Configuration Linear Programs in Online Non-linear Problems and in Algorithmic Game Theory

Nguyen Kim Thang
(IBISC, University Paris-Saclay)

Economies vs Diseconomies

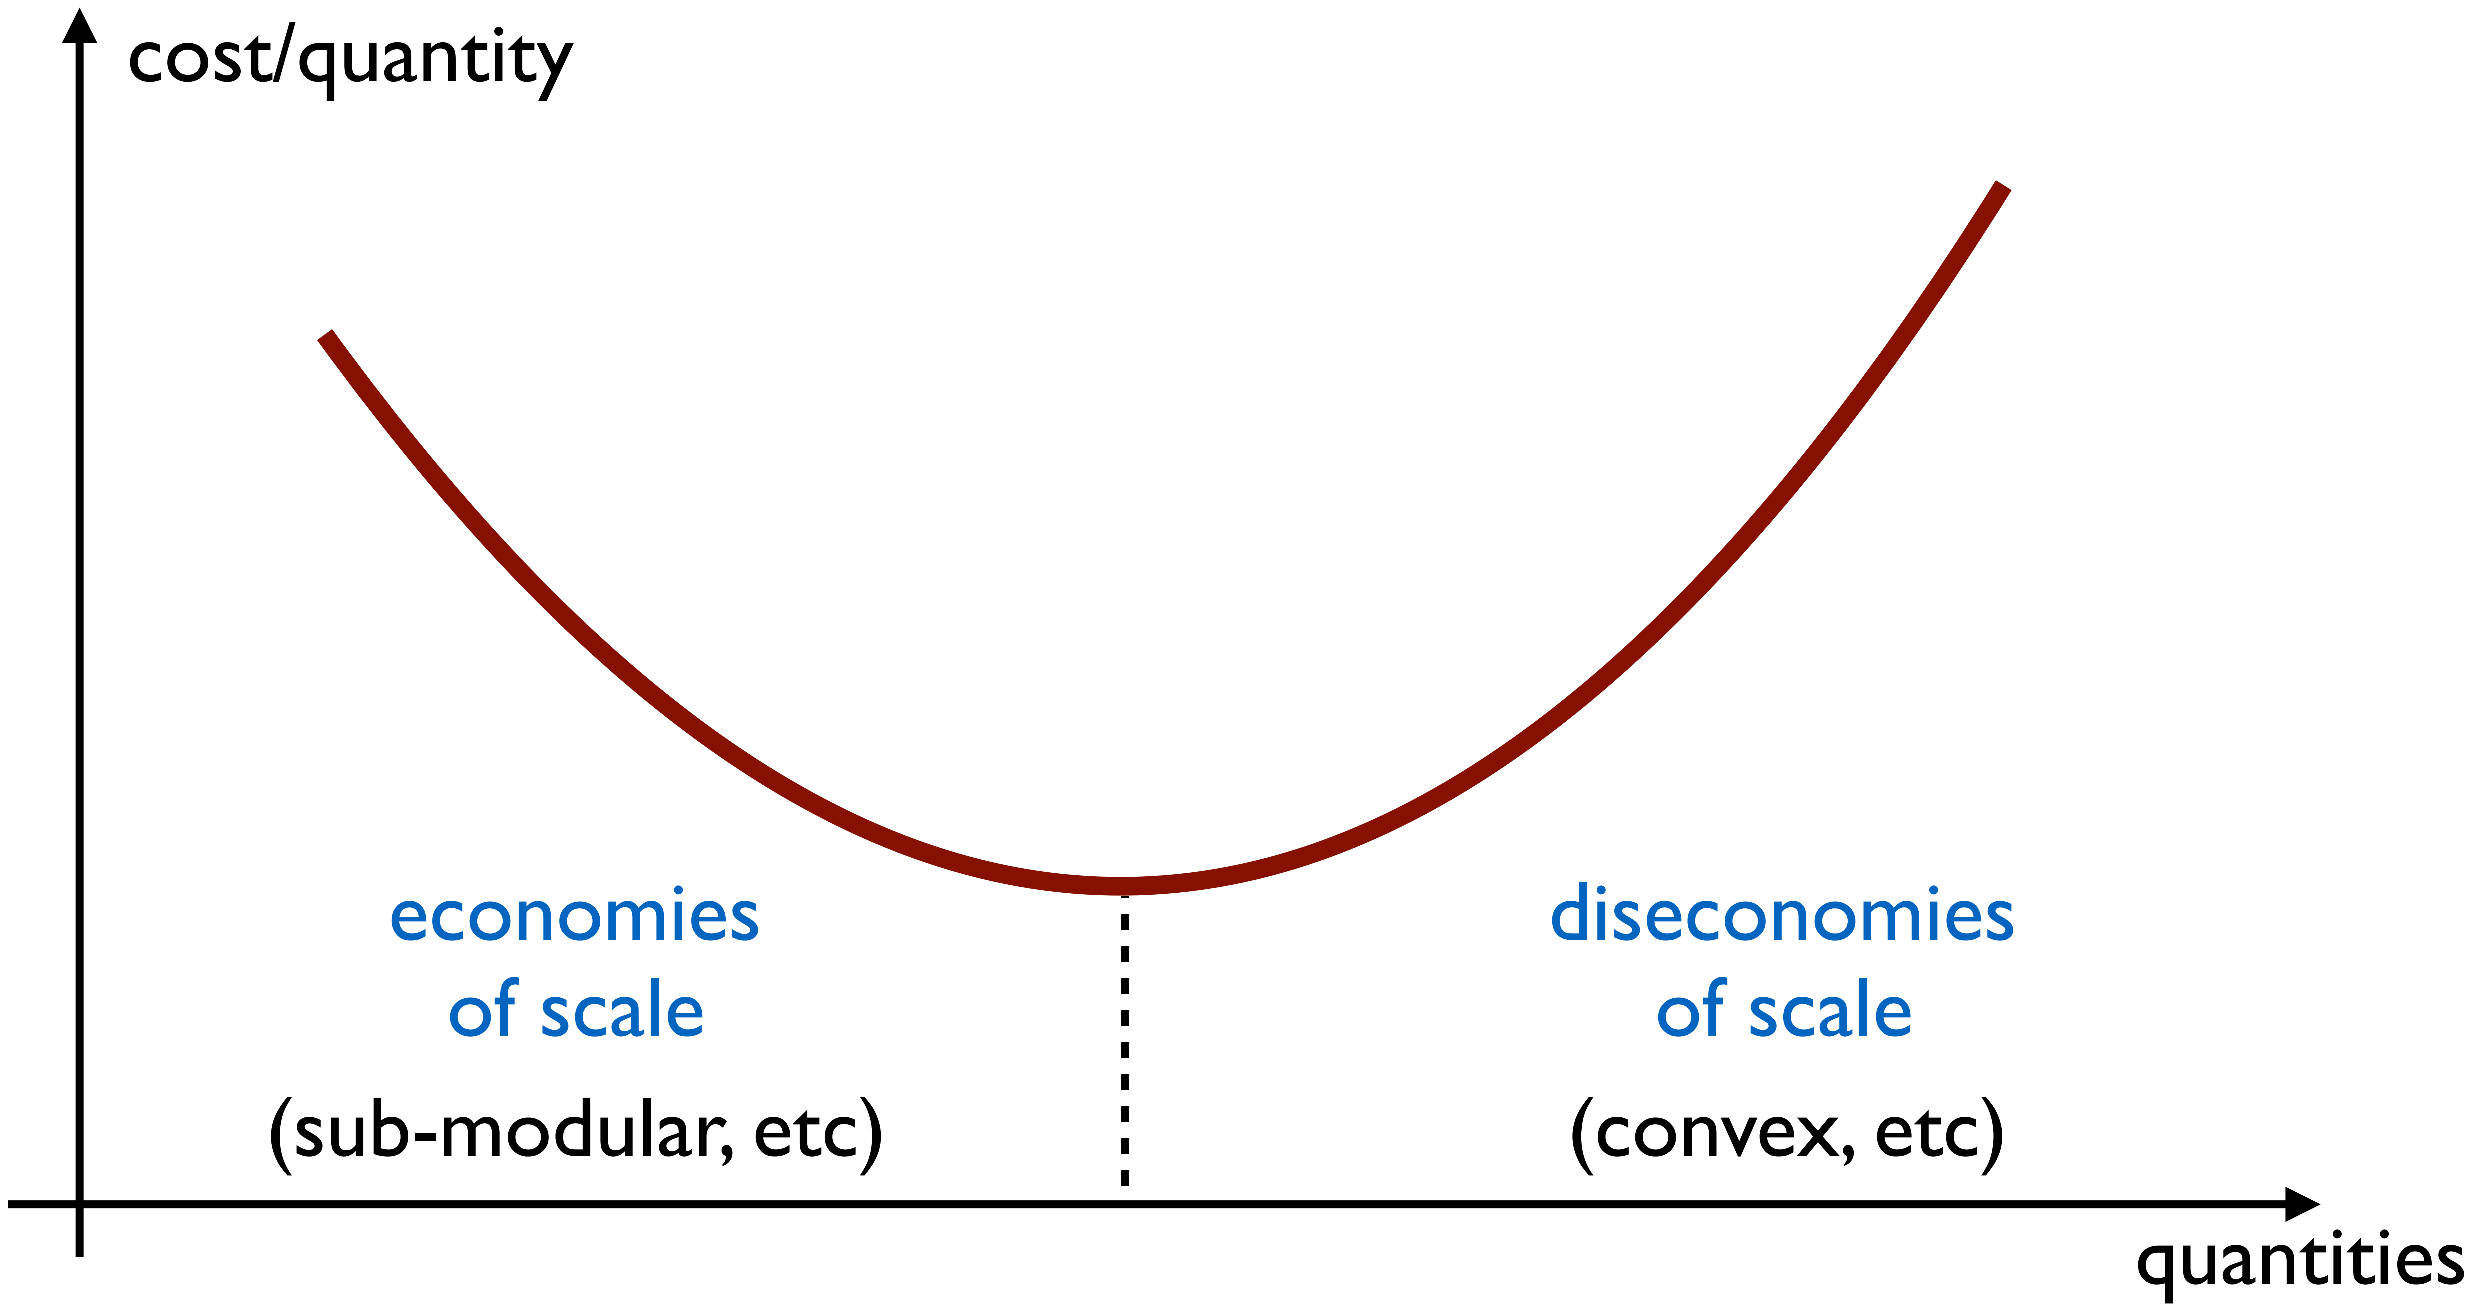

Economies vs Diseconomies

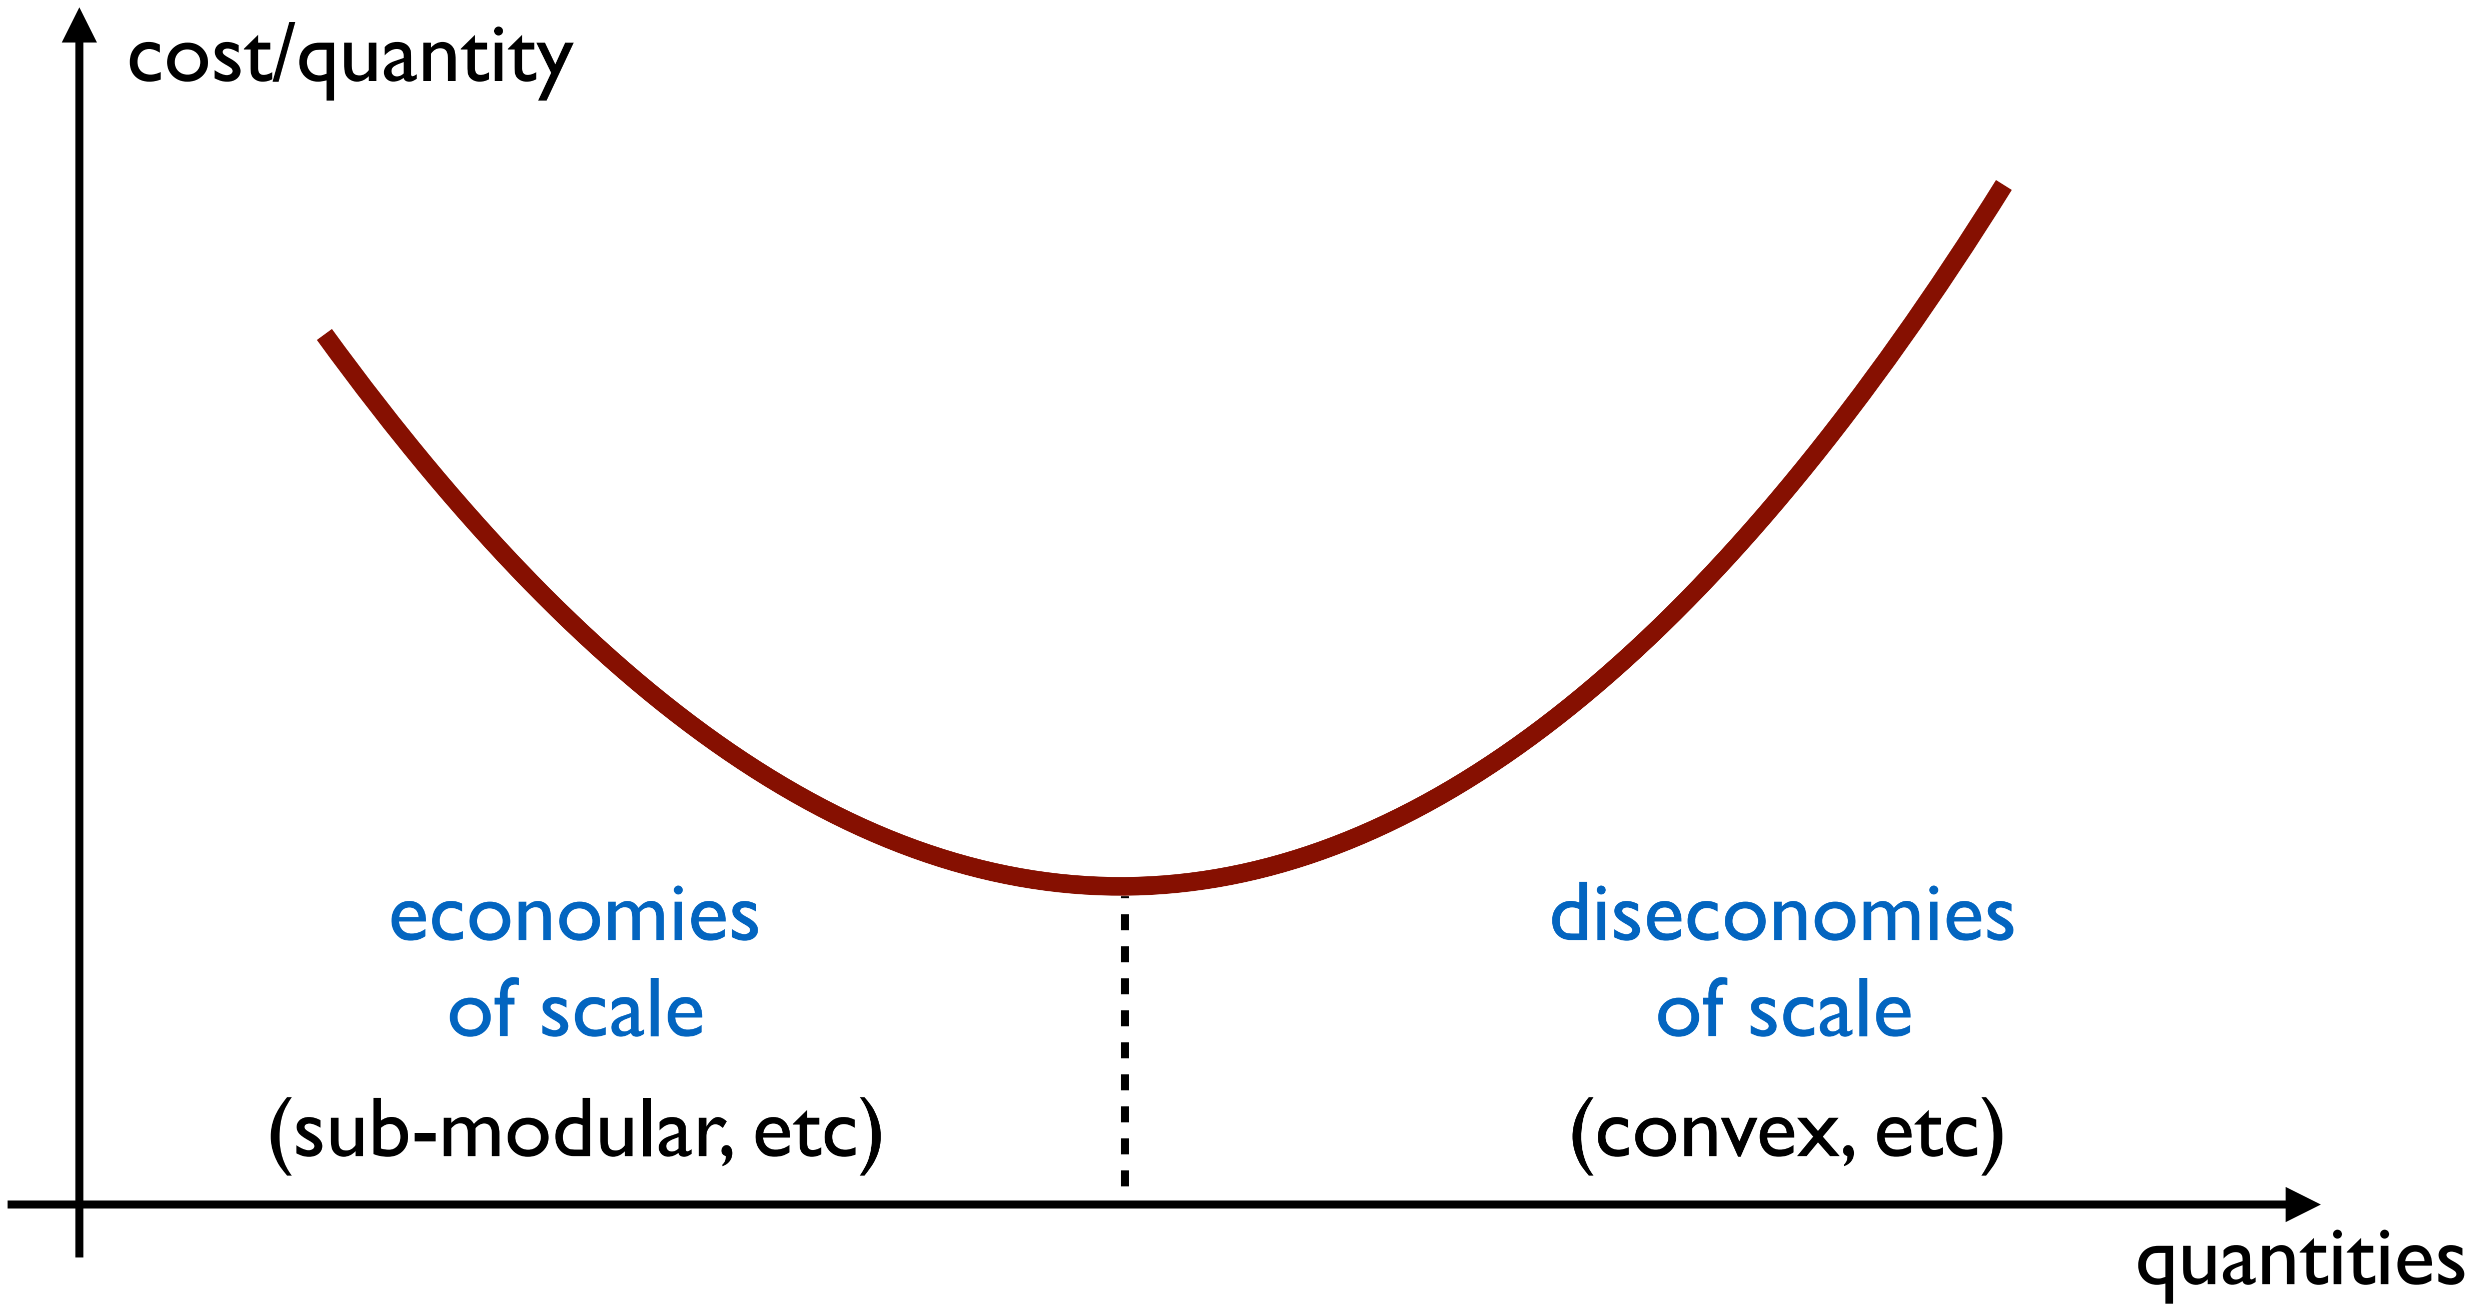

Arbitrarily-grown cost functions

Online Algs and Alg. Game Theory

- Online algorithms: requests arrive online, optimize some quality of service.

