Optimal Admission Control for a QoS-aware Service-oriented System

M. Abundo, V. Cardellini, F. Lo Presti

Talk Outline

- Background & Motivation
  - SOA- composite service orchestrated by a broker
- Broker Admission Control
- MDP Formulation
  - Infinite Horizon
  - Finite
  - Heuristic
- Numerical Examples
- Conclusions

Background & Motivation

- We consider a broker that offers a composite service with multiple QoS classes to several users each generating a flow of requests over time

Flow 1 – G-QoS class

Flow 2 – Y-QoS class

Flow 3 – R-QoS class

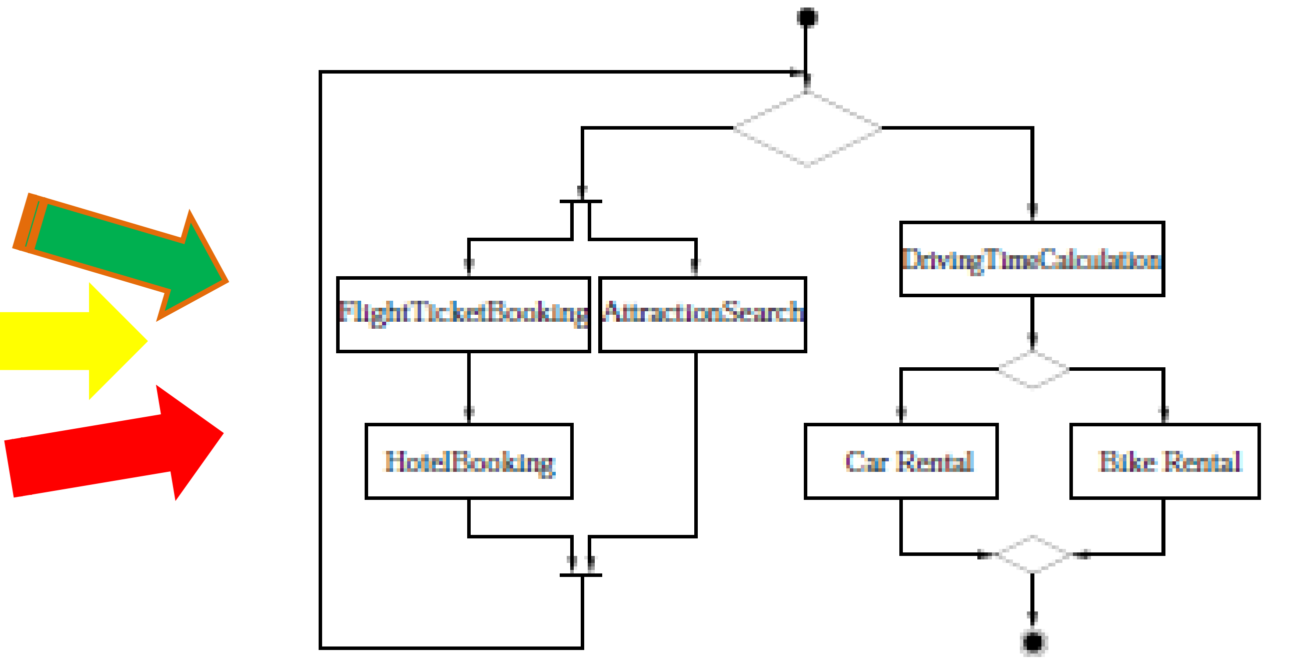

Broker Admission Control

- Broker
  - Finite amount of providers with own QoS and with limited capacity

⇒ It might not be possible to:

1. Serve all users requests; and/or
2. Serve them at their requested QoS

- Solutions
  1. Discover new providers/negotiate more capacity QoS with existing ones
  2. Renogotiate SLA with users
  3. Admission Control

