

1: Introduction

Last Modified:
1/27/2003 1:25:26 PM

-1

Why study computer networks?

- They are engineering marvels!
 - Scalability, layered protocols, lots of subtleties- worthy of study
- They are all around us!
 - Understanding the nuts and bolts behind a technology you use every day is exciting
- They are changing the world!
 - Profound societal changes taking place in our lifetime

-2

How much do you know already?

- How much do you know about what happens when you use networking applications?
- What do you hope to learn in this class?

-3

Goals of this class

- Understand both basic computer networking concepts and their instantiation in the current Internet
- Question why the current Internet is the way it is
 - Appreciate good
 - Understand limitations and consider solutions
- Gain practical skills (network analysis, network programming)

-4

How do we begin to make sense of the Internet?

- Often like the blind men and the elephant?

It was six men of Indostan
To learning much inclined,
Who went to see the Elephant
(Though all of them were blind),
That each by observation
Might satisfy his mind

-5

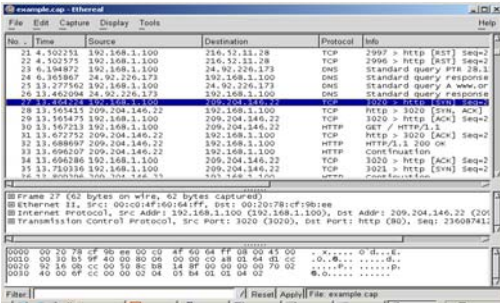
Network Trace Analysis?

- Network Trace Analysis - what happens on our local network when use a network application?

The First approached the Elephant,
And happening to fall
Against his broad and sturdy side,
At once began to bawl: "God bless me! but
the Elephant Is very like a wall!"

-6

Ethereal



Internet: Example

- Click -> get page
- page from local or remote computer
- link: <http://www.cnn.com>


specifies
 - protocol (http)
 - location
 (www.cnn.com)

- What network traffic results?
 - Request page
 - Send page
 - What else? Lots...



Internet: Locating Resource

- www.cnn.com is the name of a computer (and, implicitly, of a file in that computer)
- To find the address, the application uses a hierarchical directory service called the Domain Name System to translate human readable names to IP addresses




Internet: Port Numbers

- When a packet arrives at its destination, the operating system uses the destination port number to identify which application should receive it.
- This is called demultiplexing.

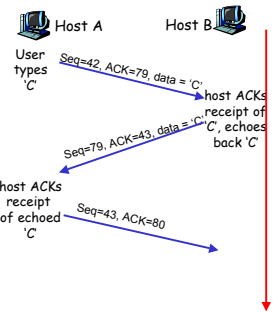
Internet: Connection

- The protocol (http) sets up a connection (another protocol, tcp) between the host and cnn.com to transfer the page
- The connection transfers the page as a byte stream, without errors: error control




Internet: End-to-end

- The end hosts worry about errors and pacing:
 - Destination sends ACKs
 - Source checks losses
- Stream of packets regulated and controlled by both ends:
 - retransmission of erroneous or missing bytes
 - Pacing, sender not overwhelming the receiver (flow control)
 - Pacing, Sender not overwhelming the network (congestion control)



Internet: Bits

- Data sent over a physical medium
- That equipment is not aware of the meaning of the bits



-13

Routing?


- Will network trace analysis let us understand the Internet?
- It gave us very little hint that between our machine and the remote machine there are many hops

The Second, feeling of the tusk,
Cried, "Ho! what have we here
So very round and smooth and sharp?
To me 'tis mighty clear
This wonder of an Elephant
Is very like a spear!"

-14

Routing

- Packets flows from end to end across many links
- Routers receive packet and try to forward one step closer to destination (**routing**)
- The packets contain all information necessary to indicate destination (**addressing**) and are handled individually
 - Example: Caravan of cars on a road trip
- Packets may take different paths