- Worst-case paradigm Competitive ratio = $\max_I ALG(I)/OPT(I)$

- Algorithmic Game Theory: players are self-interested, characterize the inefficiency of games.

- Worst-case paradigm Price of anarchy = $\max_I NE(I)/OPT(I)$

Settings

General online problem

Resources: R. The cost of a resource $f_e : 2^N \to \mathbb{R}^+$

Requests: arrive online. Set of feasible strategies of

$$\mathcal{S}_i = \{s_{ij} \subset R : 1 \leq j \leq m_i\}$$

Goal: minimize the total cost incurred on resources

Settings

General online problem

Resources: R. The cost of a resource $f_e : 2^N \rightarrow \mathbb{R}^+$

Requests: arrive online. Set of feasible strategies of

$$\mathcal{S}_i = \{s_{ij} \subset R : 1 \leq j \leq m_i\}$$

Goal: minimize the total cost incurred on resources

General game

Resources: R. The cost of a resource $f_e : 2^N \rightarrow \mathbb{R}^+$

Players: self-interested. Set of feasible strategies of

$$\mathcal{S}_i = \{s_{ij} \subset R : 1 \leq j \leq m_i\}$$

Goal: characterize the price of anarchy

Example

Energy minimization

Machine: unrelated machines, speed scalable

Jobs: arrive at r_j, deadline d_j, volume p_{ij}, preemptive non-migration

Energy: energy power function is $P(s(t))$, typically $s(t)^\alpha$

Goal: complete all jobs and minimize the total energy

Example

Energy minimization

Machine: unrelated machines, speed scalable

Jobs: arrive at r_j, deadline d_j, volume p_{ij}, preemptive non-migration

Energy: energy power function is $P(s(t))$, typically $s(t)^\alpha$

Goal: complete all jobs and minimize the total energy

☑ Known results:

online one machine (Bansal et al.'05) e^α-competitive

offline unrelated machines (Makarychev et al.) α^α-competitive

Examples

- Minimum Power Survival Routing
  - each request demands k-edge disjoint paths
  - costs on edges
  - minimize total cost

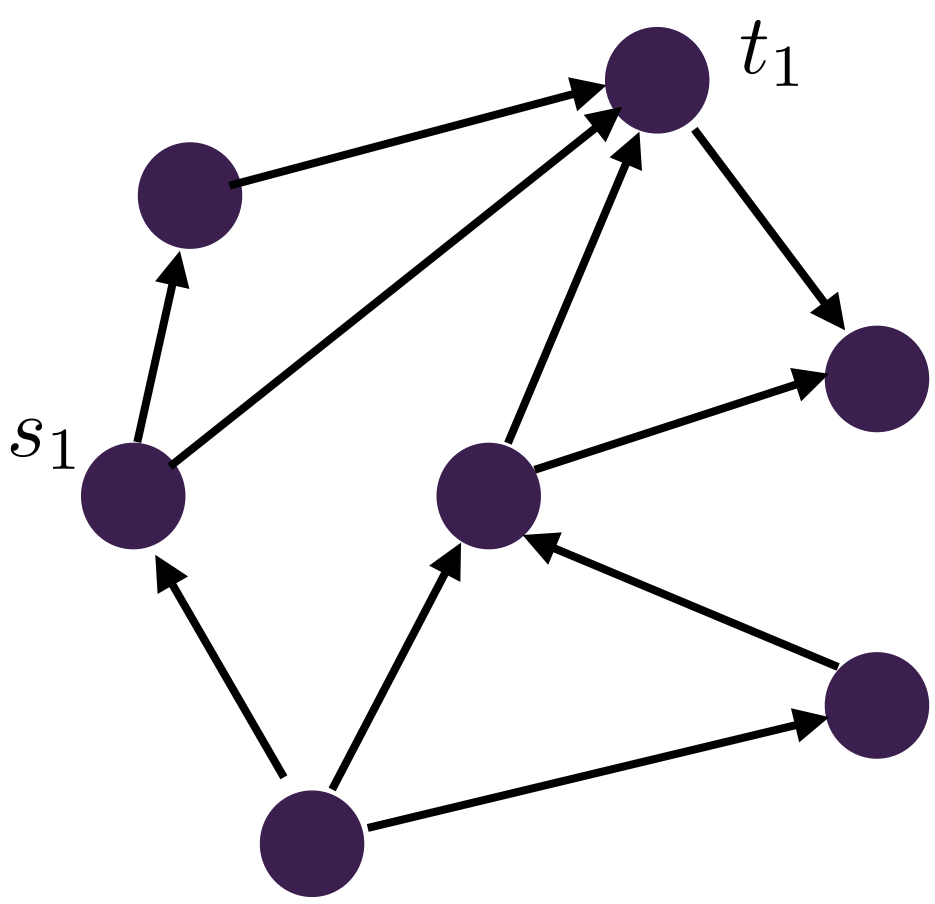

Examples

- Minimum Power Survival Routing
  - each request demands k-edge disjoint paths
  - costs on edges
  - minimize total cost

