Dynamic Service Placement in Geographically Distributed Clouds

Qi Zhang[1] Quanyan Zhu[2] M. Faten Zhani[1] Raouf Boutaba[1]

[1]School of Computer Science
University of Waterloo

[2]Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

ICDCS 2012

Outline

Introduction

- Cloud computing has become a scalable and cost-efficient model for delivering large-scale services over the Internet
- Roles in a cloud computing environment
  - **Infrastructure Providers (InPs)** own data centers and lease resources for profit
  - **Service Providers (SPs)** rent resources from Infrastructure Providers to run their services
  - **End Users** are the customers of SPs
- Data centers are built in geographically distributed locations with heterogenous characteristics
  - different capacities, electricity costs and access delays from different locations

Introduction

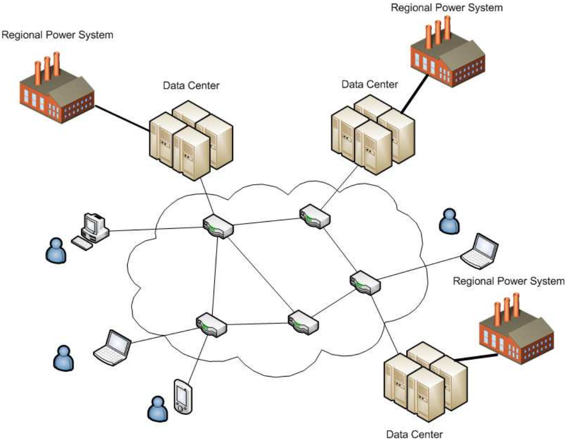

Figure 1 : Model of service placement in geo-distributed data centers

Motivation

- Every SP is facing a **Dynamic Service Placement Problem (DSPP)**:
  - Given multiple geo-distributed data centers, where should the service be placed to reduce operational cost while satisfying service level objectives (SLOs)?
- Design Challenges
  - Service demand are dynamic and originates from multiple locations (e.g. access networks)
  - Electricity prices are different from location to location and can fluctuate over time
  - There is a cost associated with reconfiguration
    - Setting up the server (e.g., VM image distribution)
    - Tearing down the server (e.g., data / state transfer)
    - Management overhead
- Limitations of existing work:
  - Early studies focus on static scenarios
  - Ignoring electricity cost and reconfguration cost

Our Contribution

- For a single SP scenario, we present an online algorithm for DSPP in geographically distributed clouds using Model Predictive Control (MPC)
- We analyze the case where multiple SPs compete for the capacity of each data center
  - Provide a formulation of the service placement game
  - Analyze the outcomes (i.e. Nash Equilibria) of the game in terms of price of stability (PoS) and price of anarchy (PoA)
  - Provide a coordination algorithm for achieving the optimal Nash Equilibrium (NE)
- Simulation results demonstrates the performance of our approach

Outline

System Architecture For A Single SP

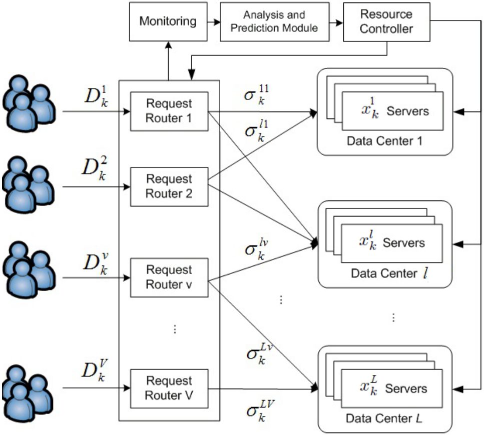

Figure 2 : System Architecture

Model and Assumptions

- Assumptions
  - All servers leased by each SP have identical size (CPU, memory, disk) and functionality (processing rate)
  - The number of servers in each data center takes continous values
- Defining $x_k^{lv} \in \mathbb{R}_+$ as the number of servers at location l serving certain amount of demand from $v \in V$
- Let $u_k^{lv} \in \mathbb{R}$ denote the change in the number of servers in x_k^{lv} at time k. Therefore, the state equation for x_k^{lv} is:

$$x_{k+1}^{lv} = x_k^{lv} + u_k^{lv}, \qquad \forall l \in L, v \in V, 0 \leq k \leq K. \qquad (1)$$

Problem Formulation

- The total resource cost H_k for service hosting at time k is

$$H_k = \sum_{l \in L} x_k^l p_k^l = \sum_{l \in L} \sum_{v \in V} x_k^{lv} p_k^l, \qquad \forall 0 \leq k \leq K \tag{2}$$

where $p_k^l \in \mathbb{R}_+$ is the resource cost in DC l at time k (e.g. electricity)

