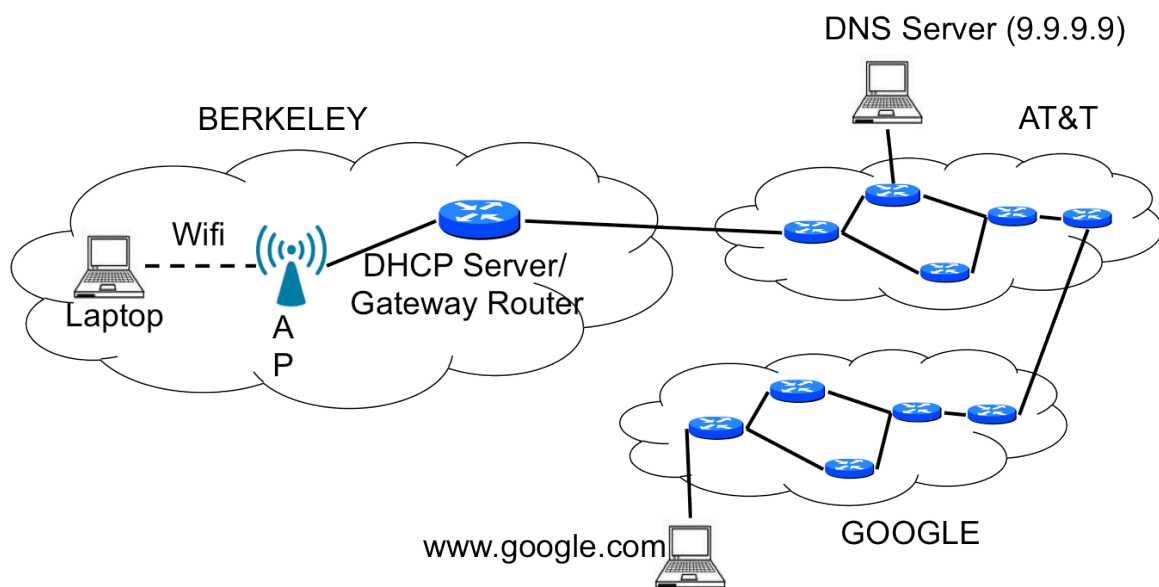


## CS168 Fall 2014 Discussion End-to-End and Switched Ethernet

### An end-to-end view

Consider the following set-up:

You have come to Soda Hall for CS168 office hours and would like to know the complete networking view of “what happens when you send a web page request”. But all the GSIs are busy with their own work, not bothering to answer your question and you decide to ask Google instead. You therefore boot-up your laptop, which is connected to an access point, through Wifi, which in turn is connected to an router serviced by AT&T as shown.



### Step 1: When you first connect to the network

When the laptop is booted up and first connects to the network, it needs an IP address.

**DHCP** is the **application** layer protocol used to obtain the IP address.

It uses **UDP** as its Transport Layer protocol with destination port **67**.

The resulting packet is encapsulated in an IP header with **0.0.0.0** as the source address and **255.255.255.255** as the destination address.

The IP datagram is encapsulated in a **802.11 / Wifi** (link layer) frame with your laptop's MAC address as the source and **FF:FF:FF:FF:FF:FF** destination MAC address.

How does the request reach the DHCP server? What does it contain?

**Broadcast the packet. DHCP server receives it. UDP packet demultiplexed to port 67. It contains the DHCP discovery message.**

What do you get back in response?

**DHCP Offer consisting of:**

The requested IP address for the laptop  
The address of the DNS server  
The address of the default gateway router  
The network mask

What are the last two steps of this 4-way protocol?

Client sends DHCP request for an offer.  
The corresponding DHCP server finally sends an ACK.

What changes if you are connected using Ethernet instead of Wifi?

- The basic protocol remains same up to network layer
- Wifi and Ethernet use the same LL header
- Preamble and CRC added by Physical layer varies
- Technology used to transmit the packets varies

### Step2: Initiating request to [www.google.com](http://www.google.com)

DNS is the **application** layer protocol used to resolve the host name to the IP address.

It uses **UDP** as its Transport Layer protocol with destination port **53**.

The resulting packet is encapsulated in an IP header with the laptop's IP address obtained from DHCP as the source address and **9.9.9.9** as the destination address. The IP frame is attached to a Wifi frame with destination MAC address of **the gateway router**.

Wait! How will the laptop know this MAC address?

**ARP**

The packet is then routed to the DNS server (we later discuss how).

Do we do a complete recursive DNS look-up every time we initiate a request for the web page?

**No, check cache first**

### Step3: Communicating with [www.google.com](http://www.google.com)

**HTTP** is the **application** layer protocol used to request a web page from the web server.

It uses **TCP** as its Transport Layer protocol with destination port **80**.

What does the first packet sent by your laptop to the destination contain?

**TCP SYN packet**

This packet is encapsulated in an IP header with your laptop's address as the source address and [www.google.com's](http://www.google.com) address obtained by DNS query as the destination address.

The IP frame is attached to a Wifi frame with destination MAC address of **the gateway router**.

How will the laptop know this MAC address?

**It is now cached. No need to do ARP again.**

List all the packet exchanges that take place between the source and the destination after this first step.

**TCP SYN (laptop -> google)**

**TCP SYN ACK (google -> laptop)**

**TCP ACK (laptop -> google)**

**HTTP GET Request (laptop -> google)**

**HTTP response (google -> laptop)**

Is the final response (containing the web page) necessarily transmitted as a single packet? If not how/why is it segmented? How is in-order delivery ensured?

**No. TCP may segment based on MTU. TCP ensures in-order delivery.**

Is the final response (containing the web page) necessarily transmitted as a single packet? If not how/why is it segmented? How is in-order delivery ensured?

**No. TCP may segment based on MTU. TCP ensures in-order delivery.**

### **The meta-step we have been ignoring so far: Routing the packets**

Which layer is responsible for packet delivery to the destination?

**Network Layer**

This layer provides different functionalities for DNS requests/response, TCP connection establishment packets and HTTP request/response packets (True/False).

**False**

The routes from an edge router in one domain to an edge router in another are computed using **BGP**. It is an **application** layer protocol.

The routes within a domain (to go from one edge-router to another or from Google's edge-router to the destination or from your laptop to the Berkeley edge-router) are computed using **DV/LS**.

A router runs the above algorithms every time it gets a new packet it is supposed to route (True/False)


**No. Done independently and stored in forwarding table**

When the packet moves from one hop to another, it has the same IP frame(which encapsulates the transport layer segment with the application data), but the Link Layer and Physical frame encapsulating the IP frame may change. (True/False)

**True. We need a different MAC address or may be different Link Layer technology to transmit data across different hops.**

### Switched Ethernet

Consider the following topology of Switched Ethernet switches. 0-4 are switches and a-f are hosts.



(a) List the root and edges of the spanning tree obtained by running the Spanning Tree protocol for Switched Ethernet.

**Root = 0. 0-2, 0-1, 1-3, 1-4. All the edges connected to hosts as well.**

(b) For this part only, suppose Switch 0 fails. What will be the new spanning tree?

**Root = 1. 1-3, 1-4, 3-2. Again, all the edges connected to hosts as well.**

(c) Suppose these switches are Learning Switches and the following sequence of packets are sent in order. For each packet, will it be flooded or unicasted? The switch tables are empty to start out.

1. b to a **flood**
2. e to b **unicast**
3. f to c **flood**
4. c to f **unicast**
5. d to c **flood at 4, unicast at 1**
6. a to e **unicast**
7. d to f **unicast**