


TCP/IP Architecture

Brighten Godfrey
CS 538 September 3, 2013



Cerf and Kahn: TCP/IP

Cerf and Kahn: TCP/IP

today design decisions

thursday architectural principles

Interconnection challenges




Heterogeneity

- Different addressing, supported packet lengths, reliability mechanism, latency, status information, routing

Must let each network operate independently

Solution:



Gateways and IP



Gateways sit at interface between networks




Fig. 2. Three networks interconnected by two GATEWAYS.

...and speak an Internetworking protocol




Fig. 3. Internetwork packet format (fields not shown to scale).

IP packet fragmentation



Allow maximum packet size to evolve

Protocol mechanisms to split packets in-transit

- byte-level sequence numbers

Reassemble at end-hosts

- Why not gateways?

Unreliable datagrams



No need for reliability in underlying network

Greatly simplifies design

- Exception handling always adds complexity
- But in IP: Any problem? Just drop the packet
 - examples?

What's not a stated reason for datagrams?

- Statistical multiplexing



Routing unspecified—but constrained!

- Hierarchical (network, host) address
- Route computed within network, hop-by-hop
- 8 bits for network: “*This size seems sufficient for the foreseeable future.*”
- Later: 32 bits in three size classes (A,B,C), and then CIDR.

Many new routing/forwarding designs need to change this address format




Fig. 4. TCP address.



Associated with a process on a host

Identify endpoints of a connection (“association”)

Rejected design:

- connection at host level
- packet may include bytes for multiple processes

What’s the difference between a port and an address?

What we now call TCP



Window-based scheme

Provides reliability, ordering, flow control

- Even though you might want only some of these

What else does it do today?

- Congestion control
- Three-way handshake




Fig. 10. The window concept.

What we now call TCP



“ *It is our expectation that the host level retransmission mechanism ... will not be called upon very often in practice. Evidence already exists [ARPANET] that individual networks can be effectively constructed without this feature.* ”

Why did they write this? Is it true now?

No congestion control in this early version!

- TCP congestion control introduces losses intentionally

Project suggestions

see separate document

Announcements



Read over syllabus

Review due Thursday:

- End-to-end arguments in system design
(Saltzer et al, 1984)