

Network Provisioning Using A Clearing House Architecture


IWQoS 2000

Chen-Nee Chuah, Lakshminarayanan Subramanian,
Randy H. Katz and Anthony. D. Joseph

EECS Department
University of California, Berkeley

1

Problem Statement




- How to support QoS for real-time applications over IP-networks?

Chen-Nee Chuah, University of California, Berkeley

2

Present Scenario



- SLAs describe acceptable traffic volume/rate, and expected performance assurance
- In practice: SLAs are not precise

Chen-Nee Chuah, University of California, Berkeley

3

Why is it hard?

- Lack of performance assurance in current IP-networks
 - SLAs are vague
- Scalability

Missing Piece !!!

Limited understanding on control plane (policy framework)

How to regulate resource provisioning across multiple domains?

Chen-Nee Chuah, University of California, Berkeley

4


Outline

- Introduction
 - Related work
- Resource Provisioning via a Clearing House Architecture
- Trace Driven Simulation Study
- Summary

Chen-Nee Chuah, University of California, Berkeley

5

Int-Serv Approach



- End-to-end per-flow reservation
- End-to-end admission control
- Pros: deliver end-to-end QoS per flow
- Cons: requires per-flow state information => potential bottleneck

Chen-Nee Chuah, University of California, Berkeley

6

Diff-Serv Bandwidth Broker Approach

- Admission control at the Edge Router
- BBs negotiate SLAs for aggregate traffic
- Pros: preserve scalability
- Cons: trade-off end-to-end QoS

Chen-Nee Chuah, University of California, Berkeley

Outline

- Introduction
- Resource Provisioning via a Clearing House Architecture
- Trace Driven Simulation Study
- Summary

Chen-Nee Chuah, University of California, Berkeley

Clearing House Approach

- Coordinates resource provisioning across multiple domains to provide statistical QoS

Hierarchical structure

Aggregate Scheduling of Requests

Chen-Nee Chuah, University of California, Berkeley

Design Goals

- Deliver finer-grained statistical QoS
 - Set ups and re-negotiate resource allocations for aggregate calls
 - Preserve quality of existing calls via admission control
- Retain scalability of state-less approach
- Optimize network utilization
- Continuous path monitoring for dynamic routing
- Easily extended to support multicast, mobility and secure billing

Chen-Nee Chuah, University of California, Berkeley

Proposed Clearing House Architecture

- Introduce logical hierarchy
- Distributed database
 - reservation status, link utilization, network performance


Chen-Nee Chuah, University of California, Berkeley

Assumptions

- Networks are capable of supporting differentiated service levels
- Traffic statistics and network performance can be easily measured
 - Independent monitoring system or ISPs
- Control and data paths are separate

Chen-Nee Chuah, University of California, Berkeley


Clearing House Infrastructure



Local Clearing House (LCH)

- Maintains intra-domain resource allocations
- Service intra-domain reservation requests
Local admission control
- Acts as a monitoring agent for parent CH-node
Regular updates to parent CH-node

Predictive Reservations



- Measure aggregate traffic statistics
- Advance reservations based on bandwidth predictor
- Adapt reservations based on aggregate requests

Traffic Predictor

- Local Gaussian predictor for advance reservations

Online measurement of aggregate traffic


Measurement & update interval: T_{mea}

$$R_A = m + Q^{-1}(\rho_{loss}) \cdot \sigma$$

Allow local admission control

Compensate for the coarse granularity of notifications

Resource Reservation Strategies




- Separate reservation and call-setup
- Aggregation of reservation requests
- Hierarchical approach
- Suppress immediate notifications after dynamic reservation changes

Outline

- Introduction
- Resource Provisioning via a Clearing House Architecture
- Trace Based Simulation Study
- Summary

Simulation Study

- vBNS Backbone Network Topology (1999)

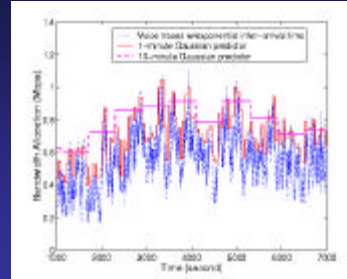


- Traffic Matrix
Weighted by population in each city

Simulation: Workload Model

- Two QoS classes
 - High priority voice calls
 - Best-effort data traffic
e.g. web, telnet, ftp
- Voice Traffic
 - Conference & telephone calls traces
 - Duration: 0.5 - 113 minutes
 - Exponential call inter-arrival time

Predictor Characteristics



1-min predictor
- 0.4 % Loss
- 7 % Over-Prov.

10-min predictor
- 0.7% Loss
- 33 % Over-Prov.

More BW for BE
traffic than
pre-partitioning
- avg. 286 Kbps
- max 857.2 Kbps

CH-node Characteristics

- RxW Aggregate Scheduling
 - Throughput = 3500 calls/s
(71% more than FIFO)
 - Call-blocking rate < 10%
 - Response time
 - negligible for small loads
 - increases piecewise linearly after 2000 calls/s
 - between 50-220 ms for most requests (>80%)
 - at 2000 calls/s

Outline

- Introduction
- Resource Provisioning via a Clearing House Architecture
- Trace Driven Simulation Study
- Summary

Summary

- A Clearing House Architecture
 - Coordinates resource reservations and admission control across multiple domains
 - Performs aggregate scheduling in a hierarchical manner
- Predictive Reservations
 - Shorten reservation setup time
 - Achieve loss rate < 1% without extensive over-provisioning

Advantages of CH Approach

- Maintains scalability
 - Core routers only maintain coarse-grained network state information
- Delivers statistical QoS assurance
 - Advance and dynamic reservations
 - Hierarchical coarse-grained admission control
- Reduces setup time
 - Advance reservations allow fast admission control decisions

Network Provisioning Using A Clearing House Architecture

C-N. Chuah, L. Subramanian,
R. H. Katz and A. D. Joseph
{chuah, lakme, randy, adj}@cs.berkeley.edu

URLs:

- <http://www.cs.berkeley.edu/~chuah/research/>
- <http://iceberg.cs.berkeley.edu/>

25