Network Design and Management

Edo Biagioni

esb@hawaii.edu

Department of Information and Computer Sciences

ICS 351: Today's plan

- introductions
- course overview
- overview of the Internet and routing
- initial lab visit

Introductions

ICS 351 Course Overview: Contents

practical, hands-on knowledge and understanding of networks and network equipment:

- end-systems (hosts, computers)
- switches
- routers
- cabling
- configuration
- understanding of network parameters and design
- address assignment
- network monitoring
- protocol headers

ICS 351 Course Overview: Structure

instructor: Edo Biagioni

Office: POST 305F

Office Hours: Mon 5pm-6pm, Thu 9:30am-10:30m

esb@hawaii.edu, 808-956-3891

instructor will open and monitor lab

each prelab must be completed individually **prior to** the scheduled lab time

the lab reports must be turned in the day after the lab

since the computers in the lab are not connected to the wider Internet, students must carry a personal USB drive and use it to save or transport data.

the course web page is at www2.hawaii.edu/~esb/2015fall.ics351/index.html

The Internet

- at one level, a collection of boxes and cables
- boxes are called: hosts, routers, switches, hubs, and many variations
- cables are called serial, ethernet, fiber, wifi, etc.
- logically, cables and boxes are grouped into networks
- the Internet has this name because it is designed to carry data across networks
- i.e. from one network to another

How the Internet works

- hosts originate data (routers may behave as hosts in this respect)
- data is grouped in limited-size packets (e.g. max size 576 or 1500 bytes)
- each packet's header carries a destination address and other information (e.g., what port, or application, this packet is for)
- the originating host sends the packet towards its destination
- each intermediate box must route the packet by selecting the next box

How the Internet works, part II

- the final destination may accept or discard the packet
  - many more details will be covered later
- the header with the destination address is the Internet Protocol (IP) header
- the header with the port number is the Transmission Control Protocol (TCP) header
- a box may get more packets than it can forward, in which case it will have to drop (i.e., not forward) some of the packets
  - congestion leads to packet loss

How the Internet works, part III

- the Domain Name system is also required for correct functioning of most of the Internet
  - but not for packet transport
- DNS is a distributed database
  - names are in many servers around the world
- Servers have many different functions
  - email, web, certificate authority, name server, game server...

How the Internet works, summary

- the most essential part of the Internet is packet forwarding
- each box forwards packets closer to their destination
- different clients, servers, and different protocols are what makes the Internet useful

- see this example

Routing a packet, part I

- a host usually has a single network to which it is connected
- that network is usually connected to the Internet by a single router
- so a host can send all its non-local traffic to this **default router**
- in contrast, a router is usually connected to more than one network
- so a router needs a mechanism to decide where to send each packet

Routing a packet, part II

- each router keeps a table of destinations to which each packet can be sent: a **routing table**
  - hosts do this too
- one of those table entries can be a default route to another router which knows more about the Internet: the **default router**
- the collection of routers which have specific routing table entries to every part of the Internet is known as the **default-free zone** or backbone
- every other router in the Internet has a **default route** and more specific routes

Routing protocols: overview

- maintaining the routing tables is very labor-intensive if done manually
  - so routing tables are maintained automatically:
  - each router knows what networks it is connected to
  - and communicates that to other routers
  - a routing protocol defines how routing information is communicated among routers
  - popular routing protocols are OSPF, BGP, RIP
- with the information from the routing protocol, each router can build and maintain its routing tables

Routing protocols: properties

- routing is not a perfect process: sometimes the routing tables are inconsistent, because it takes time for a router to discover changes and it takes time to communicate the new information
  - if a router doesn't have a route to a packet's destination, it will drop the packet, i.e., not forward it
- most of the time routing protocols maintain the routing tables correctly
  - much faster than with manual updates

Router hardware and software

- big, expensive Cisco routers
- inexpensive Linux boxes with multiple network interfaces
- a Linux general-purpose computer can:
  - use routing software to route packets
  - just as the expensive Cisco router can
  - but not as fast
  - and perhaps not on the same media
- an expensive router should have hardware acceleration for
  - looking up routes in a routing table, and
  - forwarding packets from one interface to another
- the software to run the routing protocol might be very similar on a generic box and on an expensive specialized router