Algorithms for Right-Sizing Data Centers

Susanne Albers

TU Munich

Data centers

Electricity costs

Dominant, rapidly growing expense

18–28% of budget invested into energy

1.8% of total electricity worldwide
(90 million households)

Server utilization

Servers used 20–40% of the time

When idle & active, they use half of peak power

Data center

- Solution approach for energy and capacity management
  Transition idle servers into standby / sleep states.

- Challenge: Dynamically match number of active servers with the varying demand for computing capacity.

Contributions

- Albers. On Energy Conservation in Data Centers. SPAA 2017.
- Albers, Quedenfeld. Optimal Algorithms for Right-Sizing Data Centers. SPAA 2018.

Model 1

Dynamic Power Management (DPM)

- m heterogeneous servers
  Server: Several states with power consumption rates
  State transitions incur energy

- Planning horizon: $[t_1, t_n)$ with $t_1 < t_2 < \ldots < t_n$
  At least d_k servers must be active in $[t_k, t_{k+1})$

- Goal: Schedule minimizing energy consumption

Albers, SPAA ’17

Dynamic Power Management (DPM)

$\mathcal{I} = (\mathcal{S}, \mathcal{D})$

- $\mathcal{S} = \{S_1, \ldots, S_m\}$ power-heterogeneous servers

  S_i has states $s_{i,0}, \ldots, s_{i,\sigma_i}$

  $r_{i,j}$ = power consumption rate in $s_{i,j}$ $\quad\quad r_{i,0} > \ldots > r_{i,\sigma_i}$

  $\Delta_{i,j,j'}$ = energy to transition from state j to state j'

Dynamic Power Management (DPM)

$\mathcal{I} = (\mathcal{S}, \mathcal{D})$

- $\mathcal{D} = (T, D)$

  $T = (t_1, \ldots, t_n) \qquad t_1 < \ldots < t_n$

  $D = (d_1, \ldots, d_{n-1})$

  d_k servers must be in active state in $[t_k, t_{k+1})$

Schedule Σ specifies for each S_i and every $t \in [t_1, t_n)$ which state to use.

Goal: Minimize total energy

Previous work

- Power-down mechanisms on a single processor

  Minimize energy in an idle period

  Online algorithms with optimal competitiveness

  Augustine, Irani, Swamy '08;   Irani, Shukla, Gupta '03

  Ski-rental: Resource needed for a certain time.   Rent $\longleftrightarrow$ Buy

Previous work

- Machine activation: Activation cost budget
  Schedule activated machines so as to minimize makespan
  Algorithms approximating makespan / budget
  Activation cost non-decreasing function of load
  Khuller, Li, Saha ’10; Li, Khuller ’11

- Gap scheduling: Schedule jobs with release times, deadlines and processing times on homogeneous machines with one sleep state so as to minimize gaps.
  Demaine, Ghodsi, Hajiaghayi, Sayedi-Roshkhar, Zadimoghaddam ’13

Our contributions

- Definition of DPM:

  $\longrightarrow$ Power-heterogeneous servers

  $\longrightarrow$ Time horizon with varying demand for computing capacity

- Offline problem: Comprehensive study

  $\longrightarrow$ Data centers use past workload traces to identify demand in future time windows

Our results

- Each server has 1 sleep state

  Optimal solutions in poly-time using combinatorial algorithm

  Min-cost flow

- Each server has multiple standby/sleep states

  τ-approximation τ = # server types

  2-commodity min-cost flow

Model 2

m homogeneous servers, one sleep state

$t = 1, 2, \ldots, T$ in time horizon $[1, T]$

- operating cost: $f_t(\cdot)$ non-negative convex function
  $f_t(x_t)$ $\quad x_t$ = # active servers

- switching cost: $\beta(x_t - x_{t-1})^+$ $\quad \beta \in \mathbf{R}^+$
  $(x)^+ = \max\{0, x\}$

Schedule $X = (x_1, \ldots, x_T)$

$$\text{minimize} \sum_{t=1}^{T} f_t(x_t) + \beta \sum_{t=1}^{T} (x_t - x_{t-1})^+$$

Lin, Wierman, Andrew, Thereska '11 and '13

Previous work

Fractional problem: $x_t \in [0, m]$ real numbers

- Offline: optimal algorithm based on convex programming
  Online: 3-competitive algorithm LCP (Lazy Capacity Provisioning)
  Lin, Wierman, Andrew, Thereska ’11 and ’13

- Online: 2-competitive algorithm
  3-competitive memoryless algorithm
  $c \geq 1.86$, for any algorithm
  Bansal, Gupta, Krishnaswamy, Pruhs, Schewior, Stein ’15

Previous work

Fractional problem: $x_t \in [0, m]$ real numbers

- Online: $c \geq 2$, for any algorithm
  Antoniadis, Schewior '17

- Convex body chasing
  Antoniadis, Barcelo, Nugent, Pruhs, Schewior, Scquizzato '16

Our results

Discrete problem: $x_t \in [0, m]$ integers

- Offline: Polynomially solvable

  $O(T \log m)$ graph-based approach

- Online: Adaption of LCP 3-competitive

  Analysis: exploit discrete problem structure

  $c \geq 3$, for any algorithm $\quad c \geq 2$ in fractional problem

  Extensions:

  $\min \sum_{t=1}^{T} x_t f(\lambda_t / x_t) + \beta \sum_{t=1}^{T} (x_t - x_{t-1})^+$

  λ_t = incoming load at time t

  Finite lookahead

Open problems

Model 1

- Servers with multiple low-power states
  Improve approximation
- Online setting

Model 2

- Heterogeneous servers