Data Center Network Design for Internet-Related Services and Cloud Computing

Runyu Tang

with Yong Liang, Mengshi Lu, Zuo-Jun (Max) Shen

Oct. 14, 2021 @ XJTU

Background

Communication

Streaming Media

Online Shopping

Social Media

Background


Google Cloud
HUAWEI
华为云
Alibaba Cloud
aws
Microsoft Azure
腾讯云

Background

High spending:

CRN NEWS, ANALYSIS AND PERSPECTIVE FOR SOLUTION PROVIDERS AND TECHNOLOGY INTEGRATORS

Why Data Center Spending Will Hit $200B In 2021: Gartner

'Much of the reduced demand in 2020 is expected to return in 2021 when staff can physically be onsite,' says Naveen Mishra, Gartner's senior research director.

By Mark Haranas October 13, 2020, 11:44 AM EDT

$\sim$ US 1% GDP 2019, (Qatar 2019 GDP, 53rd worldwide.)

Research question:

How to design a data center network?

Background

Traditional facility location	• Fixed Cost • Transportation cost
Data center network design	• Power cost, Resource provisioning • Fabric latency cost, Endhost latency cost • Colocation, Interdependent demand, Non-linear power consumption, Network congestion

Data center network design is more challenging.

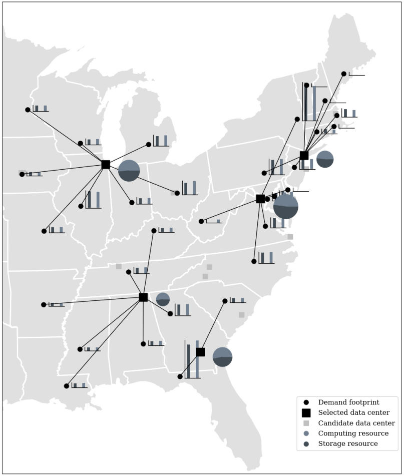

(a) Benchmark hierarchical model.

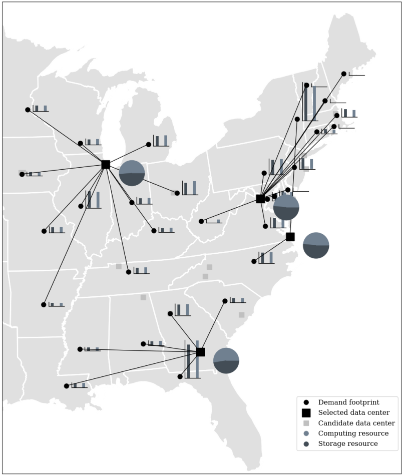

(b) Proposed integrated model.

The proposed integrated model saves more than half in endhost latency and cuts a quarter of the total cost.

Related Literature

- The importance of data center network design: [Iyoob et al., 2013, Greenberg et al., 2008, Larumbe and Sansò, 2012, Larumbe and Sansò, 2013]
- Facility Location: [Berman and Krass, 2015, Wang et al., 2004, Elhedhli, 2006, Aboolian et al., 2008, Castillo et al., 2009, Abouee-Mehrizi et al., 2011, Paraskevopoulos et al., 2016]
- MISOCP: [Baron et al., 2011, Atamtürk et al., 2012, Mak et al., 2013, Kong et al., 2013, Mak et al., 2015, Kong et al., 2017, He et al., 2017, He et al., 2018, Sen et al., 2017]

Main Takeaway

To the literature

- A novel data center network design model;
- Non-trivial MISOCP reformulation and two solution approaches;
  – Shed lights to similar integrated models.

To practise

- Structural properties
- Real-world data and design guidelines;

Table of Contents

Base model

Objective function

$$\min_{\mathbf{x},\mathbf{y},\mathbf{z}} \underbrace{\sum_{j\in\mathcal{J}} f_j x_j + \alpha \sum_{j\in\mathcal{J},k\in\mathcal{K}} c_j w_k z_{jk}}_{\text{fixed and power costs at data centers}} + \underbrace{\sum_{i\in\mathcal{I},j\in\mathcal{J}} d_i t_{ij} y_{ij} + \sum_{j\in\mathcal{J}} L_j(\mathbf{y}_{\cdot j}, \mathbf{z}_{j\cdot})}_{\text{fabric and endhost latency costs}}$$

Decision variables

- **x**: Location choice, integer variable;
- **y**: Demand assignment, integer variable;
- **z**: Resource provisioning, continuous variable.

Constraints

- Each footprint is assigned to one opened data center;
- Total required resource $\leq$ total resource provisioning;
- Total power provisioned $\leq$ power capacity;
- Maximum ratio between resource types.

Endhost Latency Cost

How a job is processed in DC?

- Processor sharing
- Tandem queue
- Service time for each stage: Coxian-2 distribution

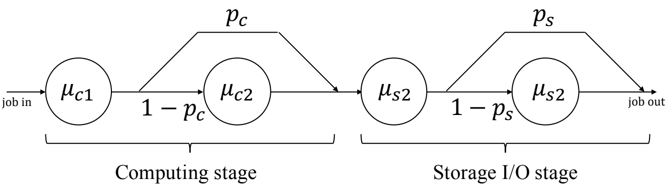

Figure: Sojourn time approximation with tandem queue and Coxian-2 distributed stage service times.

Endhost Latency Cost

Recall the **Objective function**

$$\min_{\mathbf{x},\mathbf{y},\mathbf{z}} \underbrace{\sum_{j\in\mathcal{J}} f_j x_j + \alpha \sum_{j\in\mathcal{J},k\in\mathcal{K}} c_j w_k z_{jk}}_{\text{fixed and power costs at data centers}} + \underbrace{\sum_{i\in\mathcal{I},j\in\mathcal{J}} d_i t_{ij} y_{ij} + \sum_{j\in\mathcal{J}} L_j(\mathbf{y}_{\cdot j}, \mathbf{z}_{j\cdot})}_{\text{fabric and endhost latency costs}}$$

Proposition

(Endhost Latency Cost)

$$L_j(\mathbf{y}_{\cdot j}, \mathbf{z}_{j\cdot}) = \sum_{k\in\mathcal{K}} \frac{\sum_{i\in\mathcal{I}} \tau_i d_i u_{ik} y_{ij}}{z_{jk} - \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij}}.$$

$$= \frac{\textit{unit latency cost} \times \textit{demand}}{\textit{provisioned resource} - \textit{total demand}}$$

The latency term is non-linear in decision variables $\mathbf{y}$, $\mathbf{z}$.

Extension to a queuing network

More generally:

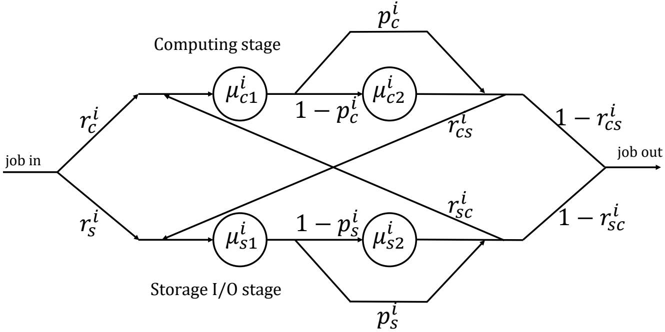

Figure: Sojourn time approximation with queuing network and Coxian-2 distributed stage service times.

Extensions — Colocation

Colocation:
A colocation (colo) is a data center facility in which a business can rent space for servers and other computing hardware to support **local** demand.

The corresponding model:

$$\textbf{(P-CO)} \quad \min_{\mathbf{x},\mathbf{y},\mathbf{z},\tilde{\mathbf{x}},\tilde{\mathbf{z}}} \quad \text{fixed and power cost} + \text{fabric and endhost latency cost}$$

$$+ \underbrace{\sum_{i\in\mathcal{I}} \tilde{f}_i \tilde{x}_i + \alpha \sum_{i\in\mathcal{I},k\in\mathcal{K}} \tilde{c}_i w_k \tilde{z}_{ik} + \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}} \frac{\tau_i d_i u_{ik} \tilde{x}_i}{\tilde{z}_{ik} - d_i u_{ik} \tilde{x}_i}}_{\text{costs associated with colocations}}$$

Extensions — Interdependent Footprints

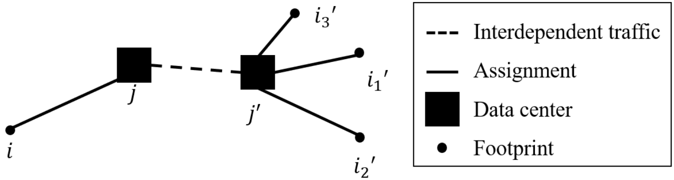

The arrival rate :

$$\lambda_{ij} = d_i y_{ij} + \sum_{i' \neq i} d_i P_{ii'} y_{i'j} \triangleq d_i \sum_{i'} P_{ii'} y_{i'j},$$

Note. $P_{ii'}$: the probability that a job from footprint i is routed to footprint i''s dedicated data center,

$P_{ii} \triangleq 1$, for all $i \in \mathcal{I}$.