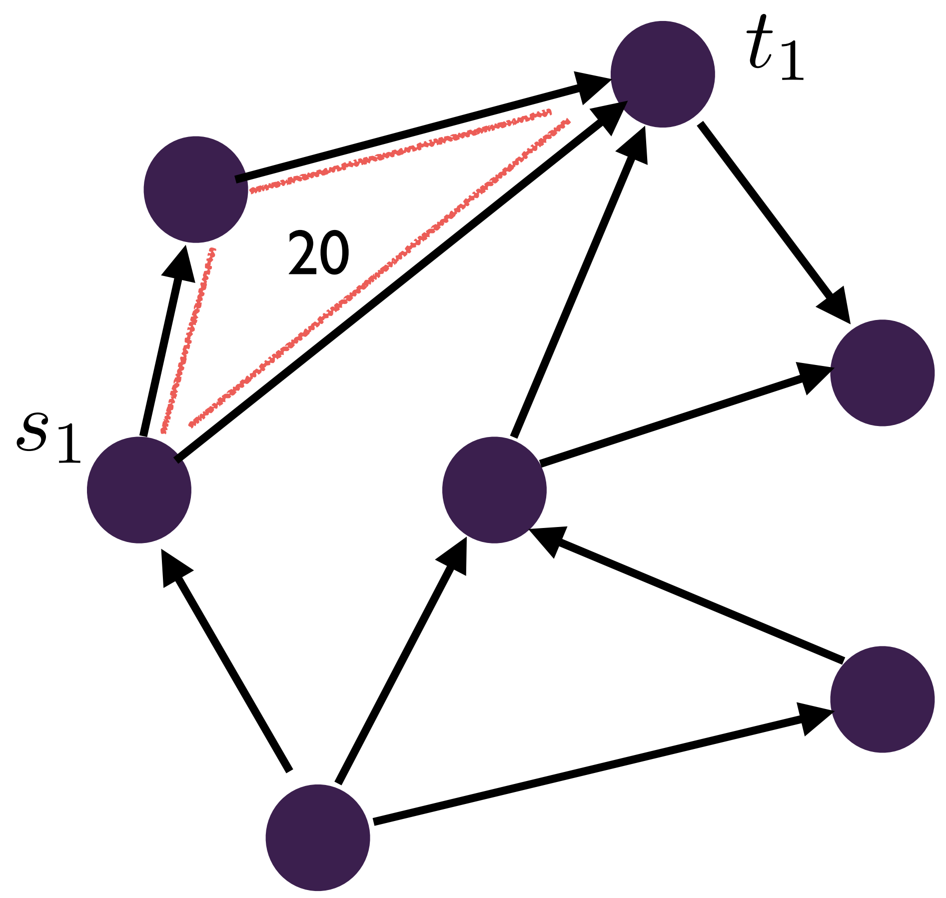

Examples

- Minimum Power Survival Routing
  - each request demands k-edge disjoint paths
  - costs on edges
  - minimize total cost

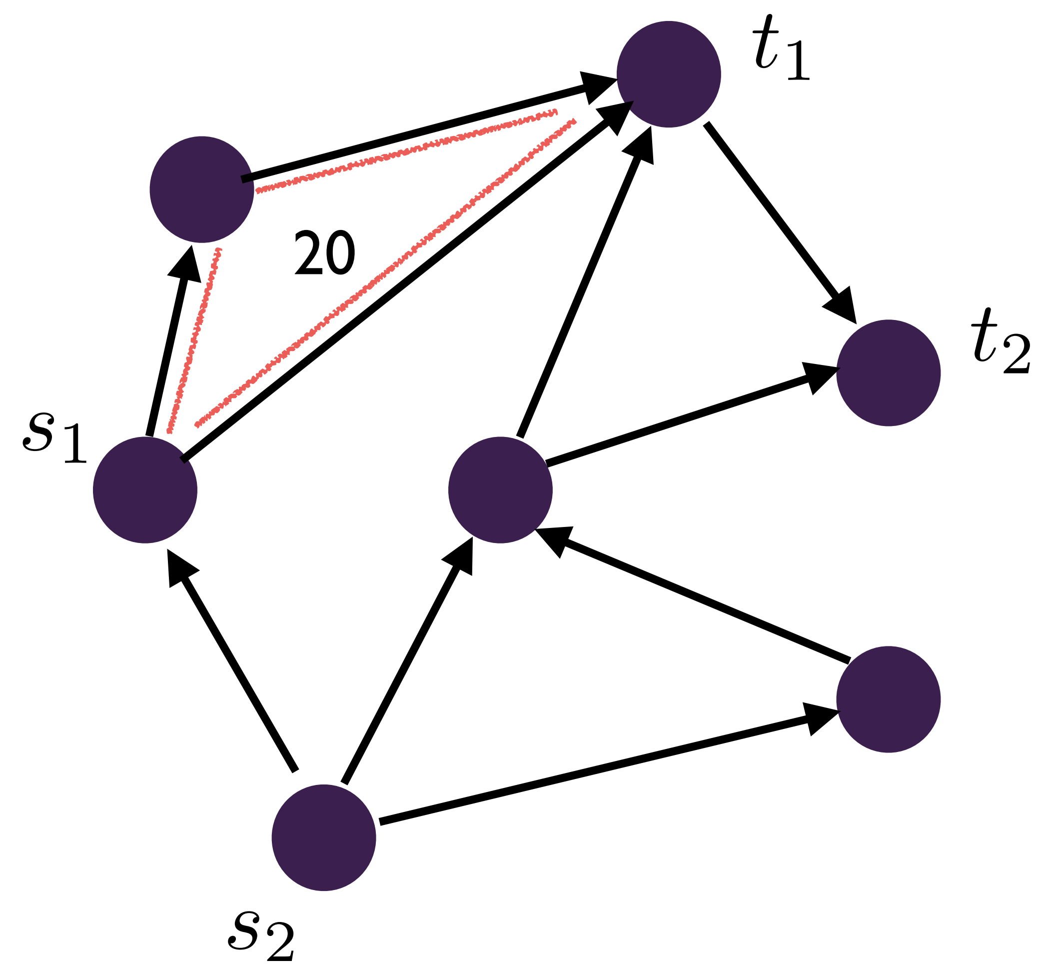

Examples

- Minimum Power Survival Routing
  - each request demands k-edge disjoint paths
  - costs on edges
  - minimize total cost

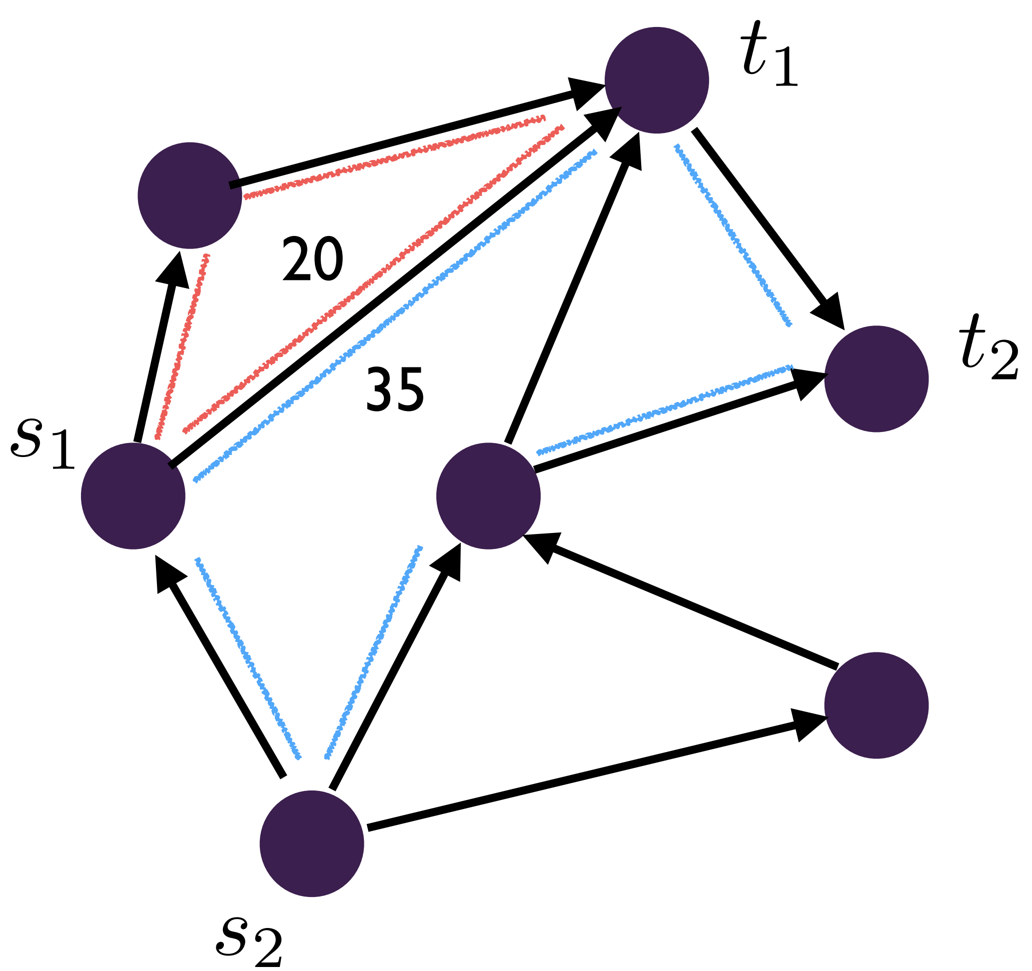

Examples

- Online Vector Scheduling
  - multiple machines
  - online multi-dimensional jobs
  - minimize the norms of the load vector.

