ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Information-Centric Networks

Section # 2.1: Internet Evolution

Instructor: George Xylomenos

Department: Informatics

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Funding

- These educational materials have been developed as part of the instructors educational tasks.
- The **"Athens University of Economics and Business Open Courses"** project only funded the reformatting of these educational materials.
- The project is being implemented as part of the Operational Program "Instruction and Lifelong Learning" and is co-financed by the European Union (European Social Fund) and national funds.

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF
2007-2013
programme for development
EUROPEAN SOCIAL FUND

Licencing

Week 2 / Paper 1

- The Design Philosophy of the DARPA Internet Protocols
  - David D. Clark
  - ACM CCR, Vol. 18, No. 4, August 1988
- Main point
  - Many papers describe how the Internet Protocols work
  - But why do they work as they do?
  - What was the motivation for the design choices made?
  - Essentially, what was their design philosophy

Introduction

- The DARPA Internet Protocols
  - DARPA (previously ARPA) funded development of the Internet
  - The paper captures 15 years of experience
  - Robert Kahn and Vinton Cerf designed TCP/IP in 1973
    - DARPA eventually made TCP/IP mandatory
  - David Clark assumed responsibility in 1981
  - Van Jacobson redesigned TCP congestion control in 1987
- Why are these protocols the way they are?
  - The design has evolved over time
    - The datagram was not always as prominent as it became
    - The Transport / Network split did not even exist
  - This paper outlines the original design principles
    - It also acknowledges that some of them were not met!

Fundamental Goal

- Multiplexed utilization of *existing* interconnected networks
  - Originally, the ARPANET and the ARPA packet radio network
  - Local area networks did not even exist!
  - But it was always assumed that other networks would show up
  - The alternative would be a unified network
    - It may have led to better performance
    - But it was not deemed to be practical
    - Separate networks better reflected administrative boundaries
  - The multiplexing technique chosen was packet switching
    - It fit existing applications better than circuit switching
    - Existing networks were also packet switched
  - Networks would be interconnected by gateways
    - Essentially these were store and forward packet switches

Second level goals

- The interconnection technique should be effective
- Effective means, in order of importance
  1. Communication continues despite failures
  2. Multiple types of service must be supported
  3. A variety of networks must be accommodated
  4. Distributed resource management must be permitted
  5. The architecture must be cost effective
  6. Host attachment must require a low level of effort
  7. The resources used must be accountable
- The order is *very* important!
  – A military network places survival on top
  – A commercial network could place accountability on top
  – Cost effectiveness is below the multiple types of service

Survivability

- Survivability in the face of failure
  - Two entities must continue communicating despite failures
    - If there is a path between them, they should keep communicating
    - Synchronization is lost only if there is total partition
  - This implies that the network should avoid maintaining state
    - Conversation state must not be kept inside the network
      - E.g. packets transmitted or acknowledged, flow control data
    - Otherwise failures would require the applications to be reset
  - Therefore, conversation state is kept at the endpoints
    - Fate sharing: conversation state is lost only if an endpoint fails
    - Much easier to engineer than in-network state replication
  - Consequences of fate sharing
    - Packet switches are stateless (datagram model)
    - Endpoints are responsible for the transport layer

Types of Service

- Support of multiple types of service
  - Different requirements in speed, latency, reliability
- The traditional network service was a virtual circuit
  - Bi-directional reliable data delivery
  - Remote login: low delay
  - File transfer: high throughput
  - TCP attempts to provide both
- But TCP cannot handle everything
  - XNET: cross-Internet debugger
    - How can you expect reliable transmission in a buggy endpoint?
  - Real time delivery of speech (teleconferencing)
    - No point in retransmissions and they also introduce delays
  - The network should work well with other services

Types of Service

- The split between TCP and IP
  - Originally a single protocol existed, but TCP did not fit everything
  - IP kept the best effort datagram delivery service
  - UDP exported this service to higher layers
  - TCP added the virtual circuit service
- What do the underlying networks provide?
  - No assumptions are made, datagrams are enough
  - But the architecture does not exploit better services
  - It proved hard to support multiple types of service
  - Each underlying network worked best with a specific service
    - X.25 works better for reliable services
    - Ethernet works better for unreliable services

Varieties of networks

- The Internet supports all kinds of networks
  - Slow and fast, reliable and unreliable, wired and wireless
  - Many more have been added since the paper was written
- Minimal requirements for Internet support
  - Ability to transport reasonably sized datagrams (>100 bytes)
  - Reasonable (not perfect) reliability
  - Some suitable for of addressing nodes is needed
  - No need for anything else
    - Sequenced delivery
    - Broadcast or multicast
    - Packet priorities
    - Error reporting
  - Everything else is engineered at the transport layer

Other Goals

- The three top goals were met quite well
  - The rest were not met or engineered that well!
- Distributed management partially works
  - Routing is distributed and differs at each domain
  - Policy routing is tough to do (see later lectures on BGP)
- Cost effectiveness depends on the circumstances
  - A 40 byte TCP/IP header is overkill for remote login
  - End-to-end retransmissions are wasteful
- Host attachment is complicated
  - The endpoint needs to implement complex transport protocols
  - Bad transport implementations hurt the network
- Accountability is basically non existent!

Architecture and Implementation

- The Internet architecture is very flexible by design
  - Network performance depends on lot on the realization
    - What networks are connected, what protocols are implemented
- How can one engineer a specific performance goal?
  - What is the required bandwidth?
  - What are the reliability requirements?
  - The issue is not verifying correctness, but predicting performance
  - Simulations are usually not enough due to complexity
  - The operating system is very critical due to reliance on endpoints
  - Back of the envelope calculations are common!
- Performance goals have not been addressed well
  - Even though it was one of the main goals of the designers
  - Hard to specify network parameters for procurement contracts!

Datagrams

- Datagrams are the fundamental unit of transport
  - No need for connection state in the network
  - Can be used to build different services
  - Allow many different networks to be incorporated
  - The use of datagrams turned out to be very important
- Datagrams do not reflect the intended Internet usage
  - UDP does export datagram services to higher layers
  - But very few applications actually are content with it!
    - Simple query/response service (e.g. DNS)
  - UDP is nearly always used as a basis
    - Reliability or delay smoothing are commonly added
  - The datagram is a building block, not a service in itself

TCP

- How did TCP evolve to its present state?
  - Note that it has hardly changed since 1988!
  - TCP uses byte based flow control for many reasons
    - Insertion of control data in the stream (dropped, ad hoc solutions)
    - Fragmentation of packets (dropped, IP does that)
    - Grouping of many small transmissions (used)
    - More exact flow control (used)
    - In retrospect, packet based flow control would also be useful
  - TCP uses the PSH flag as an indicator
    - Originally an EOL flag was used to delimit records
    - But it was insufficient for packets with many records
    - A lot of different solutions were debated and tried
    - In the end EOL was dropped as it was not generalized enough

Retrospective

- The Internet works well for its top priorities
  - But these priorities do not always match user needs
    - This became even clearer in the 1990s
- The datagram is a good case
  - It made the network very flexible and popular
  - But it makes accounting and resource management complex
- Packet flows
  - Sequences of related datagrams
  - Ideally recognized by the network
  - But without requiring state that must be replicated
    - The network should keep working if that state was lost
    - This is called soft state
  - Lack of flow support shows how hard it is to change the Internet!

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

End of Section # 2.1

Course: Information-Centric Networks, **Section # 2.1: Internet Evolution**

Instructor: George Xylomenos, **Department:** Informatics

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND