


# Design of a Network Service Architecture

PI: Tilman Wolf

Department of Electrical and Computer Engineering  
University of Massachusetts Amherst

wolf@ecs.umass.edu



## Network Services

- Traditional Internet abstractions:
  - Protocols, applications, processing on end-system
  - Simple store-and-forwarding on routers
- Network service architecture abstractions:
  - Decomposition of communication into basic services
    - Examples: data encoding, reliability, security, scheduling, ...
  - Custom end-to-end communication assembled from services
  - Routers provide implement services in data path
- Goals:
  - End-system application can choose best communication setup
  - Packet processing functionality not limited to end-systems
  - Network can evolve by introducing new services
- Example Services:
  - Example 1: reliable communication
    - Information encoding → service service → service service → service service → information decoding
  - Example 2: web caching
    - Information encoding → reliability service → reliability service → reliability service → reliability service → information decoding
  - Example 3: content transcoding
    - Information encoding → reliability service → reliability service → reliability service → reliability service → information decoding
- API: service socket
  - Similar to UNIX socket
  - Specification of service requirements
    - Example 1:  $\{N1|TCP\_TX > *|ENC > *|DEC > N2|TCP\_RX\}$
    - Example 2:  $\{N1|TCP\_RX < *|TCP\_TX&CACHE&TCP\_RX < N2|TCP\_TX\}$
- Challenges:
  - What services should be provided?
  - How can services be located?
  - How can connections be mapped to service nodes?
  - How is state managed across nodes and connections?

Tilman Wolf, "Service-centric end-to-end abstractions in next-generation networks," in Proc. of Fifteenth IEEE International Conference on Computer Communications and Networks (ICCCN), Arlington, VA, Oct. 2006, pp. 79-86.

## Service Architecture

- Hierarchical inter-network and intra-network design
  - Scalable service management (similar to PNNI)
  - Match with administrative boundaries of Internet
- Design assumptions:
  - Fixed routes
  - Connection reservations using service sockets
- Connection setup example:
  - End-System → Service Controller → Service Node 1 → Service Node 2
  - Sequence: connection request, mapping, service setup, setup ack, connection ack, data transmission, service processing, data transmission, service processing
- Prototype design of service controller and service node:
  - Service Controller: Connection Management, Request Parser, Controller View of Local Network and Resources, Flow Manager, Control Interface to Service Nodes
  - Service Node: Control Interface to Local Service Controller, Resource Monitor, Flow Manager, Service Manager, Processing Service, Flow-independent Processing State, Packet Demultiplexer, Packet Scheduler, I/O System

Sivakumar Ganapathy and Tilman Wolf, "Design of a network service architecture," in Proc. of Sixteenth IEEE International Conference on Computer Communications and Networks (ICCCN), Honolulu, HI, Aug. 2007.

## Task Placement Problem

- Infrastructure maps services to nodes
  - Mapping considerations
    - Node capabilities
    - Link and processor resources
- Mapping problem appears on all levels
  - End-system, Router, End-system
  - Service Provisioning on Routers, Mapping Computation on Network Processors
- Mapping problem example:
  - Nodes: 1, 2, 3, 4, 5, 6
  - Services:  $s_1, s_2, s_3, s_4$
- Mapping algorithms:
  - Randomized placement
  - Layered graph
  - Service step search
- Evaluation of mapping algorithms:
  - Simulation on 300-node network
  - Qualitative comparison and run time:
 

Algorithm	link resource usage	drop rate	end-to-end delay	running time
Randomized placement	-	-	-	++
Layered graph	o	o	o	-
Service step search	-	-	-	..
  - Network resource usage:
 

Algorithm	300 nodes	3000 nodes
Randomized placement	3 ms	169 ms
Layered graph	30 ms	2100 ms
Service step search	70 ms	5480 ms
  - End-to-end delay of connections:
    - Graphs showing network usage in percent vs number of connections in network and end-to-end connection delay in ms.
- Project status:
  - Implementation plans for Emulab prototype
  - Next steps:
    - Distributed mapping algorithm
    - State management

Sivakumar Ganapathy, Xin Huang, and Tilman Wolf, "Evaluating Algorithms for Composable Service Placement in Computer Networks," under review.