Examples

- Online Vector Scheduling
  - multiple machines
  - online multi-dimensional jobs
  - minimize the norms of the load vector.
- Online Non-Convex Facility Location
  - clients assigned online to facility
  - facilities: opening cost + serving cost
  - minimize total cost.

Integrality gap

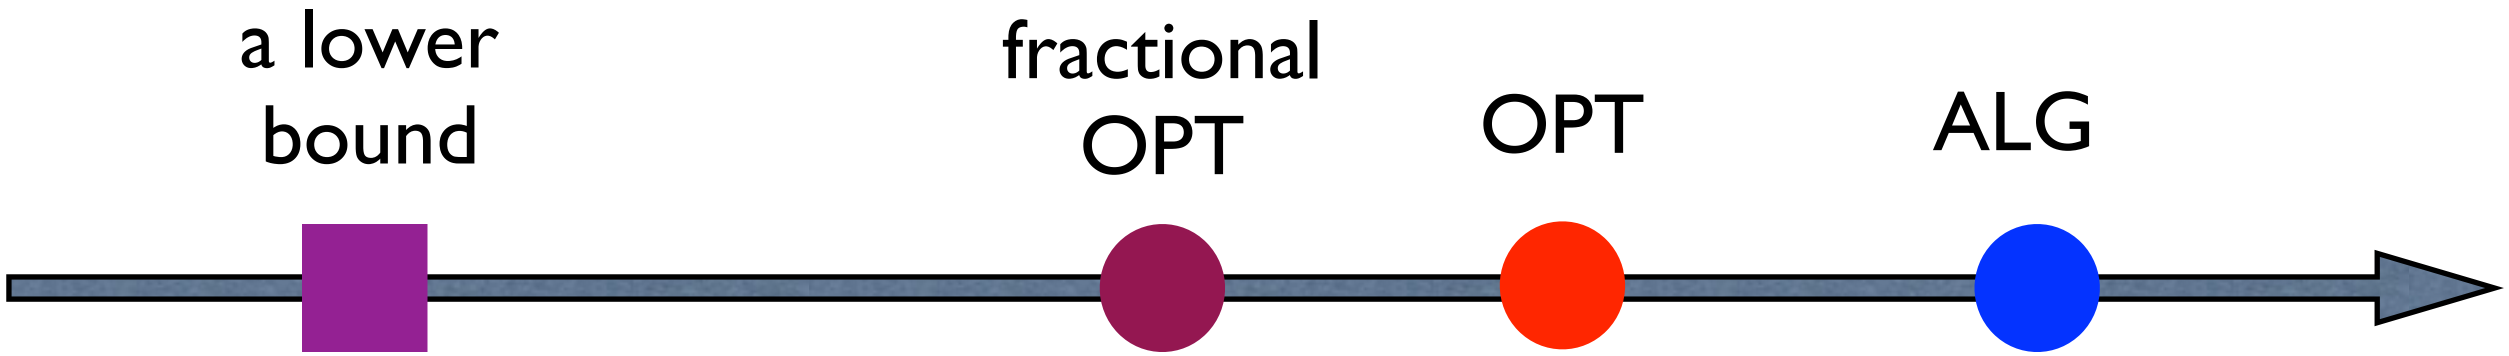

Integrality gap

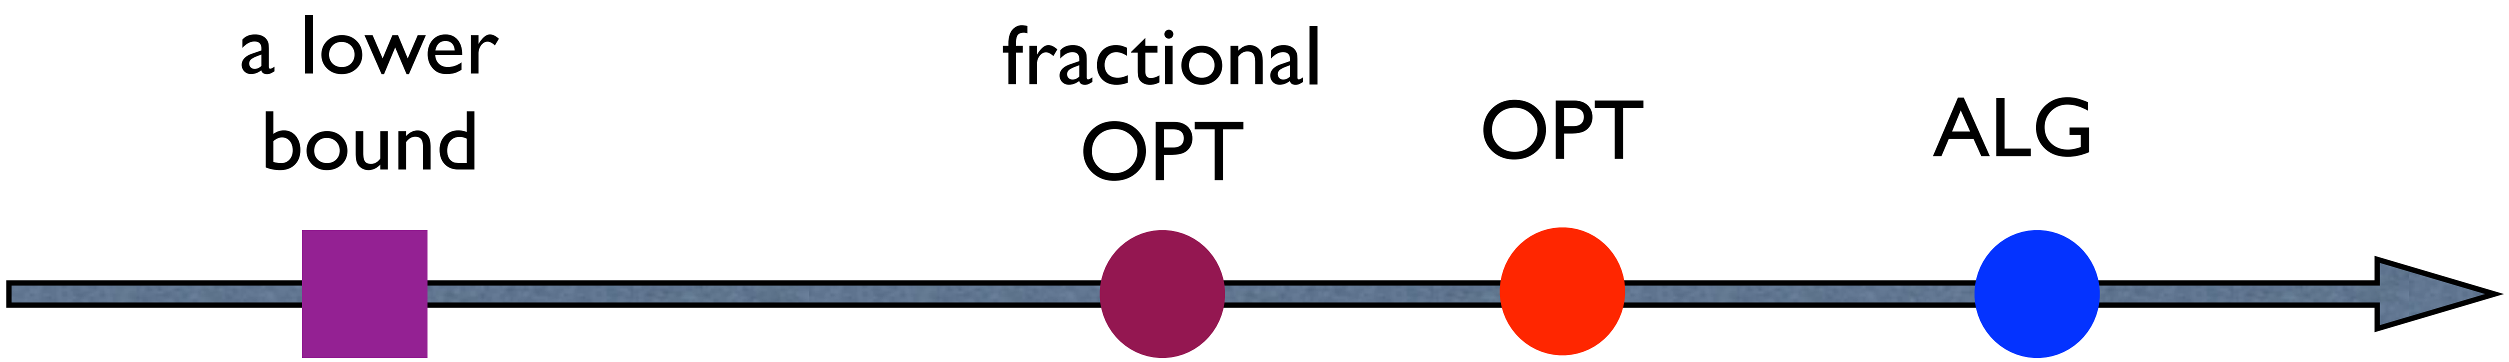

- Natural linear formulation: one request

$$\min \sum_{e=1}^{m} x_e^{\alpha}$$

$$\sum_{e=1}^{m} x_e = 1$$

$$x_e \in \{0,1\}$$

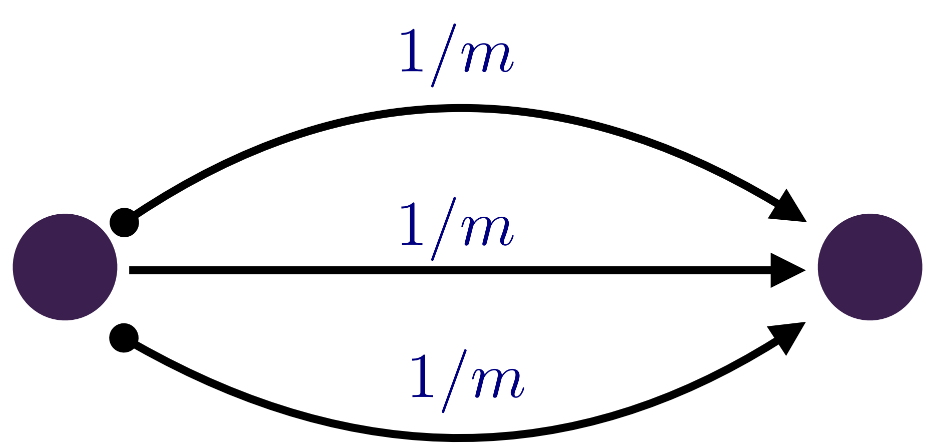

$$OPT = 1$$

$$OPT_f = m \cdot \frac{1}{m^{\alpha}}$$

Configuration LPs: a new way

- Systematically reduce integrality gap for (non-linear) problems.

- Design (online) primal-dual algorithms
  - No need of separation oracles and rounding (typical approaches for configuration LPs)
  - Light-weight algorithms.

Smoothness

- Definition: a function f is (λ, μ)-smooth if

$$\forall A_1 \subset A_2 \subset \ldots \subset A_n = A, B = \{b_1, \ldots, b_n\}$$

$$\sum_{i=1}^{n} [f(A_i \cup b_i) - f(A_i)] \leq \lambda \cdot f(B) + \mu \cdot f(A)$$

- Similar notion in algorithmic game theory (Roughgarden'15)

Configuration LP

A configuration A is subset of requests

$x_{ij} = 1$ if request i selects strategy $s_{ij} \in \mathcal{S}_i$

$z_{eA} = 1$ iff for every request $i \in A$, $x_{ij} = 1$
for some strategy $s_{ij} : e \in s_{ij}$

Configuration LP

A configuration A is subset of requests

$x_{ij} = 1$ if request i selects strategy $s_{ij} \in \mathcal{S}_i$

$z_{eA} = 1$ iff for every request $i \in A$, $x_{ij} = 1$
for some strategy $s_{ij} : e \in s_{ij}$

$$\min \sum_{e,A} f_e(A) z_{e,A}$$

$$\sum_{j: s_{ij} \in \mathcal{S}_i} x_{ij} = 1 \qquad \forall i$$

$$\sum_{A: i \in A} z_{eA} = \sum_{j: e \in s_{ij}} x_{ij} \qquad \forall i, e$$

$$\sum_{A} z_{eA} = 1 \qquad \forall e$$

$$x_{ij}, z_{eA} \in \{0, 1\} \qquad \forall i, j, e, A$$

