CS644
Advanced Networks

Lecture 3

Internet Architecture

Andreas Terzis

Overview

- Layering
- End-to-End Arguments
- A Case Study: the Internet

What is Layering?

- A technique to organize a network system into a succession of logically distinct entities, such that the service provided by one entity is solely based on the service provided by the previous (lower level) entity

Why Layering?

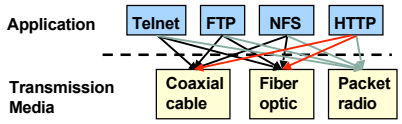

- No layering: each new application has to be re-implemented for every network technology!

Why Layering?

- Solution: introduce an intermediate layer that provides a unique abstraction for various network technologies

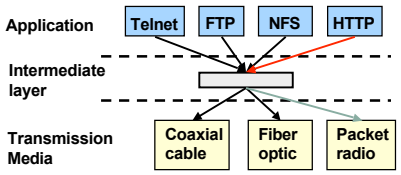

Layering

- Advantages
  - Modularity – protocols easier to manage and maintain
  - Abstract functionality –lower layer can be changed without affecting the upper layer
  - Reuse – upper layer can reuse the functionality provided by lower layer
- Disadvantages
  - Information hiding – inefficient implementations

ISO OSI Reference Model

- ISO – International Standard Organization
- OSI – Open System Interconnection
- Started to 1978; first standard 1979
  - ARPANET started in 1969; TCP/IP protocols ready by 1974
- Goal: a general open standard
  - Allow vendors to enter the market by using their own implementation and protocols

ISO OSI Reference Model

- Seven layers
  - Lower three layers are peer-to-peer
  - Next four layers are end-to-end

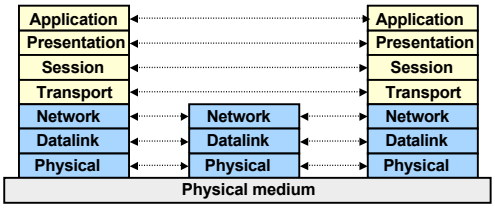

Data Transmission

- A layer can use only the service provided by the layer immediate below it
- Each layer may change and add a header to data packet

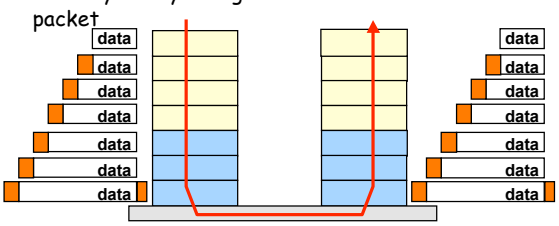

OSI Model Concepts

- Service – says what a layer does
- Interface – says how to access the service
- Protocol – says how is the service implemented
  - A set of rules and formats that govern the communication between two peers

Data and Control Planes

- Data plane: concerned with
  - Packet forwarding
  - Buffer management
  - Packet scheduling
- Control Plane: concerned with installing and maintaining state for data plane

Example: Routing

- Data plane: use Forwarding Table to forward packets
- Control plane: construct and maintain Forwarding Tables (e.g., Distance Vector, Link State protocols)

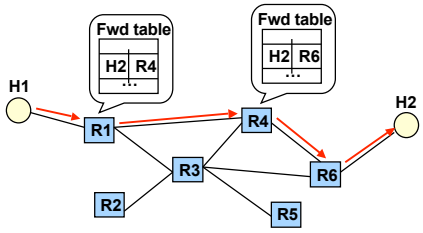

OSI vs. TCP/IP

- OSI: conceptually define: service, interface, protocol
- Internet: provide a successful implementation

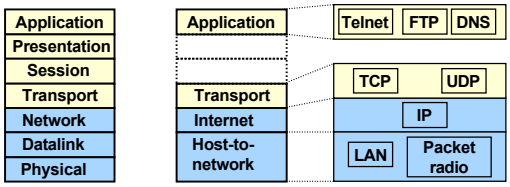

Key Design Decision

- How do you divide functionality across the layers?

Overview

- Layering
- ➢ End-to-End Arguments
- A Case Study: the Internet

End-to-End Argument

- Think twice before implementing a functionality that you believe that is useful to an application at a lower layer
- If the application can implement a functionality correctly, implement it at a lower layer only as a performance enhancement

Example: Reliable File Transfer

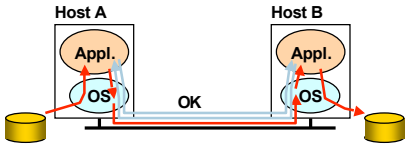

- Solution 1: make each step reliable, and then concatenate them
- Solution 2: end-to-end check and retry

Discussion

- Solution 1 not complete
  - What happens if the sender or/and receiver misbehave?
- The receiver has to do the check anyway!
- Thus, full functionality can be entirely implemented at application layer; no need for reliability from lower layers
- Is there any need to implement reliability at lower layers?

Discussion

- Yes, but only to improve performance
- Example:
  - Assume a high error rate on communication network
  - Then, a reliable communication service at datalink layer might help

Trade-offs

- Application has more information about the data and the semantic of the service it requires (e.g., can check only at the end of each data unit)
- A lower layer has more information about constraints in data transmission (e.g., packet size, error rate)
- Note: these trade-offs are a direct result of layering!

Rule of Thumb

- Implementing a functionality at a lower level should have minimum performance impact on the application that do not use the functionality

Other Examples

- Secure transmission of data
- Duplicate message suppression
- RISC vs. CISC

Overview

- Layering
- End-to-End Arguments
- A Case Study: the Internet

Goals

0. **Connect existing networks**
   - Initially ARPANET and ARPA packet radio network
1. Survivability
   - Ensure communication service even in the presence of network and router failures
2. Support multiple types of services
3. Must accommodate a variety of networks
4. Allow distributed management
5. Must be cost effective
6. Allow host attachment with a low level of effort
7. Allow resource accountability

Connect Existing Networks

- Existing networks: ARPANET and ARPA packet radio
- Decision: packet switching
  - Existing networks already were using this technology
- Packet switching → store and forward router architecture
- Internet: a packet switched communication network consisting of different networks connected by store-and-forward routers

Survivability

- Continue to operate even in the presence of network failures (e.g., link and router failures)
  - As long as the network is not partitioned, two endpoint should be able to communicate…moreover, any other failure (excepting network partition) should be transparent to endpoints
- Decision: maintain state only at end-points (fate-sharing)
  - Eliminate the problem of handling state inconsistency and performing state restoration when router fails
- Internet: stateless network architecture

Services

- At network layer provides one simple service: best effort datagram (packet) delivery
- Only one higher level service implemented at transport layer: reliable data delivery (TCP)
  - performance enhancement; used by a large variety of applications (Telnet, FTP, HTTP)
  - does not impact other applications (can use UDP)
- Everything else implemented at application level

Key Advantages

- The service can be implemented by a large variety of network technologies
- Does not require routers to maintain any fine grained state about traffic. Thus, network architecture is
  - Robust
  - Scalable

Other goals

- Distributed Management
  - Not there yet
- Cost effectiveness
  - Packet Headers: Do we really care?
  - E2E Retransmissions
- Accounting
  - No mechanisms for (fine-grain) packet counting mechanisms
    Flat-Rate pricing
  - Datagrams don't help

Other Points

- Architecture vs. Implementation
  - The architecture allows different implementations
  - Different implementations will provide different services
  - A lot of engineering required
- Diverse administration is a more difficult problem
  - No one has complete control of the network
  - Maybe not everyone folows the same rules
  - Different users have different goals