Broker architecture overview: Moses

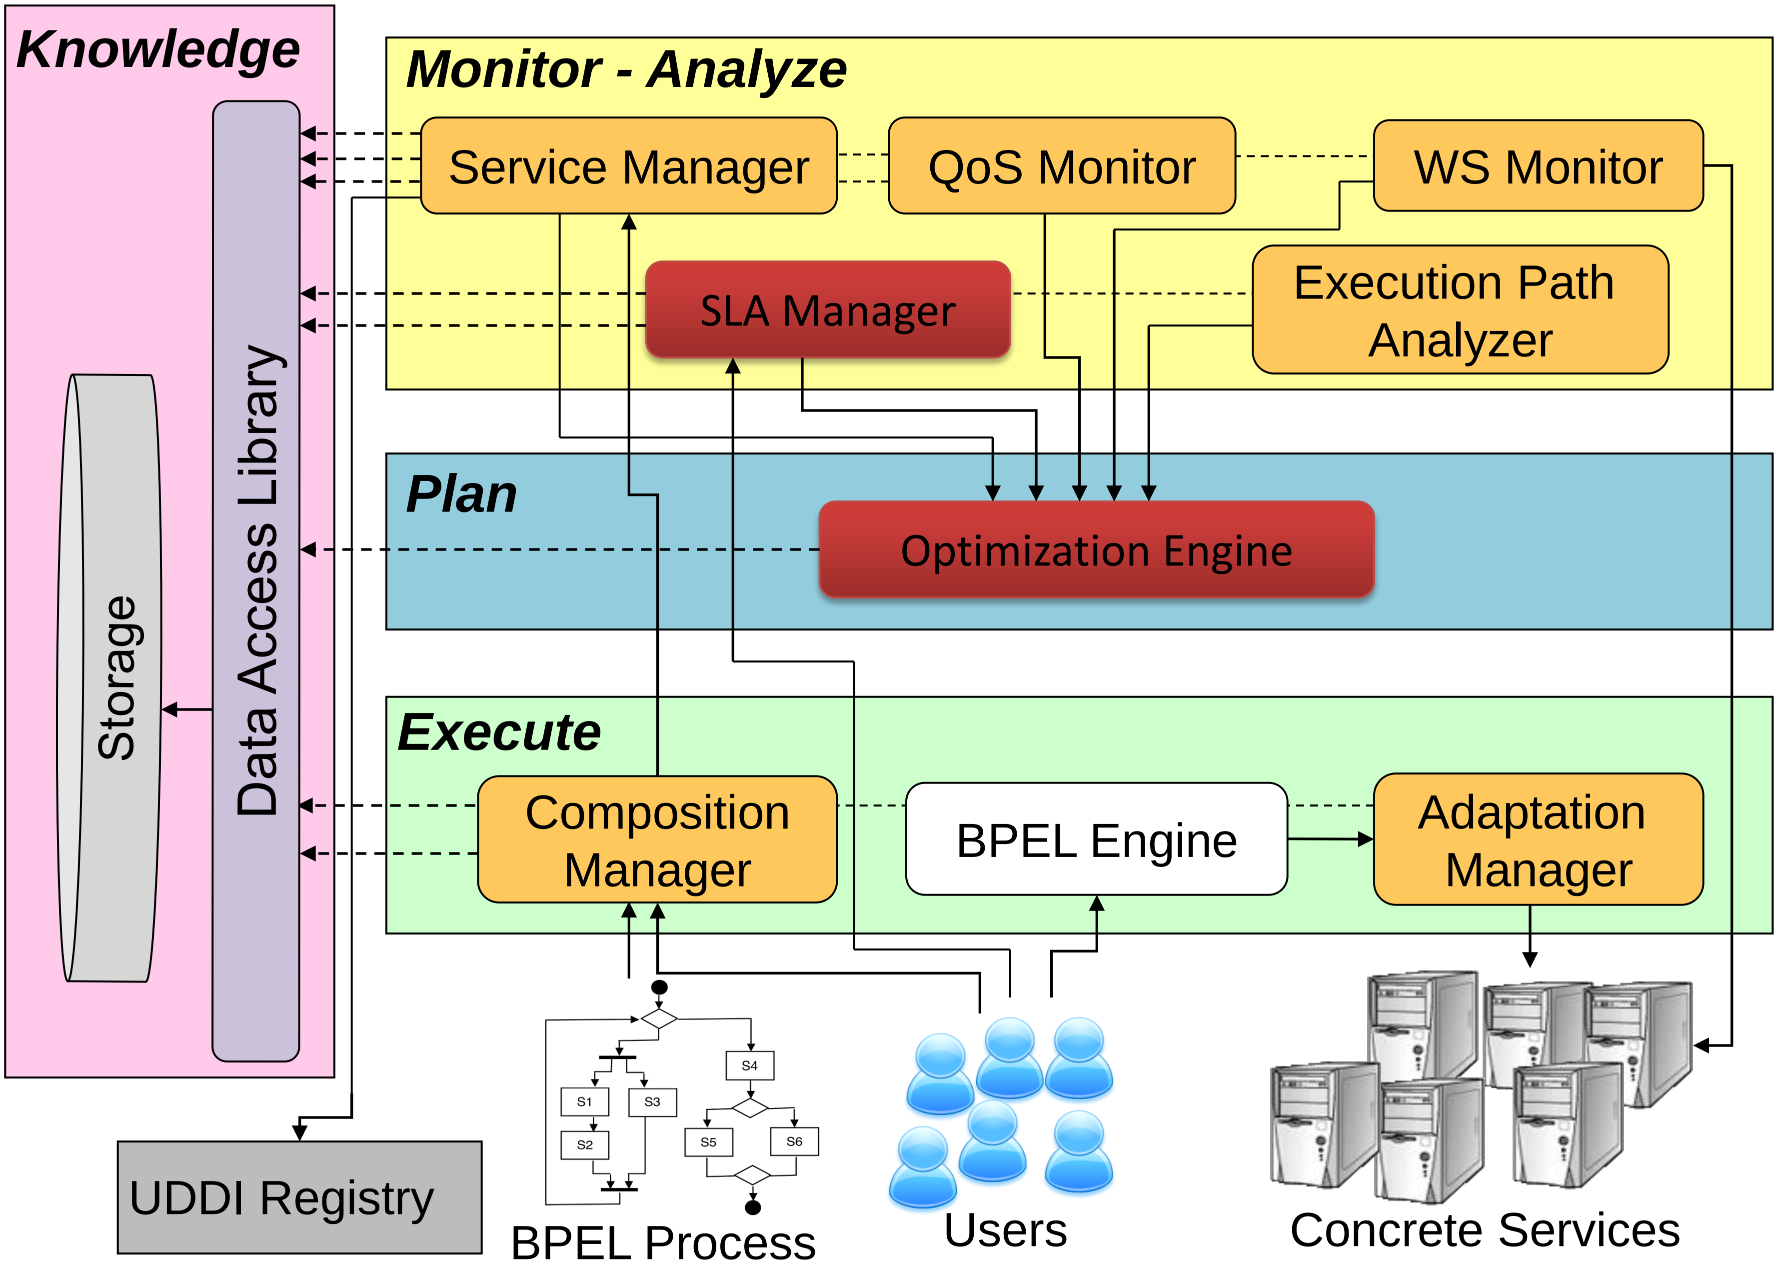

Broker Functionalities

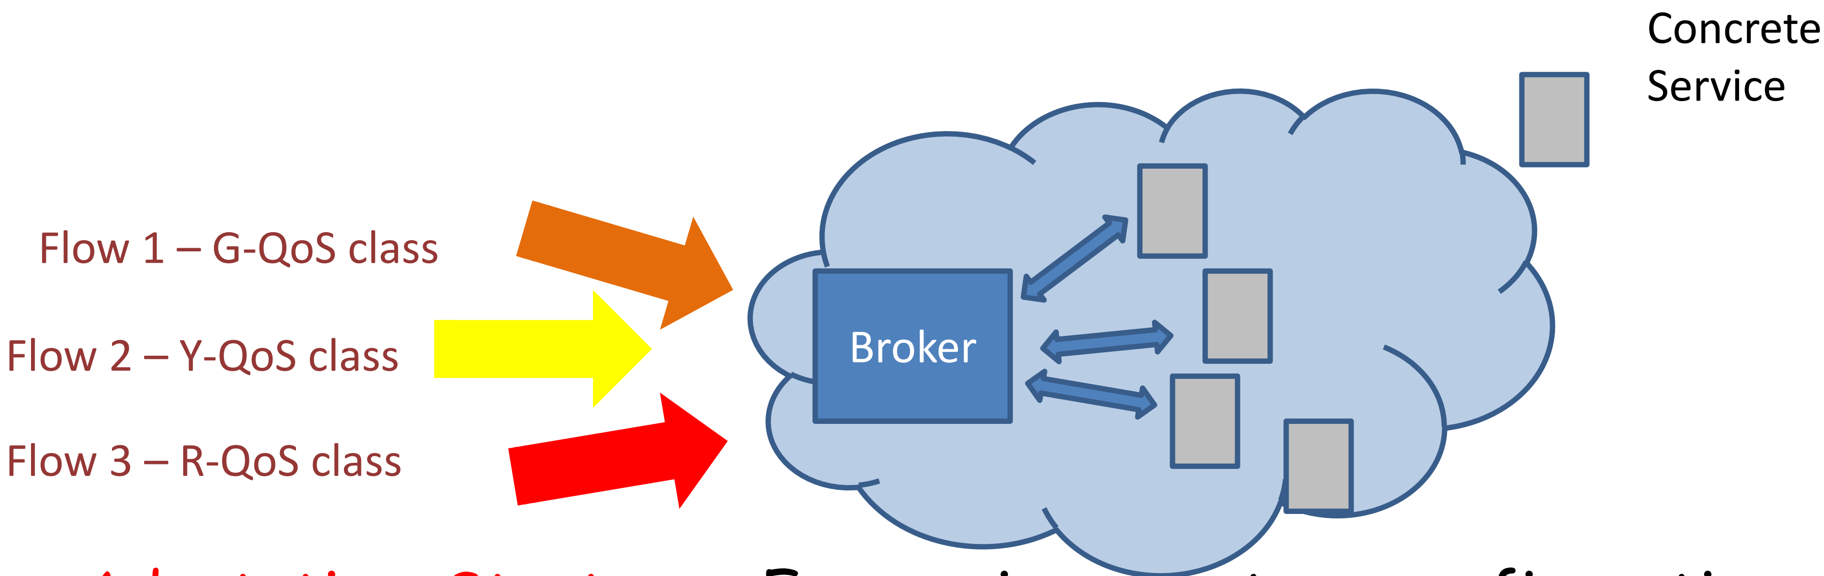

Adaptation Strategy: For a given system configuration, we determine for each service class the (set of) concrete services that implement each abstract service as to maximize the broker reward while guaranteeing user QoS using available concrete services

Admission Control: Upon a user arrival, decide whether to admit or refuse it as to optimize the expected infinite-horizon reward

Service Selection Problem

Given

- Set of classes of service with QoS parameters
- Workflow with n abstract services: S1, S2,…, Sn
- for each Si a set of concrete implementation $s_{i1},\ldots,s_{im}$, each characterized by its own QoS values

Determine for each Si and class of service k:

- Probabilities x^k_{ij} of assigning s_{ij} to a class k request which satisfy user QoS and maximizes the broker profit

Admission Control: MDP Model

- Assumption:
  - User arrive according to a Poisson Process
    - Characterized by:
      - service class k
      - contract duration
        - » exponential length
- Broker must decide whether to admit or refuse candidate users
  - Goal: Maximize broker expected infinite horizon discounted profit

$\Rightarrow$ Determine the Optimal Policy

