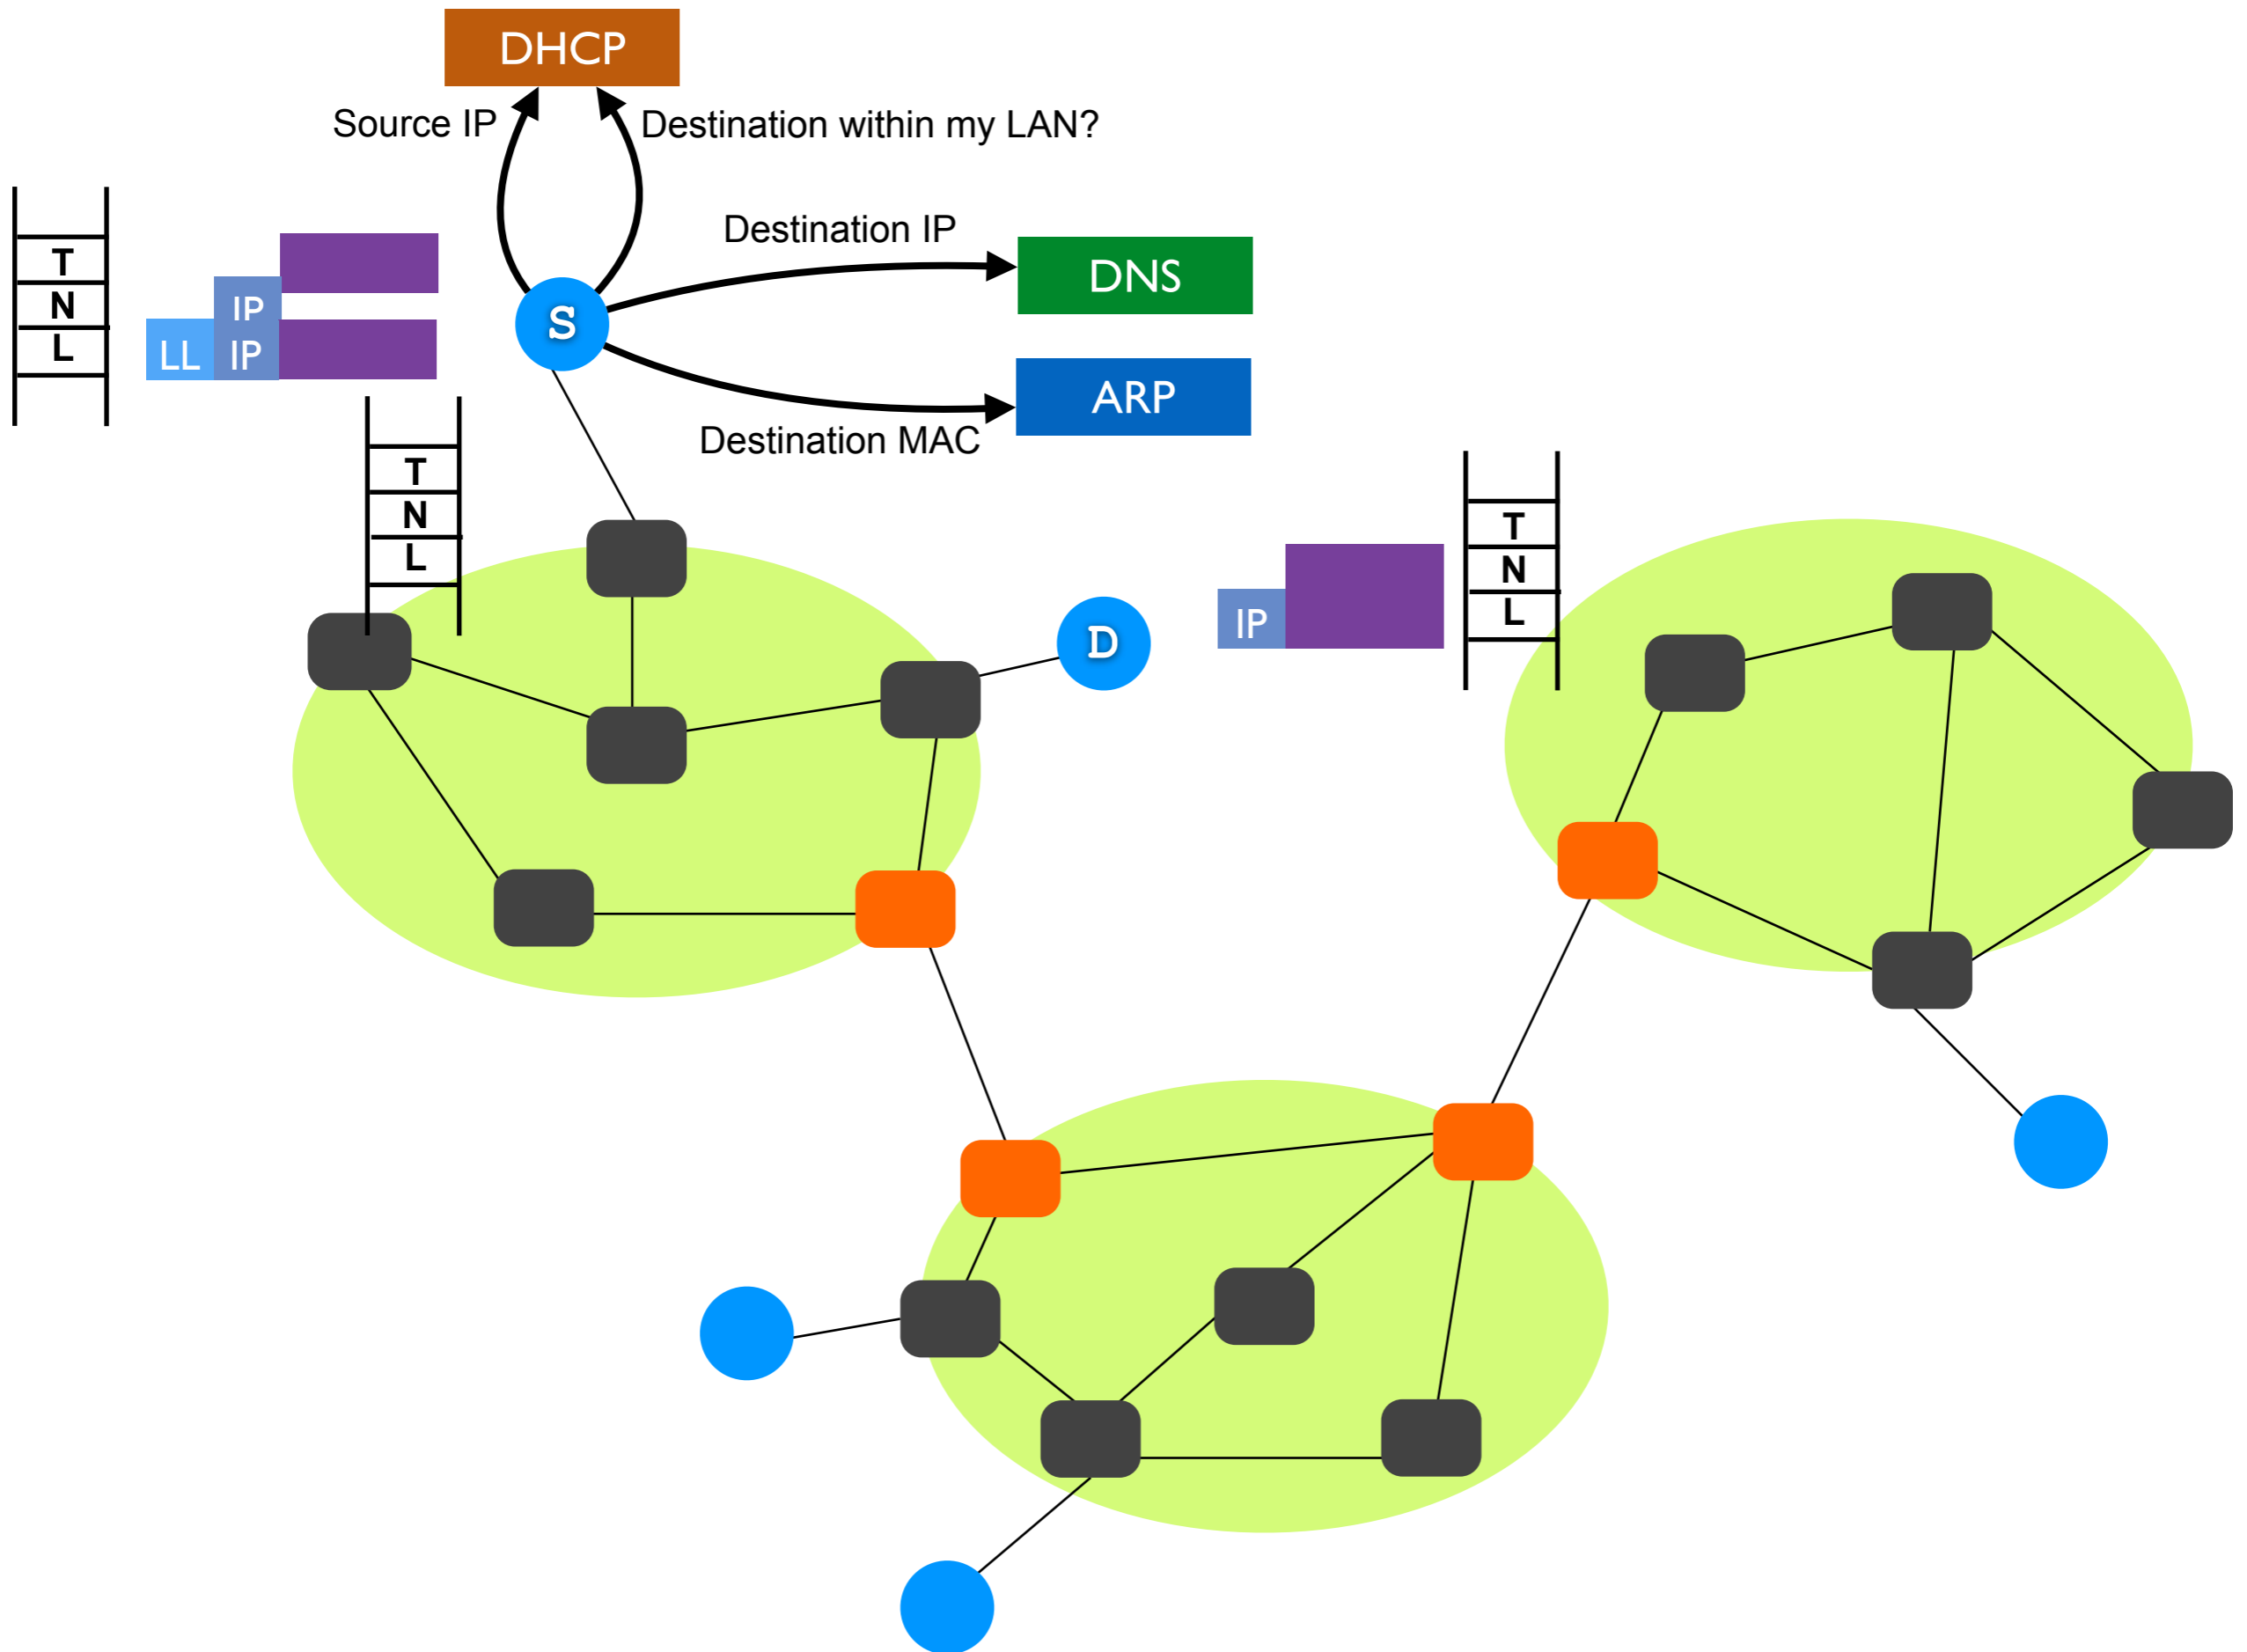
Are you ready?

