4/9/02 Lecture, Part 2. Notes taken by Dean Campbell

Project Discussion:

- Implement a basic TCP sliding window protocol
- Do a tree lookup structure (something like a k-d tree?) asymmetrical

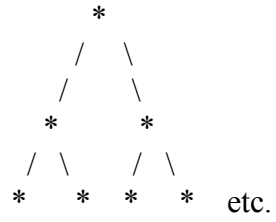

- Goal of project is to maximize the ratio of Learning to Work
- must provide specification of protocol to be designed
- any size group is OK
- target turning the project in 3 weeks after class discussion of relevant area
- project can be anything the Internet does
- if you have little Internet background, it may be easier and more structured to do something much like the existing Internet; if more experienced can go further afield
- it is not a requirement to use Fishnet
- code does not need to build on GCC under Linux, but needs to be buildable and executable. Could just send a binary executable along with source, or bring in and demo
- Can modify Fishnet if necessary or desireable
- 20 hours per project is about right level of effort
- Can choose 2 bigger projects instead of 3 "regular" size if particular area of interest
- Not reasonable to try to implement OSPF
- Two possible "real-world" applications that could be tried, but most such are difficult and too involved
    - Build a firewall
    - Mobile computing support (host level, make application mobile-aware)
- Send Fishnet bug reports to Tom, he will forward as appropriate

Routers

- How does a router know what to do with a packet?
    - Tree lookup for other routers, or
    - Use MAC address for hosts in internal ethernet
    - ARP request asks host on internal Ethernet whose IP address is in the packet for delivery. Host answers
    - What prevents another host from spoofing someone else's address?
        - If 2 hosts advertise same IP address for an ARP, then the packet is sent to the "closer" node. Happens frequently.
        - Security issues: Can configure routers to refuse all but certain IP address range

Packet Switching vs. Circuit Switching

Packets may take alternate routes: (Packet A shown going through R2, Pkt B thru R3)

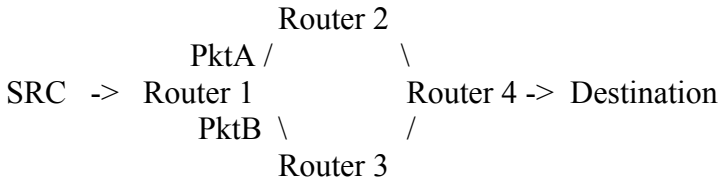

- Advertisement from destination works its way back to source
- Every packet sent from source has IP address of destination
- If we switch tables in R2 and R3 above, can make Packets A and B change paths
- What if mid-stream router is changed? Real time guarantees about packet switching times and routes are difficult to enforce
- What about establishing a "connection" packet? To establish a table entry or connection makes a "virtual circuit ID" so don't have to do switching/lookup on each packet
    - ATM does this, so does MPLS (multi protocol label switching)
    - Sets up a table entry for each router, appends table ID to each transmitted packet
    - Virtual circuit is established for each router path, so # of table entries is proportional to # of active circuit connections in router
    - Internet DOESN'T do it this way
    - MPLS widely used inside ISPs (like Global Crossing) to connect between "edge" routers - tunnels within ISP network

Fragementation

- What would we lose if disallow fragmentation?
    - If we send too big a packet, router sends ICMP command indicating acceptable packet size (assuming "don't fragment" bit is set)
    - ICMP = Internet Control Message Protocol. Every error/informational message generated by Internet is an ICMP message, for example, "destination does not exist".
    - Only about 10% of all IP packets get frag'd, so seems like a lot of extra work for routers to fragment and reassemble
    - On errors, why not selectively re-send the dropped fragments? Because this would require ability to ACK fragment by fragment.
    - IP is supposed to hand complete TCP packets to the TCP layer. The TCP header is only in the first of (potentially) multiple IP packets if fragmented.