MDP Model

$s=(n,\omega)$ System state [Bertsekas99]

- n =number of users for each service class/contract length
  - n is a matrix
- ω is the last event (user arrival/departure)

a broker action (admit/refuse)

$\beta(s,a)$ transition rate out of s under action a

$q_{ss'}$ transition rate from s to s'

MDP Optimal Policy

- Admission control policy π defines for each state s whether the broker should admit or refuse a new user
  - Function π: $\mathcal{S} \rightarrow \mathcal{A}$
    - $\mathcal{S}$ set of states
    - $\mathcal{A}$ set of actions

MDP: Optimal Policy

Broker expected discounted profit under policy π starting from stase s

– Infinite Horizon Discount Reward

$$v_\alpha^\pi(s) = E\left\{\sum_i \int_{\sigma_i}^{\sigma_{i+1}} e^{-\alpha u} c(s_i, a_i)\, du\right\}$$

α discount factor

σ_i i-th event (arrival/departure)

$c(s,a)$ broker profit (per unit of time) while in state s under action a

- Given by the solution of the service selection problem

Admission Control: Optimal Policy

The optimal policy π^* satisfies the optimality equation

$$\boxed{v_\alpha^{\pi^*}(s) = \sup_a \left\{ \frac{c(s,a)}{\alpha+\beta(s,a)} + \sum_{s'} \frac{q_{ss'}}{\alpha+\beta(s,a)} v_\alpha^{\pi^*}(s') \right\}}$$

Expected discounted reward til next event under a

Expected discounted reward from next event onward under the optimal policy

- Numerically computed
  - LP problem
  - Value Iteration

Admission Control: Issues

- Optimal Policy computational cost
  - State Space Explosion
    - as the number of classes, contract length, max level of class population grow
- Stationarity Assumption…
  - …policy re-computation!!

- Solutions
  - Optimize Finite Horizon reward
  - Heuristics
    - Sampling
    - Repeated Random Walk
  - Orders of magnitude faster

Experimental Results: Summary

- Broker offering 4 service classes
  - Class 1: Expensive, Stringent QoS requirements
  - …
  - Class 4: Cheap, Relaxed Qos requirements
- Compared the Optimal Policy to a Blind Policy
  - Blind Policy: Accept a new user as long as resources are available to support him and the already admitted users at the proper QoS levels

⇒Up to 100% reward improvement

- Class 4 users refused even when resources available (resources "reserved" for higher class users)
- Very good results (80% improvement) with the Finite Horizon and Heuristics

Conclusions & Future Works

- We studied and implemented an Admission Control for broker offering composite service
  - Formulated as MDP problem
  - On-line operation requires heuristics

- Future Works
  - Multiple Broker: Competitive Scenario