UNIVERSITY OF NEBRASKA AT OMAHA

Heterogeneous resources cost-aware geo-distributed data analytics

Minmin Zhang

Department of Computer Science, University of Nebraska at Omaha, Omaha, NE 68182

Abstract

Many popular cloud service providers deploy tens of data centers (DCs) around the world to reduce user-perceived latency for better user experiences, in which a large amount of data is generated and stored in a geo-distributed manner. To collectively analyzing these data, Geo-distributed Data Analytics (GDA) has gained great popularity in meeting the growing demand to mine meaningful and timely knowledge from such highly dispersed data across scientific, commercial, and social domains. Many existing works invested significant effort to optimize data transfer strategies to efficiently use limited WAN by considering the network pricing policies on the base of infinite compute resources. However, the compute capacities and pricing policies, the limited and heterogeneous resources at different data centers, were ignored in most of the previous. To avoid both performance- and cost- bottlenecks, we propose a heterogeneous cloud resources cost-aware GDA system that exploits heterogeneous cloud resource capacities with a consideration of heterogeneous costs to meet cost-performance goals.

Research Problem

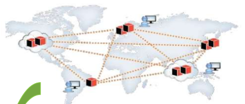

- Cloud service providers deploy datacenters (DCs) around the world
- User-oriented internet applications run their services on the geo-distributed DCs
- Geo-distributed Data Analytics (GDA) has gained great popularity for mining meaningful and timely knowledge from the dispersed data

The data transfer and compute cost are heterogeneous

Up to 7 times cost difference for data transfer

Nearly 2 times cost difference for computation on different DC locations and compute resource types (C4.4xlarge)

Region	compute cost ($/Hr)	network cost ($/GB)
US EAST (Virginia)	0.796	0.02
US WEST (California)	0.997	0.02
EU WEST (Ireland)	0.905	0.02
ASIA SE (Singapore)	0.924	0.09
ASIA SE (Sydney)	1.042	0.098
ASIA NE (Tokyo)	1.008	0.09
ASIA SOUTH (Mumbai)	0.8	0.086
SOUTH AMERICA (Sao Paulo)	1.235	0.138

Problem model

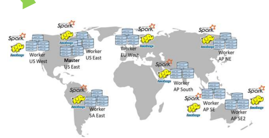

- 8 DCs located at different regions
- Each DC has diverse cloud resources and cost policies
- Cloud resources are heterogenous and may fluctuate due to resource contentions
- HiBench will be used to evaluate the system

Problem Model

Definition of variables

Variable	Definition
T_t	Time for data transfer across DCs
T_c	Time for computation
I	Total input data size
F_{xy}	The fraction of tasks assigned for DC y, but need to read data from DC x
NB_{xy}	The network bandwidth for transferring data from DC x to DC y
D	Set of datacenters (DCs)
I_x	Input data size at DC x
tc	Computation time for tasks
S_x	Computation cores at DC x
N_t	Total tasks of a job
NC_{xy}	The data transfer price from DC x to DC y
CC_x	The price for each computation slot per second

Example equation for Max cost

$$T_t \geqslant Max(\sum_y \frac{I \cdot F_{xy}}{NB_{xy}}, \sum_y \frac{I \cdot F_{yx}}{NB_{yx}}), \forall x, y \in D$$

$$T_c \geqslant t_c \cdot \frac{N_t \cdot (\frac{I_x}{I} - \sum_y F_{xy} + \sum_y F_{yx})}{Sx}, \forall x, y \in D$$

The process for getting min cost is similar and the tradeoff space between min and max cost can be chosen by users based on their budget.

Example Scenario

Initial settings for DCs

Parameter	DC 1	DC 2	DC 3
Input data size GB	20	30	50
Number of compute slots	40	10	20
Upload bandwidth GB/s	5	1	2
Download bandwidth GB/s	5	1	5

Data transfer for three different strategies

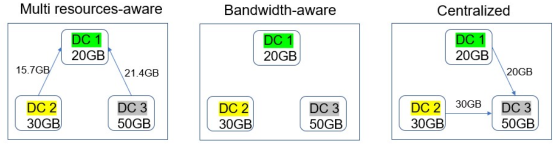

- Three task placement strategies for map stages are applied in the example.
- There are three DCs in the geo-distributed environment and the different initial settings are shown in Table.
- DC 2 has the computation and network bandwidth bottlenecks. DC1 has the least input data and largest compute capacity.

Example Scenario

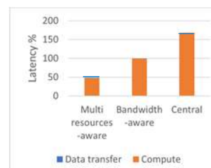

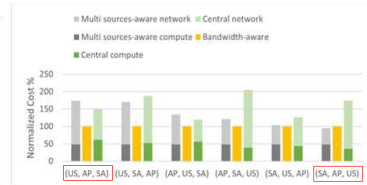

- To achieve the same performance, the cost may have to be doubled because of the heterogenous resources and cost policies.
- The compute and network resources and pricing policies are heterogeneous across the environment, AP and SA have more expensive data transfer and computation costs. Choose the last case can minimize the overall cost without affecting the performance.

Preliminary results

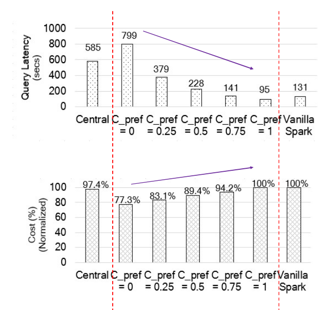

- Deployed Kimchi on CloudLab with 1 Spark master and 8 workers
- Applied TPC-DS as benchmark and 100GB data as input data
- The value of C_pref represents the desired network cost-performance tradeoff preference.
- Richer Cost-performance Tradeoff Space

Conclusion and Future Directions

- None of the current solutions have considered heterogeneous compute cost, which can lead to an overall cost bottleneck based on given workloads.
- Butler can determine optimal take placement based on given inputs and achieve best performance by avoiding cost bottleneck.
- More compute resources, e.g., serverless, have high performance could be applied in in future research.

Nebraska UNIVERSITY OF Omaha