


Resource Provisioning Using A Clearing House Architecture

Chen-Nee Chuah, Lakshminarayanan Subramanian
 Randy H. Katz and Anthony D. Joseph
 {chuah, lakme, randy, adj}@cs.berkeley.edu
 EECS Department, U. C. Berkeley

ACM SIGCOMM 2000 - Poster Session
 August 31, 2000

Problem Statement



Future integrated network architecture: a single IP-based core for wireless, data and voice

Question: how to regulate resource allocation within and across multiple domains efficiently to achieve satisfactory QoS?

2

State of The Art

Present work

Int-Serv, Diff-Serv Bandwidth Brokers, Core Stateless Fair Queuing

Resource Allocation: Missing Pieces

Limited understanding on intra- and inter-domain resource control policy framework

Lack of algebra to define traffic regulation policy and service-level composition

Lack of a distributed resource control system

Research Challenges

Better end-to-end QoS assurance

Efficient resource utilization

Scalability and reduced state maintenance

Incremental deployment

Inter-operability with existing routing protocols

3

Our Approach: Clearing House Architecture

The Clearing House is a distributed architecture where each CH-node

Monitors traffic and network statistics continuously

Adapts aggregate reservations within and across multiple domains

Maintains intra- or inter-domain reservation status

Performs admission control

Components of CH architecture

Distributed resource controllers to allocate resources (tokens) fairly among a set of nodes


Admission control points at local CH-nodes close to host networks

Traffic predictors used to dynamically adapt intra- and inter-domain reservations

Aggregate scheduling mechanisms to increase throughput

4

CH: Flat + Hierarchical Structure




Local hierarchical structure within large domains
 Distribute network state information to various CH-nodes to reduce the amount of states maintained

Flat structure for peer-to-peer provisioning across multiple domains

5

CH: Distributed Resource Controller



Model as an economic bidding system to distribute resources

Control Points

Collectively emulate LCH-nodes

Service requests locally

Monitor traffic patterns & local resources

Send inter-domain resource requests to Resource Broker

Resource Broker

Partitions resource tokens in a fair manner

6

CH: Admission Control Modules

Trunks Reservation
Flows
ERs aggregate reservation requests

Reservation requests
 Admission control decisions

Pipe-model (edge-to-Edge)
 Advance reservations, R_A , are set up based on traffic predictors
 When a new request arrives,
 Check (peak rate + aggregate bandwidth usage for admitted flows)
 If less than R_A => admit
 Else, if less than a preset threshold
 => forward requests to upper CH-nodes, which performs admission control at different hierarchical levels: admit if successful, otherwise reject.
 Else => reject

7

Simulation Results of 2-Level Clearing House

Adaptive reservation using **Gaussian traffic predictor** achieves **loss rate < 1%**, with only **7% over-provisioning**.
RxW and Kinetic Tournaments scheduling increase CH throughput by 70% from 1800 requests/s (FIFO) to 3500 requests/s

8

Comparison to Diff-Serv Bandwidth Broker

	Bandwidth Broker	Clearing House
Architecture	Flat, pair-wise	Hierarchical + flat
Admission Control	Per flow at the edge	Nearest local CH-nodes
Reservation	Hose-model (pair-wise)	End-to-end aggregate resv.
State maintained	Per flow at the edge; Stateless at the core	Aggregate at the edge; Stateless at the core
Distributed resource control	None (one bandwidth broker per domain)	Deployed for large domains
Traffic regulator algebra	None	Present, with SLA composition & renegotiation

9

Strengths of the Clearing House Approach

- Maintains scalability**
 Core routers are stateless + edge routers maintain aggregate states
 Distributed resource controllers provision resources within large domains with reduced state information
- Provides statistical end-to-end QoS assurance**
 Adaptive aggregate reservation based on Gaussian predictor achieves loss rate < 1% without extensive over-provisioning
 CH-architecture performs edge-to-edge admission control
- Reduces setup time**
 Advance reservations allows fast admission control decisions
 RxW and Kinetic Tournament aggregate scheduling reduces mean setup-time and preserve fairness
- Facilitates interoperability with**
 Routing protocols: MPLS, OSPF
 Queuing mechanisms: CSFQ

10