Extensions — Interdependent Footprints

The corresponding model:

$$
\textbf{(P-ID)} \quad \min_{\mathbf{x},\mathbf{y},\mathbf{z}} \quad \text{fixed cost} + \text{power cost}
$$

$$
+ \begin{array}{l}\text{fabric latency cost between} \\ \text{footprints and data centers}\end{array}
$$

$$
+ \underbrace{\sum_{i,i' \in \mathcal{I}} \sum_{j,j' \in \mathcal{J}} \psi_i t_{jj'} d_i P_{ii'} y_{ij} y_{i'j'}}_{\text{fabric latency cost between data centers}}
$$

$$
+ \underbrace{\sum_{j \in \mathcal{J}, k \in \mathcal{K}} \frac{\sum_i \left(\sum_{i'} \tau_{i'} d_{i'} u_{i'k} P_{i'i}\right) y_{ij}}{z_{jk} - \sum_i \left(\sum_{i'} d_{i'} u_{i'k} P_{i'i}\right) y_{ij}}}_{\text{endhost latency cost}}
$$

Extensions — Convex Power Consumption

The power consumption is typically **increasing and convex in utilization** [Chen et al., 2013]. The high workload at one location requires extra power for cooling facility and results in worse power usage effectiveness (PUE).

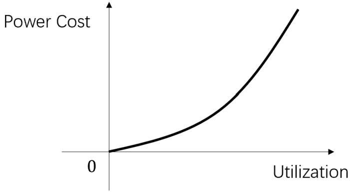

The corresponding model:

$$\textbf{(P-CP)} \quad \min_{\mathbf{x},\mathbf{y},\mathbf{z}} \text{ fixed cost} + \alpha \sum_{j\in\mathcal{J},k\in\mathcal{K}} c_j w'_{jk}$$
$$+ \text{ fabric and endhost latency cost}$$

Extensions — Network Congestion

Intense data exchanges may cause congestion in the network. The fabric latency cost i to j is given by:

$$d_i y_{ij} t_{ij} \left(1 + \left(\frac{d_i}{\chi_{ij}}\right)^{\sigma_{ij}}\right).$$

Extensions — Network Congestion + Convex Power Consumption + Interdependent Footprints

Putting all together:

(P-CC)

$$\min_{\mathbf{x},\mathbf{y},\mathbf{z}} \underbrace{\sum_{j\in\mathcal{J}} f_j x_j}_{\text{fixed cost}} + \alpha \underbrace{\sum_{j\in\mathcal{J},k\in\mathcal{K}} c_j w'_{jk}}_{\text{convex power cost}}$$

$$+ \underbrace{\sum_{i\in\mathcal{I},j\in\mathcal{J}} \left(d_i t_{ij} s_{ij} + \kappa_{ij}\chi_{ij}\right)}_{\substack{\text{congested fabric latency cost} \\ \text{between footprints and data centers}}} + \underbrace{\sum_{jj'\in\mathcal{I},j\in\mathcal{J}} \left(d_{jj'} s_{jj'} + \kappa_{jj'}\chi_{jj'}\right)}_{\substack{\text{congested fabric latency cost} \\ \text{between data centers}}}$$

$$+ \underbrace{\sum_{j\in\mathcal{J},k\in\mathcal{K}} \frac{\sum_i \left(\sum_{i'} \tau_{i'} d_{i'} u_{i'k} P_{i'i}\right) y_{ij}}{z_{jk} - \sum_i \left(\sum_{i'} d_{i'} u_{i'k} P_{i'i}\right) y_{ij}}}_{\text{endhost latency cost with interdependent footprints}}$$

Highly non-linear & non-convex, an MILP approximation may cause high waste.