Primal-Dual

$$\min \sum_{e,A} f_e(A) z_{e,A}$$

$$\sum_{j:s_{ij} \in \mathcal{S}_i} x_{ij} = 1$$

$$\sum_{A:i \in A} z_{eA} = \sum_{j:e \in s_{ij}} x_{ij}$$

$$\sum_{A} z_{eA} = 1$$

$$x_{ij}, z_{eA} \geq 0$$

$$\max \sum_{i} \alpha_i + \sum_{e} \gamma_e$$

$$\alpha_i \leq \sum_{e:e \in s_{ij}} \beta_{ie}$$

$$\gamma_e + \sum_{i \in A} \beta_{ie} \leq f_e(A)$$

Primal-Dual

$$\min \sum_{e,A} f_e(A) z_{e,A}$$

$$\sum_{j: s_{ij} \in \mathcal{S}_i} x_{ij} = 1$$

$$\sum_{A: i \in A} z_{eA} = \sum_{j: e \in s_{ij}} x_{ij}$$

$$\sum_{A} z_{eA} = 1$$

$$x_{ij}, z_{eA} \geq 0$$

$$\max \sum_{i} \alpha_i + \sum_{e} \gamma_e$$

$$\alpha_i \leq \sum_{e: e \in s_{ij}} \beta_{ie}$$

decision rule

$$\gamma_e + \sum_{i \in A} \beta_{ie} \leq f_e(A)$$

Primal-Dual

$$\min \sum_{e,A} f_e(A) z_{e,A}$$

$$\sum_{j:s_{ij} \in \mathcal{S}_i} x_{ij} = 1$$

$$\sum_{A:i \in A} z_{eA} = \sum_{j:e \in s_{ij}} x_{ij}$$

$$\sum_{A} z_{eA} = 1$$

$$x_{ij}, z_{eA} \geq 0$$

$$\max \sum_{i} \alpha_i + \sum_{e} \gamma_e$$

$$\alpha_i \leq \sum_{e:e \in s_{ij}} \beta_{ie}$$

decision rule

$$\gamma_e + \sum_{i \in A} \beta_{ie} \leq f_e(A)$$

smooth inequality

Primal-Dual

$$\min \sum_{e,A} f_e(A) z_{e,A}$$

$$\sum_{j: s_{ij} \in \mathcal{S}_i} x_{ij} = 1$$

$$\sum_{A: i \in A} z_{eA} = \sum_{j: e \in s_{ij}} x_{ij}$$

$$\sum_{A} z_{eA} = 1$$

$$x_{ij}, z_{eA} \geq 0$$

$$\max \sum_i \alpha_i + \sum_e \gamma_e$$

$$\alpha_i \leq \sum_{e: e \in s_{ij}} \beta_{ie}$$

decision rule

$$\gamma_e + \sum_{i \in A} \beta_{ie} \leq f_e(A)$$

smooth inequality

- **Algorithm:** at the arrival of a request, select a strategy that incurs the minimum marginal cost

Competitiveness

- Theorem: Assume that resource cost functions are (λ, μ) -smooth. Then the algorithm is $\lambda/(1-\mu)$-competitive.

Competitiveness

- **Theorem:** Assume that resource cost functions are (λ, μ)-smooth. Then the algorithm is $\lambda/(1-\mu)$-competitive.