- The total reconfiguration cost G_k for service hosting at time k is

$$G_k = \sum_{l \in L} \sum_{v \in V} c^l (u_k^{lv})^2, \ \forall\ 0 \leq k \leq K. \tag{3}$$

we adopt a quadratic penalty function in this case, where $c^l \in \mathbb{R}_+$ is a known constant

- The objective of DSPP is to minimize the sum of costs over time

$$J = \sum_{k=1}^{K} \sum_{l \in L} \sum_{v \in V} H_k + G_k \qquad \forall 0 \leq k \leq K \tag{4}$$

Model Constraints

- Define σ_k^{lv} as the demand arrival rate from v assigned to data center l at time k,

$$\sum_{l \in L} \sigma_k^{lv} = D_k^v, \qquad \forall v \in V,\ l \in L,\ 0 \leq k \leq K. \tag{5}$$

- There is a data center capacity requirement C^l for each data center $l \in L$.

$$\sum_{v \in V} x_k^{lv} \leq C^l, \qquad \forall l \in L,\ 0 \leq k \leq K. \tag{6}$$

- The demand σ_k^{lv} arriving from location v is equally split among the local servers x_k^{lv}. The queueing delay for a demand at location $v \in V$ to a server at $l \in L$ can be computed as:

$$q(x_k^{lv}, \sigma_k^{lv}) = \frac{1}{\mu - \lambda} = \frac{1}{\mu - \sigma_k^{lv}/x_k^{lv}}. \tag{7}$$

Model Constraints

- For a request from location $v \in V$ to server at $l \in L$, we want to ensure that for any $(v,l) \in E$ with $\sigma_k^{lv} > 0$, the average delay is upper-bounded by $\bar{d}_{lv}$ (specified in the SLO):

$$d_{lv} + q(x_k^{lv}, \sigma_k^{lv}) \leq \bar{d}_{lv}, \quad \forall v \in V, l \in L, 0 \leq k \leq K. \tag{8}$$

By defining

$$a^{lv} = \begin{cases} \frac{1}{\mu - (\bar{d}_{lv} - d_{lv})^{-1}}, & \text{if } \bar{d}_{lv} - d_{lv} > 0, \\ \infty, & \text{otherwise }, \end{cases} \tag{9}$$

as a known constant for all $l \in L$, $v \in V$, we can rewrite the constraint (8) as:

$$x_k^{lv} \geq a^{lv} \sigma_k^{lv}, \qquad \forall v \in V, l \in L. \tag{10}$$

Problem Formulation

- DSPP can be formally represented as

$$
\begin{aligned}
\min_{\{\mathbf{u}_0,..,\mathbf{u}_{K-1}\}} \quad & J = \sum_{k=0}^{K} \mathbf{p}_k^\top \mathbf{x}_k + \mathbf{u}_k^\top \mathbf{R} \mathbf{u}_k \\
\text{s.t.} \quad & \mathbf{a}_k^\top \mathbf{x}_k \geq \mathbf{D}_k, && \forall 0 \leq k \leq K, \\
& \mathbf{s}^\top \mathbf{x}_k \leq \mathbf{C}, && \forall 0 \leq k \leq K, \\
& \mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{u}_k, && \forall 0 \leq k \leq K-1, \\
& \mathbf{x}_k \in \mathbb{R}_+^{LV}, \mathbf{u}_k \in \mathbb{R}^{LV}, && \forall 0 \leq k \leq K-1.
\end{aligned}
$$

Controller Design for DSPP

- DSPP is a convex optimization problem that can be solved optimally offline
- However, we need to solve the problem dynamically at run-time
- We adopt the Model Predictive Control (MPC) framework
  - At time k, predict the future demand $\mathbf{D}^v_{k+1}, ..., \mathbf{D}^v_{k+K}$ and electricity price $\mathbf{p}^l_{k+1}, .., \mathbf{p}^l_{k+K}$ over the horizon $[k+1, ...k+K]$ using methods such as ARIMA
  - Solve the DSPP to find $\mathbf{u}^{lv}_k, ..., \mathbf{u}^{lv}_{k+K-1}$
  - Perform the first action $\mathbf{u}^{lv}_k$