Extensions — Network Congestion + Convex Power Consumption + Interdependent Footprints

Putting all together:

$$
\textbf{(P-CC)} \quad \min_{\mathbf{x},\mathbf{y},\mathbf{z}} \underbrace{\sum_{j\in\mathcal{J}} f_j x_j}_{\text{fixed cost}} + \alpha \underbrace{\sum_{j\in\mathcal{J},k\in\mathcal{K}} c_j w'_{jk}}_{\text{convex power cost}}
+ \underbrace{\sum_{i\in\mathcal{I},j\in\mathcal{J}} (d_i t_{ij} s_{ij} + \kappa_{ij}\chi_{ij})}_{\substack{\text{congested fabric latency cost}\\ \text{between footprints and data centers}}} + \underbrace{\sum_{jj'\in\mathcal{I},j\in\mathcal{J}} (d_{jj'} s_{jj'} + \kappa_{jj'}\chi_{jj'})}_{\substack{\text{congested fabric latency cost}\\ \text{between data centers}}}
+ \underbrace{\sum_{j\in\mathcal{J},k\in\mathcal{K}} \frac{\sum_i \left(\sum_{i'} \tau_{i'} d_{i'} u_{i'k} P_{i'i}\right) y_{ij}}{z_{jk} - \sum_i \left(\sum_{i'} d_{i'} u_{i'k} P_{i'i}\right) y_{ij}}}_{\text{endhost latency cost with interdependent footprints}}
$$

Highly non-linear & non-convex, an MILP approximation may cause high waste.
All of the models can be reformulated into **MISOCP**.

Reformulations

An example of reformulation — endhost latency cost:

$$\min_{\mathbf{y},\mathbf{z}} \quad \sum_{k\in\mathcal{K}} \frac{\sum_{i\in\mathcal{I}} \tau_i d_i u_{ik} y_{ij}}{z_{jk} - \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij}}$$

$$\Longleftrightarrow$$

$$\min_{\mathbf{y},\mathbf{z},\mathbf{v}} \quad \sum_{j\in\mathcal{J},k\in\mathcal{K}} v_{jk}$$

$$s.t. \quad \left\| \begin{pmatrix} 2\Lambda_k \mathbf{y}_{\cdot j} \\ v_{jk} - z_{jk} + \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij} \end{pmatrix} \right\|_2 \le v_{jk} + z_{jk} - \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij}, \ \forall j\in\mathcal{J}, k\in\mathcal{K}$$

$$v_{jk} \ge 0, \quad \forall j \in \mathcal{J}, k \in \mathcal{K}$$

Reformulations

An example of reformulation — endhost latency cost:

$$\min_{\mathbf{y},\mathbf{z}} \sum_{k\in\mathcal{K}} \frac{\sum_{i\in\mathcal{I}} \tau_i d_i u_{ik} y_{ij}}{z_{jk} - \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij}}$$

$$\Longleftrightarrow$$

$$\min_{\mathbf{y},\mathbf{z},\mathbf{v}} \sum_{j\in\mathcal{J},k\in\mathcal{K}} v_{jk}$$

$$s.t. \quad \left\| \begin{pmatrix} 2\Lambda_k \mathbf{y}_{\cdot j} \\ v_{jk} - z_{jk} + \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij} \end{pmatrix} \right\|_2 \le v_{jk} + z_{jk} - \sum_{i\in\mathcal{I}} d_i u_{ik} y_{ij}, \ \forall j\in\mathcal{J}, k\in\mathcal{K}$$

$$v_{jk} \ge 0, \quad \forall j \in \mathcal{J}, k \in \mathcal{K}$$

Still hard to solve, NP-hard $\rightarrow$ better solution approaches.

Table of Contents

Lagrangian Approach

- Relax the assignment fulfillment constraint $\rightarrow$ not useful;
- Relax the power capacity constraint.

Lagrangian Approach

- Relax the assignment fulfillment constraint $\rightarrow$ not useful;
- Relax the power capacity constraint.

$$\textbf{(P-L)}\ Z_D(\boldsymbol{\lambda}) = \min_{\mathbf{x},\mathbf{y},\mathbf{z},\mathbf{v}} \ \text{fixed and power cost} \\ + \text{fabric and endhost latency cost} \\ + \underbrace{\sum_j \lambda_j \Big(\sum_k w_k z_{jk} - p_j\Big)}_{\text{Capacity violation punishment}}$$