-15

Traceroute/tracert

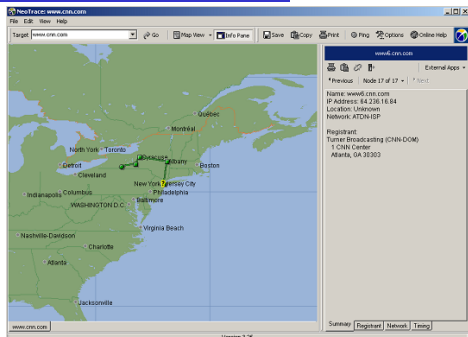
```

C:\>tracert www.cnn.com

Tracing route to cnn.com [64.236.16.84]
over a maximum of 30 hops:
  0  0 ms  0 ms  0 ms  csgate6.cc.cornell.edu [128.84.76.3]
  1  1 ms  1 ms  1 ms  cerbhus.cc.cornell.edu [128.84.76.4]
  2  1 ms  1 ms  1 ms  holl-8540-v1339.cit.cornell.edu [128.84.154.4]
  3  1 ms  1 ms  1 ms  core1-msfc-v17.cit.cornell.edu [128.253.222.174]
  4  1 ms  1 ms  1 ms  cornellnet4-dm1.cit.cornell.edu [128.253.222.4]
  5  2 ms  3 ms  2 ms  at-gsr2-syr-1-2.cornell-oc3.appliedtheory.net [169.130.253.5]
  6  6 ms  6 ms  6 ms  at-gsr1-alb-4-0-oc3.appliedtheory.net [169.130.253.41]
  7  9 ms  8 ms  9 ms  at-gsr3-nyc-3-0-oc12.appliedtheory.net [169.130.253.29]
  8  9 ms  9 ms  8 ms  pop1-nyc-p0-0.atdn.net [66.185.141.29]
  9  8 ms  7 ms  8 ms  hbl01-nyc-p0-0.atdn.net [66.185.141.18]
 10 14 ms 13 ms 14 ms  hbl2-vic-p0-0.atdn.net [66.185.152.201]
 11 13 ms 14 ms 15 ms  hbl1-vic-p11-0.atdn.net [66.185.152.206]
 12 55 ms 24 ms 22 ms  hbl1-cha-p7-0.atdn.net [66.185.152.28]
 13 27 ms 26 ms 30 ms  hbl1-atn-p0-0.atdn.net [66.185.152.182]
 14 23 ms 25 ms 27 ms  pop1-atl-p0-0.atdn.net [66.185.136.14]
 15 23 ms 25 ms 27 ms  pop1-atl-p0-0.atdn.net [66.185.136.14]
    
```

-16

Graphical Traceroute (plus DNS information 😊)



-17

Internet Topology?

- Traceroute gave us one slice through the Internet
- What does that tell us about the routes in the entire Internet?

The Third approached the animal,
And happening to take
The squirming trunk within his hands,
Thus boldly up and spake:
"I see," quoth he, "the Elephant
Is very like a snake!"

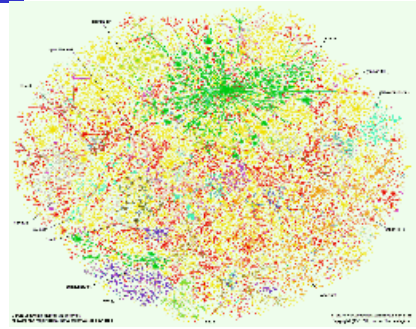
-18

Internet Map

- What does the Internet really look like?
 - That is actually a hard question to answer
 - Internet Atlas Project
 - <http://www.caida.org/projects/internetatlas/>
 - Techniques, software, and protocols for mapping the Internet, focusing on Internet topology, performance, workload, and routing data


-19

CAIDA: Layout showing Major ISPs



-20

CAIDA's skitter plot



-21


Internet growth and change

- Maps not even a complete "snapshot"
- Even if they were, what about change over time?
 - Growth trends (years)?
 - Cycles through months/weeks/days?
 - Failures? Stability?
- Summarizing the dynamic Internet

The Fourth reached out an eager hand,
 And felt about the knee.
 "What most this wondrous beast is like
 Is mighty plain," quoth he;
 "'Tis clear enough the Elephant
 Is very like a tree!"


-22

The Internet around 1990



-23


The Internet in 1997 (Post NSFNET)



-24

Internet Domain Survey Number of Internet Hosts

"Source: Internet Software Consortium (<http://www.isc.org/>)".



-25

Periodic Cycles

- Cycles of a day
 - People come to work in one part of the globe, go home in another
- Week
 - Weekends vs weekdays
- Holidays
- ...

-26

Stability?

- Despite any growth trends or periodic cycle, we also never have the Internet as a whole
- Machines disconnecting/connecting
- Stability of routes?
- Viruses? Attacks?

-27

Internet principles?

- Besides the details of protocols, topology and growth trends, are there fundamental "timeless" principles of the Internet?

*The Fifth, who chanced to touch the ear,
Said: "E'en the blindest man
Can tell what this resembles most;
Deny the fact who can
This marvel of an Elephant
Is very like a fan!*

-28

Packet Switching

- We've already seen this
- Packets indicate their destination
- No predetermined path for a packet to take
- Each intermediate node routes the packet closer to its destination

-29

Protocol

- If two entities are going to communicate, they must agree on the expected order and meaning of messages they exchange.
 - Hi ...Hi...Got the time?...two o'clock
SUCCESSFUL PROTOCOL EXCHANGE
 - Hi...Don't bother meXX ABORTED PROTOCOL
 - Allo...Hello..Quelle heuere a'tilXX<blank
stare> PROTOCOL MISMATCH

-30

Protocol

A protocol defines the format and the order of messages exchanged between communicating entities as well as the actions taken on the receipt or the transmission of a message.