Outline

Competition Among Multiple SPs

- So far we have only studied the case for single SP
- We now analyze the case where multiple SPs compete for resources in multiple data centers
  - All SPs share the capacities of each data center
- Define a set of SPs $\mathcal{N} = \{1, 2, ..., N\}$ as the set of SPs who participate in the game
- Define $\mathbf{u}^i = \{\mathbf{u}^{i1}, \mathbf{u}^{i2}_1, ..., \mathbf{u}^{iV}_{K-1}\}$ as the strategy played by SP $i \in \mathcal{N}$, the goal is to optimize

$$J^i(\mathbf{u}^i, \mathbf{u}^{-i}) = \sum_{k=0}^{K} \sum_{v \in V} \mathbf{p}_k \mathbf{x}_k^{iv} + \mathbf{u}_k^{iv\top} \mathbf{R}^i \mathbf{u}_k^{iv}$$

Competition Among Multiple SPs

- The solution from the game approach is usually characterized by Nash equilibrium (NE), which in our context, is defined as:

$$J^i(\mathbf{u}^*) = \min_{\mathbf{u}^i \in \mathbb{R}^{LV}} J^i(\mathbf{u}^i, \mathbf{u}^{-i*}) \qquad \forall i \in \mathcal{N}$$

- The price of stability (PoS) ξ_{MPC}

$$\xi_{\text{MPC}} = \inf_{\mathbf{u}^* \in \mathscr{U}^*} \frac{\sum_{i \in \mathcal{N}} \sum_{v \in V} J_v^i(\mathbf{u}^{i*})}{\sum_{i \in \mathcal{N}} \sum_{v \in V} J_v^i(\mathbf{u}^{i\circ})}, \tag{11}$$

- The price of anarchy (PoA)

$$\rho_{\text{MPC}} := \sup_{\mathbf{u}^* \in \mathscr{U}^*} \frac{\sum_{i \in \mathcal{N}} \sum_{v \in V} J_v^i(\mathbf{u}^{i*})}{\sum_{i \in \mathcal{N}} \sum_{v \in V} J_v^i(\mathbf{u}^{i\circ})}, \tag{12}$$

- where $\{\mathbf{u}^{i\circ}, i \in \mathcal{N}\}$ is the optimal solution that minimizes the social welfare (i.e. the total cost over all SPs)

Analysis

Theorem

(PoS) Assume that the prediction horizon of each SP $i, i \in \mathcal{N}$, is the same, i.e., $W^i = \bar{W}$ Then, the price of stability ξ_{MPC} of the game Ξ is always equal to 1

Theorem

(PoA) The price of anarchy ρ_{MPC} of the game Ξ is unbounded

Mechanism for Achieving the Optimal NE

- The analysis suggests that it is necessary for the InP to participate in the game to improve the efficiency of the outcome
- We consider a case where an InP coordinates the resource allocation among SPs
  - Assume InP charge a fair return price, the InP attempts to improve social welfare as a measure of service quality
- We adopt an optimization decomposition approach for designing InP's control algorithm

Outline

Experiments: The Single SP Case

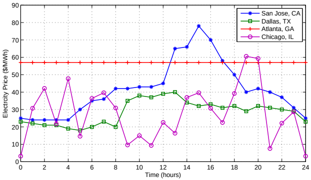

Figure 3 : Prices of electricity used in the experiments

Figure 4 : Impact of demand change on resource allocation

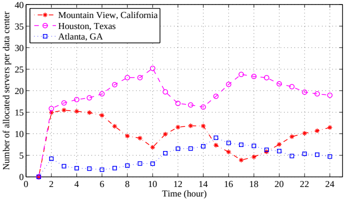

Figure 5 : Impact of price on resource allocation

Experiments: The Multiple SP Case

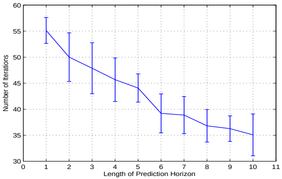

Figure 6 : Impact of prediction horizon length on the speed of convergence

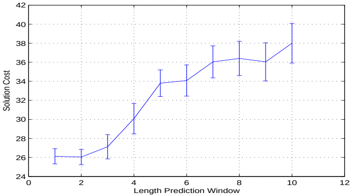

Figure 7 : Impact of prediction horizon length on the cost

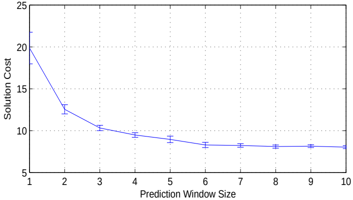

Figure 8 : Impact of prediction horizon with constant price and demand

Outline

Conclusion and Future Work

- We have presented a framework for the dynamic service placement problem using control theoretic models.
  - optimizes the desired objective dynamically over time according to both demand and resource price fluctuations.
- We have considered the case where multiple service providers compete for resource capacities in a dynamic manner
  - Analyzed the outcome of the competition and provided a mechanism for achieving the optimal NE
- Future work
  - More realistic experiments and accuracy of the prediction methods
  - Consider more realistic competition scenarios