Convert the hard constraint $\rightarrow$ soft punishment.
Shadow price

Lagrangian Algorithm

Initialize $m = 0$, dual gap tolerance ϵ, Flag $=$ True.
Initialize the Lagrangian multipliers $\boldsymbol{\lambda}^{(m)}$
while Flag $=$ True **do**
  Update the lower bound by calculating the Lagrangian subproblem's objective value $Z_D(\lambda^{(m)})$.
  if The solution to **(P-L)** is feasible to the original problem **(P)** **then**
    Update the upper bound $\bar{Z}$
  end if
  if $|\bar{Z} - Z_D(\lambda^{(m)})|/Z_D(\lambda^{(m)}) < \epsilon$ **then**
    Flag $=$ False
  end if
  Update $\boldsymbol{\lambda}^{(m+1)}$ according to [Fisher, 2004],
  $m = m + 1$.
end while

Swapping heuristic for the subproblem

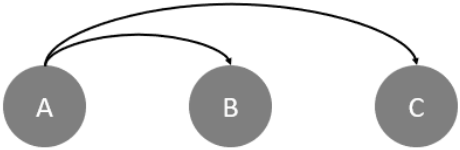

Proposition

(Bounds of the Incumbent Feasible Solution) Suppose we swap the host of footprint s *from the over-provisioned data center* t *to* t'*, which is guaranteed to be feasible after the swap. The cost increment imposed by the swap is bounded no matter* $t' \in \mathcal{B}$ *or* $t' \in \mathcal{C}$.

Lagrangian with *Extremal Extended Polymatroid Cuts*

Relax the power capacity and resource ratio constraints:

$$\min_{\mathbf{x},\mathbf{y},\mathbf{z},\mathbf{v}} \text{ fixed and power cost} + \text{fabric and endhost latency cost}$$

$$+ \underbrace{\sum_j \lambda_j(\sum_k w_k z_{jk} - p_j)}_{\text{Capacity violation punishment}} + \underbrace{\sum_{j\in\mathcal{J}}\sum_{k\in\mathcal{K}}\sum_{l\in\mathcal{K}} \zeta_{jkl}\left(z_{jk} - \bar{r}_{kl} z_{jl}\right)}_{\text{Resource ratio violation punishment}}$$

For any given $\mathbf{y}$, for data center j and resource k, optimizing z_{jk} in the Lagrangian subproblem becomes a solvable separate problem:

$$\min_{z_{jk}\geq 0} \phi_{jk} z_{jk} + \frac{\sum_{i\in\mathcal{I}} \tau_i d_i u_{ik} y_{ij}}{z_{jk} - \sum_i d_i u_{ik} y_{ij}},$$

where $\phi_{jk} = (\alpha c_j + \lambda_j)\, w_k + \sum_{l\in\mathcal{K}} (\zeta_{jkl} - \bar{r}_{lk}\zeta_{jlk})$.

Lagrangian with *Extremal Extended Polymatroid Cuts*

$$\textbf{(P-LC)} \quad \min_{\mathbf{x},\mathbf{y},\mathbf{v}} \quad \text{obj with } \mathbf{z}^*$$
$$\text{s.t.} \quad g_{jk}(\mathbf{y}_{\cdot j}) \leq v_{jk}, \quad \forall j \in \mathcal{J}, k \in \mathcal{K}$$

where $g_{jk}(\mathbf{y}_{\cdot j}) = \sqrt{\phi_{jk} \sum_{i \in \mathcal{I}} \tau_i d_i u_{ik} y_{ij}}$ is a submodular function.

Lagrangian with *Extremal Extended Polymatroid Cuts*

$$\textbf{(P-LC)} \quad \min_{\mathbf{x},\mathbf{y},\mathbf{v}} \quad \text{obj with } \mathbf{z}^*$$
$$\text{s.t.} \quad g_{jk}(\mathbf{y}_{\cdot j}) \leq v_{jk}, \quad \forall j \in \mathcal{J}, k \in \mathcal{K}$$

where $g_{jk}(\mathbf{y}_{\cdot j}) = \sqrt{\phi_{jk} \sum_{i \in \mathcal{I}} \tau_i d_i u_{ik} y_{ij}}$ is a submodular function.

Extremal Extended Polymatroid Cuts:

$$(\hat{\pi}_{jk})^{\mathsf{T}} \mathbf{y}_{\cdot j} \leq v_{jk}.$$

Structural Properties

Proposition

(Supermodularity for Endhost Latency) For each $j \in \mathcal{J}$*, the latency function* $\tilde{L}_j(\mathbf{y}_{\cdot j}, \mathbf{z}_{j\cdot}^-)$ *is increasing and supermodular in* $(\mathbf{y}_{\cdot j}, \mathbf{z}_{j\cdot}^-)$ *on* $\{0,1\}^{|\mathcal{I}|} \times \mathbb{R}_-^{|\mathcal{K}|}$.