- **Proof:**

$$\max \sum_i \alpha_i + \sum_e \gamma_e$$

$$\alpha_i \leq \sum_{e: e \in s_{ij}} \beta_{ie} \quad \forall i, j$$

$$\gamma_e + \sum_{i \in A} \beta_{ie} \leq f_e(A) \quad \forall e, A$$

$\alpha_i = \frac{1}{\lambda}$ (increase of the total cost due to the request)

$\beta_{i,e} = \frac{1}{\lambda}$ (increase of the cost on the resource if the request uses this resource)

$\gamma_e = -\frac{\mu}{\lambda}$ (the total cost of the resource)

Price of anarchy

- **Theorem:** Assume that resource cost functions are (λ, μ) -smooth. Then the price of anarchy is $\lambda/(1-\mu)$-competitive.

- **Proof:** Fix a Nash equilibrium

$$\max \sum_i \alpha_i + \sum_e \gamma_e$$

$$\alpha_i \le \sum_{e:e\in s_{ij}} \beta_{ie} \quad \forall i, j$$

$$\gamma_e + \sum_{i\in A} \beta_{ie} \le f_e(A) \quad \forall e, A$$

$\alpha_i = \frac{1}{\lambda}$ (cost of player i)

$\beta_{i,e} = \frac{1}{\lambda}$ (cost of player i on resource e)

$\gamma_e = -\frac{\mu}{\lambda}$ (cost of the Nash equilibrium)

Applications

- ☑ Corollary: If the cost functions are $f(z) = z^\alpha$ then the algorithm is $O(\alpha^\alpha)$-competitive. This is optimal for several problems.

- ☐ Proof:

The functions is $\left(\Theta(\alpha^{\alpha-1}), \frac{\alpha - 1}{\alpha}\right)$-smooth

Non-Convex Packing

Non-convex packing problem

Resources: R revealed online, one by one.

Constraints: offline $\sum_e b_{i,e} x_e \leq 1 \quad \forall i$

Goal: minimize a cost function of resources subject to the constraints.

Non-Convex Packing

Non-convex packing problem

Resources: R revealed online, one by one.

Constraints: offline $\sum_e b_{i,e} x_e \leq 1 \quad \forall i$

Goal: minimize a cost function of resources subject to the constraints.

- Known results:

  linear: elegant online primal-dual framework Buchbinder and Naor

  convex: recent online primal-dual framework (Azar et al.'16)

Configuration LP

$$\max \sum_S f(S) z_S$$

$$\sum_e b_{i,e} \cdot x_e \leq 1 \quad \forall i$$

$$\sum_{S:e \in S} z_S = x_e \quad \forall e$$

$$\sum_S z_S = 1$$

$$x_e, z_S \in \{0, 1\} \quad \forall e, S$$

$$\min \sum_i \alpha_i + \gamma$$

$$\sum_i b_{i,e} \cdot \alpha_i \geq \beta_e \quad \forall e$$

$$\gamma + \sum_{e \in S} \beta_e \geq f(S) \quad \forall S$$

$$\alpha_i \geq 0 \quad \forall i$$

$$z_S = \prod_{e \in S} x_e \prod_{e \notin S} (1 - x_e)$$

Algorithm

$F : [0,1]^n \to \mathbb{R}$ is the multilinear extension of f

$$F(x) := \sum_S f(S) \prod_{e \in S} x_e \prod_{e \notin S} (1 - x_e)$$

$$d := \max_i |\{b_{ie} : b_{ie} > 0\}| \qquad \rho := \max_i \max_{e,e' : b_{ie'} > 0} b_{ie}/b_{ie'}$$

- Algorithm

while $\sum_i b_{i,e}\alpha_i \leq \frac{1}{\lambda}\nabla_e F(x)$ & $\nabla_e F(x) > 0$ **do**

  - Increase x_e with rate $\dfrac{1}{\nabla_e F(x) \cdot \ln(1 + d\rho)}$
  - Increase α_i such that $\dfrac{\partial \alpha_i}{\partial \tau} \leftarrow \dfrac{b_{i,e} \cdot x_e}{\nabla_e F(x)} + \dfrac{1}{d\lambda}$

Result

- Definition: a function F is (λ, μ)-max-locally smooth if

$$\sum_{e \in S} \nabla_e F(x) \geq \lambda F\big(\mathbf{1}_S\big) - \mu F\big(x\big) \qquad \forall S, x.$$

- Theorem: Assume that cost function is (λ, μ)-locally smooth. Then there exists an algorithm with competitive ratio

$$O\Big(\frac{2\ln(1+d\rho)+\mu}{\lambda}\Big)$$

- Corollary: If the cost function is submodular then the algorithm has competitive ratio $O\big(\ln(1+d\rho)\big)$

Non-Convex Covering

Non-convex covering problem

Resources: R given offline.

Constraints: arrive online $\sum_e a_{i,e} x_e \geq 1 \qquad \forall i$

Goal: minimize the cost subject to the constraints.

Non-Convex Covering

Non-convex covering problem

Resources: R given offline.

Constraints: arrive online $\sum_e a_{i,e} x_e \geq 1 \qquad \forall i$

Goal: minimize the cost subject to the constraints.

- Known results:

  linear: elegant online primal-dual framework Buchbinder and Naor

  convex: recent online primal-dual framework (Azar et al.'16)

Result

- Definition: a function F is (λ, μ)-min-locally smooth if

$$\sum_{e \in S} \nabla_e F(\mathbf{x}) \leq \lambda \cdot F\big(\mathbf{1}_S\big) + \mu \cdot F\big(\mathbf{x}\big) \qquad \forall S \subset R$$

Result

- Definition: a function F is (λ, μ)-min-locally smooth if

$$\sum_{e \in S} \nabla_e F(\mathbf{x}) \leq \lambda \cdot F\big(\mathbf{1}_S\big) + \mu \cdot F\big(\mathbf{x}\big) \qquad \forall S \subset R$$

- Theorem: Assume that resource cost function is $\big(\lambda, \mu / \log m\big)$-locally smooth. Then there exists an algorithm that outputs a $\lambda/(1-\mu)$-competitive fractional solution.

Result

- Definition: a function F is (λ, μ)-min-locally smooth if

$$\sum_{e \in S} \nabla_e F(\mathbf{x}) \leq \lambda \cdot F\big(\mathbf{1}_S\big) + \mu \cdot F\big(\mathbf{x}\big) \qquad \forall S \subset R$$

- Theorem: Assume that resource cost function is $\big(\lambda, \mu / \log m\big)$-locally smooth. Then there exists an algorithm that outputs a $\lambda/(1-\mu)$-competitive fractional solution.

- Corollary: If the cost functions are $f(z) = z^\alpha$ then the algorithm is $O(\alpha^\alpha \log^\alpha m)$-competitive. This competitive ratio is tight.

Conclusion

☑ Primal-dual framework for non-linear/non-convex functions.

☑ Applicable for different classes of optimization problems.

☑ Applicable for game theory problems.

Conclusion

- ☑ Primal-dual framework for non-linear/non-convex functions.
- ☑ Applicable for different classes of optimization problems.
- ☑ Applicable for game theory problems.

Thank you!