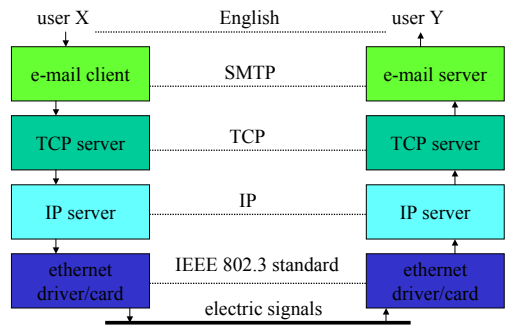
Layered Architectures

- Human beings are able to handle lots of complexity in their protocol processing.
 - Ambiguously defined protocols
 - Many protocols all at once
- How do computers manage complex protocol processing?
 - Specify well defined protocols to enact.
 - Decompose complicated jobs into layers that each have a well defined task

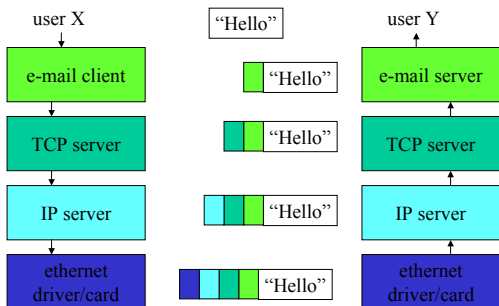
Layered Architectures

- Break-up design problem into smaller, more manageable problems.
- Modular design: easy to extend/modify.
- Difficult to implement (careful with interaction of layers for efficiency).

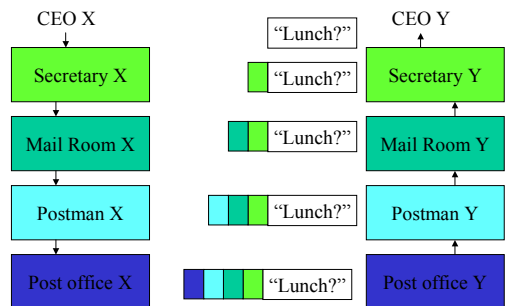
Protocol stack



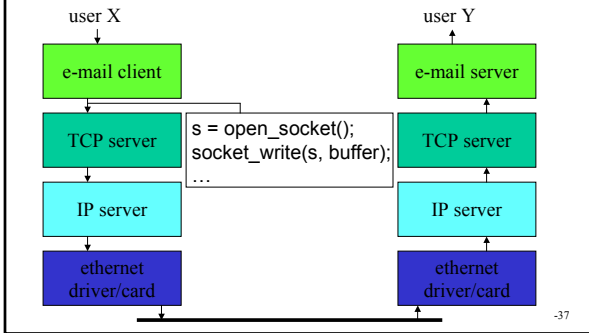
Protocol encapsulation



Physical Mail




Protocol interfaces



-37

Protocol stack: packet forwarding



-38

Internet future?

- If we could figure out what the Internet is, where is the Internet going?

*The Sixth no sooner had begun
About the beast to grope,
Than, seizing on the swinging tail
That fell within his scope,
"I see," quoth he, "the Elephant
Is very like a rope!"*

-39

How is Internet "governed"?

- "We reject kings, presidents, and voting. We believe in rough consensus and running code"
- Internet Engineering Task Force (IETF)
- Working groups
- Internet Engineering Steering Group (IESG)
- Internet Assigned Numbers Authority (IANA)

-40

Internet Standards Process

- Internet Draft - anyone
- RFC - at discretion of RFC Editor
- Internet Standard Maturity Levels
 - Proposed - IESG
 - Draft - 2 independent, implementations
 - Standard

-41

Other forces

- Proprietary technologies?
- Government funding and direction?

-42

Problems To Be Solved

- Security? Vulnerabilities to attack?
- Reliability? Enough for mission critical apps?
- Quality of service? Streaming media?
- Exhaustion of resources? IP address space? Spam mail?
- Social problems? Equal access? Privacy?

-43

Internet?

- Quick tour of the Internet
 - Introduced more questions that it answered?
 - Rest of semester is a detailed top-down tour of the Internet

And so these men of Indostan
 Disputed loud and long,
 Each in his own opinion
 Exceeding stiff and strong,
 Though each was partly in the right,
 And all were in the wrong!

-44

Outtakes

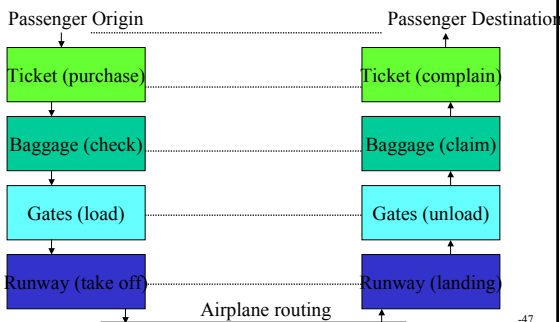
-45

Online Resources

- ISOC
- ACM
- CAIDA
- Internet History
- W3C

-46


Air travel



-47

A small Internet

Scenario:
A wants to send data to B.



-48

Distributed Control

- Where ever possible decompose the problem
- Examples:
 - No one central name to IP address data base - Domain Name System
 - No one global routing table - Hierarchical network of networks - handle routing within small autonomous systems
- Essential to Scalability

-49

Open System

- Basic Internet protocols are published as open standards
- Standards freely and readily available
- Ideal candidate for study

-50