Insights:

- Extra demand expand to the data center with lightest load;
- Extra budget to invest the data center with heaviest load.

Structural Properties

Corollary

(Free riding effect) Assigning a new demand d_i to data center j, and supplementing resource provisioning at j with $d_i u_{ik}$ for all k, then for all i' such that $y_{i'j} = 1$, the sojourn time of i' at data center j remains the same.

Structural Properties

Corollary

(Free riding effect) Assigning a new demand d_i to data center j, and supplementing resource provisioning at j with $d_i u_{ik}$ for all k, then for all i' such that $y_{i'j} = 1$, the sojourn time of i' at data center j remains the same.

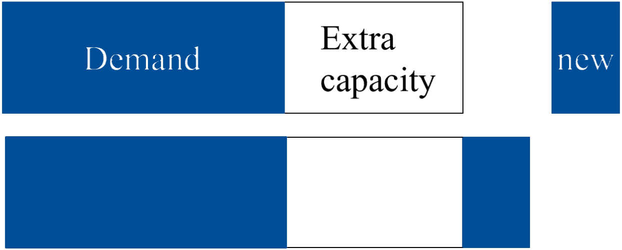

Structural Properties

Fix the location and allocation decisions, and suppose the capacity constraints and the resources ratio constraints are not binding, we have

(i) $\mathbf{z}^*$, obj^* concavely increase in the allocated demand
$\rightarrow$ economies of scale in allocated demand
$\rightarrow$ larger data centers are preferable;

(ii) $\mathbf{z}^*$ convexly decrease in unit power cost, while the obj^* concavely increases in unit power cost
$\rightarrow$ increasing marginal benefit of power efficiency;

(iii) $\mathbf{z}^*$, obj^* concavely increase in the unit endhost latency cost
$\rightarrow$ diminishing marginal effect of quality sensitivity.

Table of Contents

Basic facts

Real Dataset:

- U.S map with 49 demand notes.[Daskin, 1995]
- Demand rate is proportional to the number of households that own computers [U.S. Census Bureau, 2015].
- Each job requires only two types of resources, namely, computing and storage.
  - Computing: education level [U.S. Census Bureau, 2015].
  - Storage: high-speed Internet [U.S. Census Bureau, 2015].
- 36 candidate data center locations: 14 come from *DPR construction*, eight from Google, five from Microsoft, five from Facebook, and four from Amazon.
- The unit power costs are adopted from *the US Energy Information Administration*.

Base model

Our results reproduce many established data center sites of major warehouse-scaled cloud computing firms.

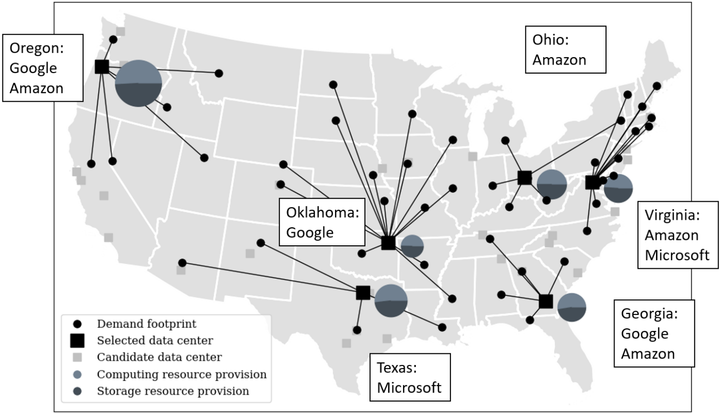

Base model — Latency Influences

Fabric	Endhost	#DC	Fabric	Endhost	#DC	Fabric	Endhost	#DC
1	1	10	5	1	16	1	1	20
1	10	9	5	10	15	10	10	20
1	80	11	5	80	16	10	80	19
1	100	12	5	100	17	10	100	19

- $\longrightarrow$: unit fabric latency cost $\nearrow$, DC build $\nearrow$:
- $\downarrow$: unit endhost latency cost $\nearrow$, DC build $\nearrow\!\searrow$:
  pooling effect for data centers with extra capacity;
  counter-pooling effect for data centers with tight capacity.

Insights:

- fabric latency: counter-pooling effect;
- endhost latency: pooling effect for data centers with extra capacity;
  counter-pooling effect for data centers with tight capacity.