(Count the number of protocols used for each packet)

How does Internet work — end-to-end?

- Network stack receives the packet from the application (roughly speaking)
- What is my IP address? (using DHCP)
- What is the destination IP address? (using DNS)
- Is destination IP address within my LAN? (using DHCP)
- **If destination IP address local:**
 - What is destination MAC address (using ARP)?
 - Convert packet into frames with correct source/destination address
 - Convert frames into bits
 - Forward the bits to the wire ...
- **Each switch:**
 - Forwards to destination (using STP/CSMA/CD)

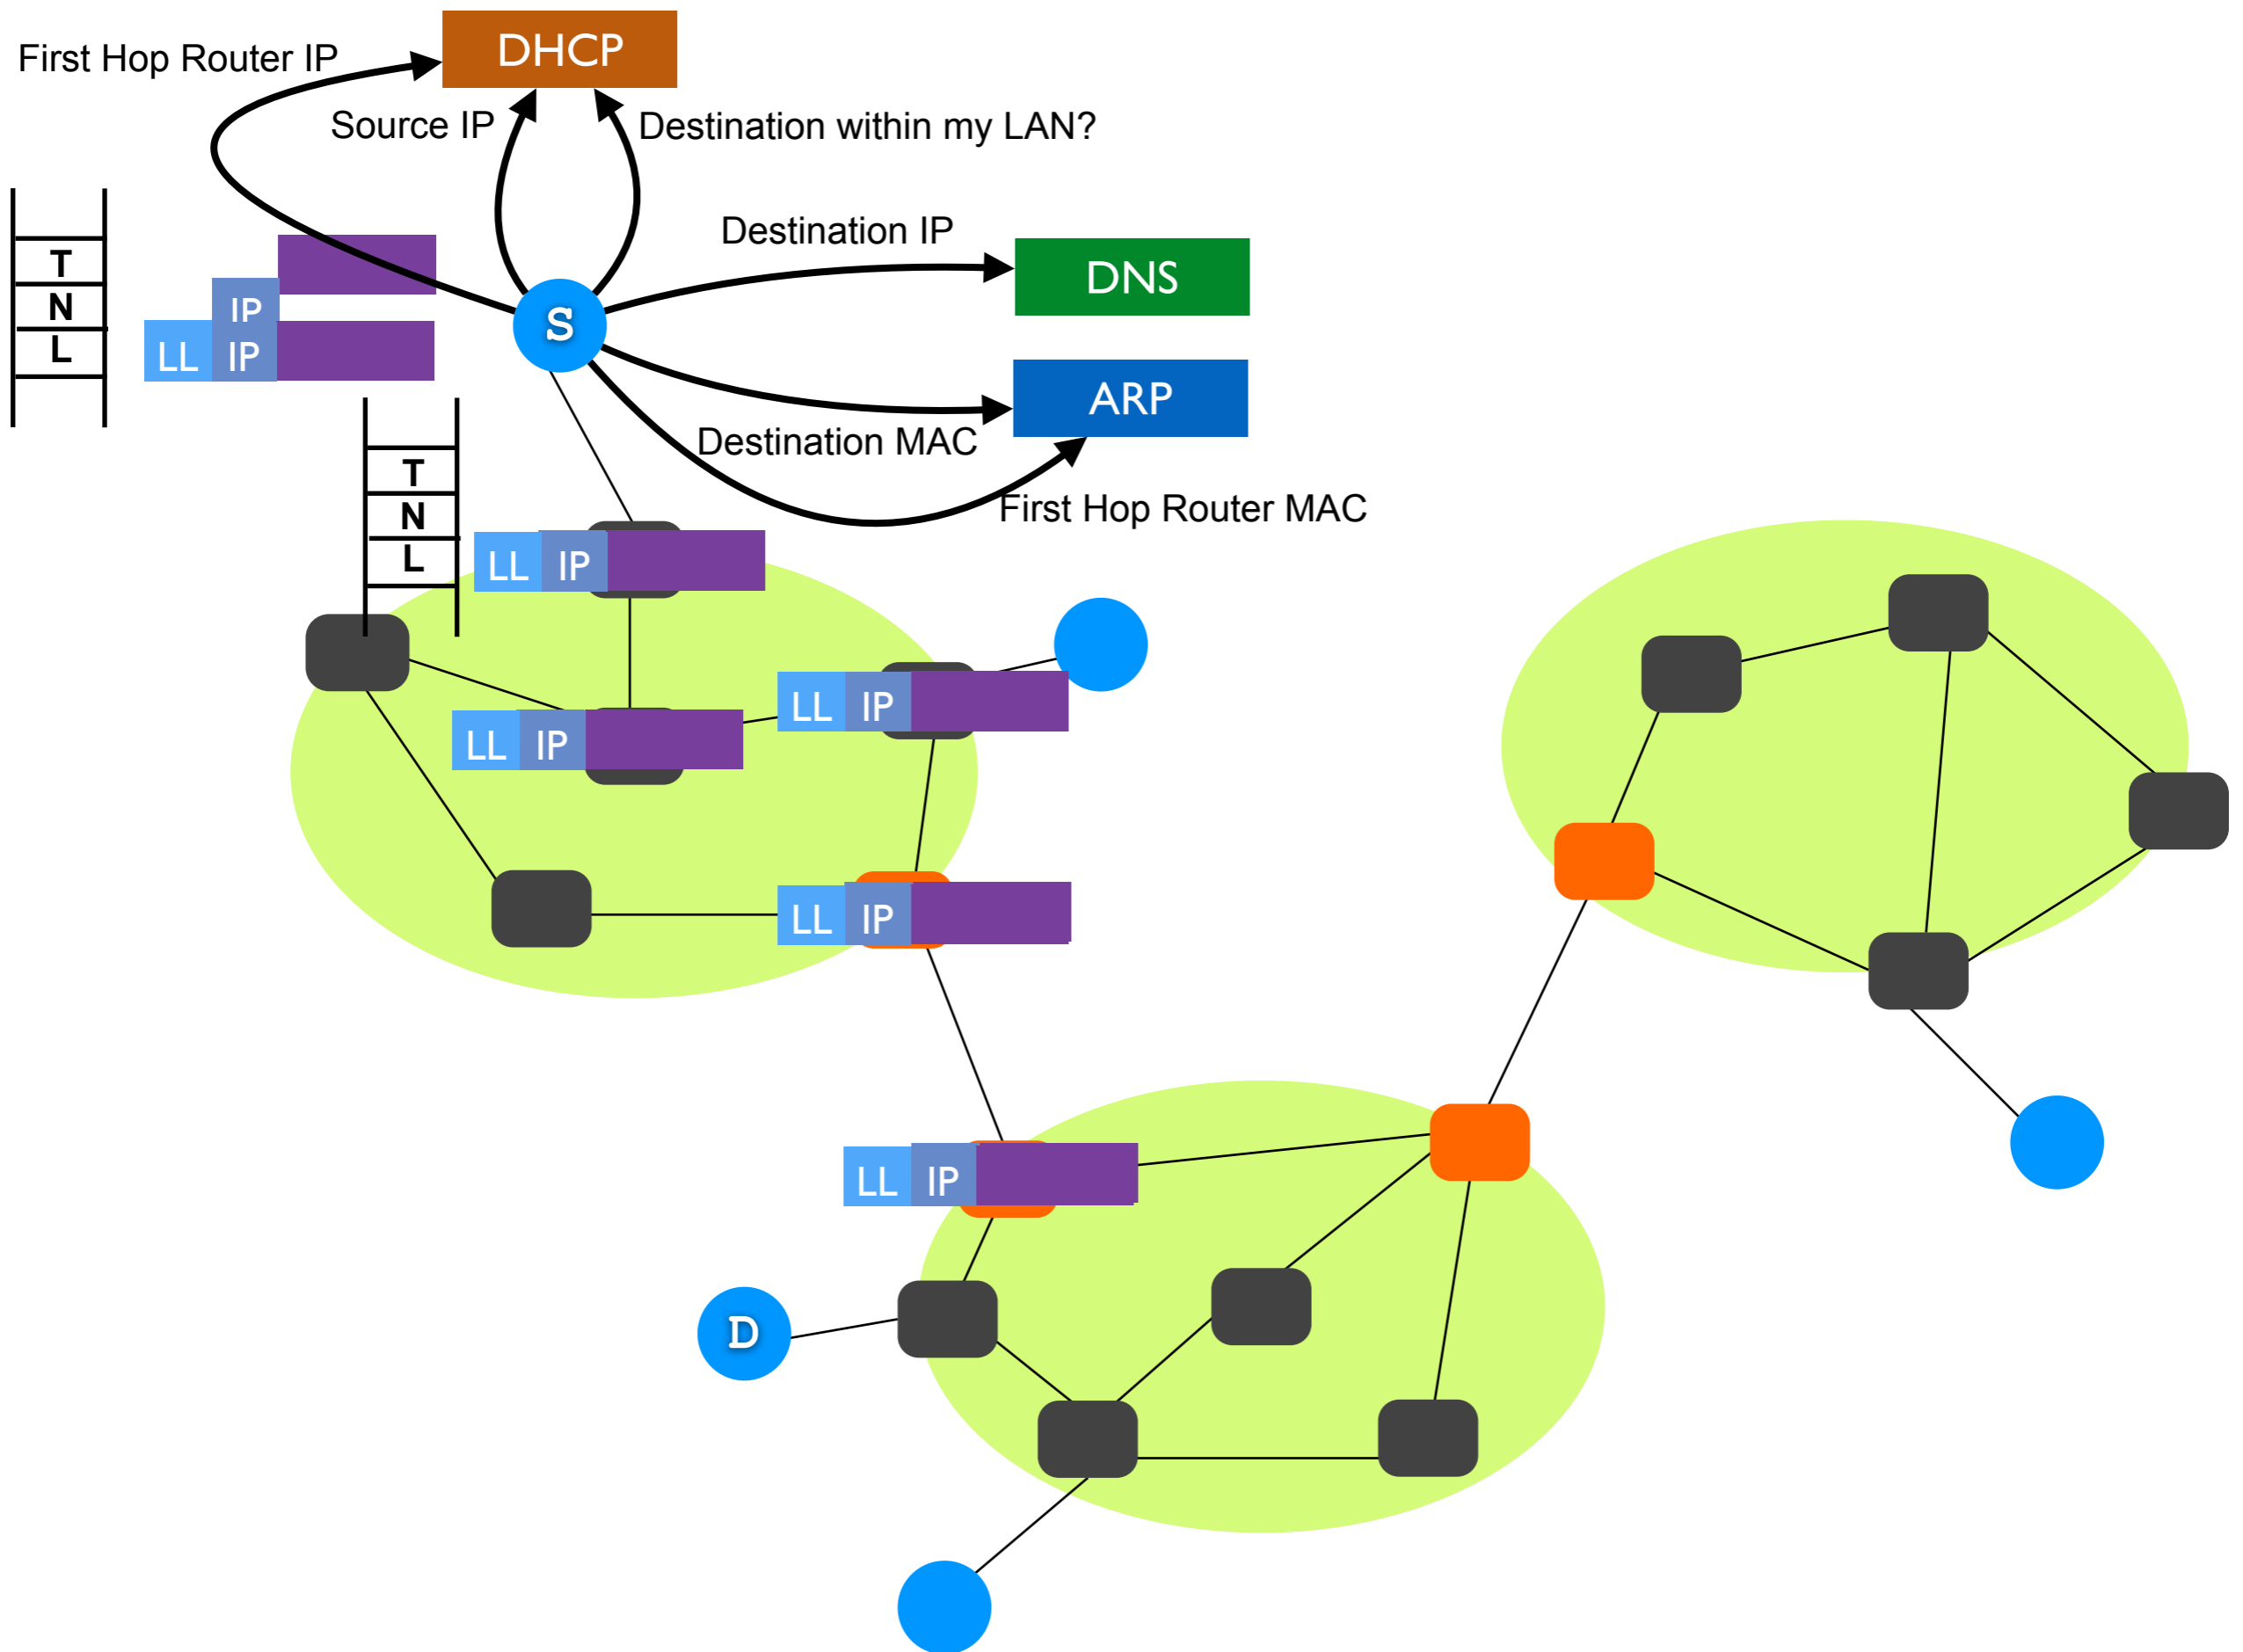
End-to-End I



How does Internet work — end-to-end?

- Network stack receives the packet from the application (roughly speaking)
- What is my IP address? (using DHCP)
- What is the destination IP address? (using DNS)
- Is destination IP address within my LAN? (using DHCP)
- **If destination IP address remote:**
 - **What is my next-hop router IP address? (using DHCP)**
 - **What is my next-hop router MAC address? (using ARP)**
 - Convert packet into frames with correct source/destination address
 - Convert frames into bits
 - Forward the bits to the wire ...
- **Each router**

End-to-End II



How does Internet work — end-to-end?

A router upon receiving a packet (implicit questions)

- **Is the destination in a LAN connected to me?**
 - Forward the packet to the destination
 - Using STP/CSMA/CD
- **Is the destination not in my LAN but in my ISP?**
 - Forward the packet to the next-hop router towards the destination
 - Using routing table entries via distance-vector routing algorithm
- **Is the destination in a different ISP?**
 - Forward the packet to the next-hop router towards the destination
 - Using routing table entries via BGP routing algorithm

You now know how the Internet works!!!!

All that is remaining:

Reliability.