Colocation

Fixed	Power	Colo	DC	Fixed	Power	Colo	DC	Fixed	Power	Colo	DC
0.01	1	39	4	0.01	1.5	31	5	0.01	2	17	6
0.05	1	20	5	0.05	1.5	7	6	0.05	2	0	7
0.1	1	8	6	0.1	1.5	0	7	0.1	2	0	7

- $\longrightarrow$: Colo unit power cost $\nearrow$, DC build $\nearrow$, Colo build $\searrow$;
- $\downarrow$: Colo fixed cost $\nearrow$, DC build $\nearrow$, Colo build $\searrow$.

Fabric	Endhost	Colo	DC	Fabric	Endhost	Colo	DC	Fabric	Endhost	Colo	DC
1	1	7	6	5	1	30	7	20	1	39	7
1	5	15	7	5	5	36	5	20	5	41	6
1	20	20	7	5	20	39	5	20	20	41	6

- $\longrightarrow$: unit fabric latency cost $\nearrow$, Colo build $\nearrow$;
- $\downarrow$: unit endhost latency cost $\nearrow$, Colo build $\nearrow$.

Colocation—Latency Influences

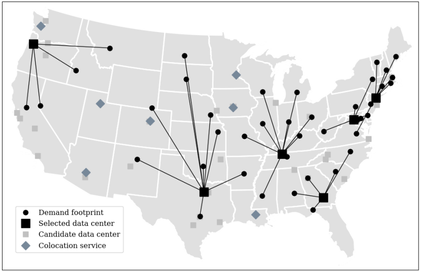

(a) Moderate latency-sensitivity

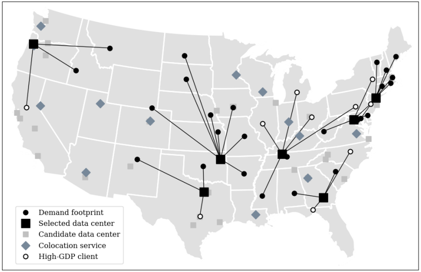

(b) High latency-sensitivity

Insights:

Colocations are more suitable for locations with low electricity rates, moderate demand volume and a long distance from regular sites.

Interdependent Footprints

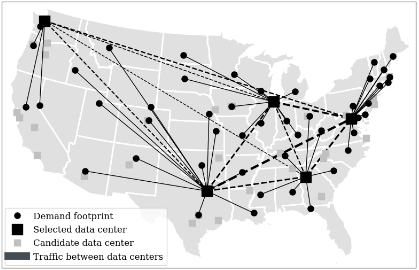

(a) Light traffic between data centers

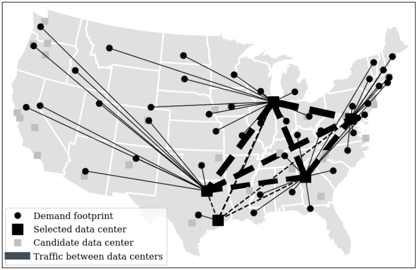

(b) Heavy traffic between data centers

Insights:

Higher interdependency → closer data centers and even pooling of multiple data centers.

Capacity Expansion with Average Shadow Price

Average shadow price [Kim and Cho, 1988, Crema, 1995].
Denote (**P**) to represent our original problem, V the objective value, and $V_j(w)$ the corresponding objective value if we increase the power capacity for the j-th candidate data center by w. Then, the average shadow price for the power constraint of data center j is defined as

$$p_j = \inf\{p \geq 0 : V - V_j(w) - pw \leq 0, \quad \forall\, w \geq 0\}.$$

# DC	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
ASP	1.02	7.74	4.41	2.92	6.20	0.00	16.79	0.00	22.55	4.75	0.00	3.24	5.47	0.00	1.49
LM	0.00	0.55	0.00	0.00	1.52	0.00	0.00	0.00	0.34	0.00	0.00	0.00	0.00	0.00	0.07

Note. "# DC" denotes the indices of data centers; "ASP" denotes the average shadow price; and "LM" denotes the Lagrangian multipliers.

Capacity Expansion with Average Shadow Price

Saved cost		Expanded capacity				
		10	100	1000	2000	5000
# DC	1	0.00	0.00	0.00	0.00	4944.29
	2	0.00	0.00	0.00	0.00	36021.29
	3	0.00	0.00	759.11	7724.34	18126.42
	4	0.00	0.00	0.00	3481.12	9952.03
	5	53.07	461.18	2184.29	11205.29	27904.51
	6	0.00	0.00	0.00	0.00	0.00
	7	170.34	1484.67	11316.64	11912.81	12561.41
	8	0.00	0.00	0.00	0.00	0.00
	9	223.24	216.33	8006.98	8595.65	8595.65
	10	0.00	0.00	0.00	7476.43	18885.93
	11	0.00	0.00	0.00	0.00	0.00
	12	0.00	0.00	0.00	5784.18	10798.03
	13	43.24	376.56	4234.69	6819.44	13143.15
	14	0.00	0.00	0.00	0.00	0.00
	15	0.00	0.00	2986.79	5157.86	5198.38

Note. "# DC" denotes the indices of data centers.

Computational Performance

# F	# S	Gurobi			Lagrangian				Lagrangian with cut			
		Ratio	Time	Gap	Ratio	Time	Iterations	Gap	Ratio	Time	Iterations	Gap
Loose Power Capacities												
10	5	100%	0.23	–	100%	0.29	3.26	–	100%	0.27	6.95	–
10	10	100%	0.45	–	100%	0.46	1.47	–	100%	0.36	2.78	–
20	10	100%	16.09	–	100%	1.48	4.00	–	100%	1.06	4.57	–
20	15	100%	11.26	–	100%	3.24	5.13	–	100%	1.27	3.40	–
30	15	100%	135.13	–	100%	5.52	7.32	–	100%	3.14	7.82	–
30	20	99%	93.31	5.98%	100%	28.36	20.00	–	100%	3.15	6.02	–
40	20	85%	260.12	7.78%	93%	131.15	79.12	8.39%	100%	50.83	81.02	–
50	25	82%	366.02	7.14%	93%	226.20	102.52	6.05%	100%	150.42	105.14	–
Medium Power Capacities												
10	5	100%	0.21	–	100%	0.32	3.93	–	100%	0.31	9.16	–
10	10	100%	0.33	–	100%	1.50	9.40	–	100%	0.47	4.72	–
20	10	100%	40.37	–	100%	5.23	22.07	–	100%	1.98	12.24	–
20	15	100%	43.73	–	100%	23.19	40.27	–	100%	2.52	9.35	–
30	15	90%	400.63	6.38%	97%	28.57	47.83	6.98%	100%	8.11	26.68	–
30	20	78%	394.39	7.55%	97%	65.11	54.61	8.28%	100%	9.16	23.67	–
40	20	17%	799.16	15.29%	88%	82.21	71.38	8.10%	100%	30.87	64.90	–
50	25	0%	–	20.44%	82%	141.42	84.27	8.73%	95%	211.39	147.61	7.92%
Tight Power Capacities												
10	5	100%	0.86	–	100%	0.39	4.89	–	100%	0.54	16.95	–
10	10	100%	2.79	–	100%	3.68	27.00	–	100%	1.12	17.49	–
20	10	88%	377.72	7.86%	100%	5.95	20.19	–	100%	3.49	17.02	–
20	15	73%	454.90	7.81%	100%	48.90	68.95	–	100%	4.05	18.34	–
30	15	12%	1099.78	11.59%	97%	51.59	106.61	7.21%	100%	15.62	58.07	–
30	20	3%	1475.89	11.07%	83%	93.81	80.72	8.53%	100%	19.16	58.02	–
40	20	0%	–	12.79%	78%	149.08	147.85	10.39%	95%	54.60	120.74	10.16%
50	25	0%	–	14.64%	63%	287.89	201.82	11.22%	75%	295.62	211.42	9.53%

Table of Contents

Summary

Modeling

- Endhost Latency—Queueing network;
- Practical extensions—Colo, Interdependency, Convex power, Congestion.

Methodology

- Non-trivial reformulations—MISOCP;
- Lagrangian with bounded heuristic and *Extremal Extended Polymatroid Cut* — better solution efficiency.

Insights

- Provable structural properties;
- Sensitivity analysis with real-world data;
  —Data center network design guidelines.

Summary

Modeling

- Endhost Latency—Queueing network;
- Practical extensions—Colo, Interdependency, Convex power, Congestion.

Methodology

- Non-trivial reformulations—MISOCP;
- Lagrangian with bounded heuristic and *Extremal Extended Polymatroid Cut* — better solution efficiency.

Insights

- Provable structural properties;
- Sensitivity analysis with real-world data;
  —Data center network design guidelines.

Thank you!