Routing for Energy Minimization in the Speed Scaling Model

Antonio Fernández Anta

Universidad Rey Juan Carlos

Joint work with Matthew Andrews, Lisa Zhang, Wenbo Zhao

Network optimization problem

Problem statement

- Undirected network $G = (V, E)$, with n vertices and m edges.
- Set of traffic demands $(s_1, t_1), \ldots, (s_d, t_d)$.
- A cost function $f_e(x)$ for each $e \in E$.
- Objective: Route each demand in G so that the total cost $\sum_e f_e(x_e)$ is minimized.

Applications:

- Installation of oil pipes.
- Communication networks bandwidth provisioning.

Usually show economy of scale.

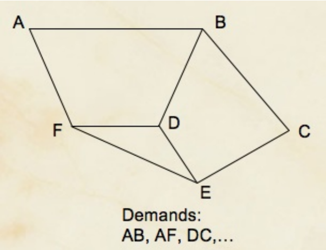

Network optimization problem

Formulation with linear constraints

$$
\begin{array}{rcll}
(P_0) & & \min \displaystyle\sum_e f_e(x_e) & \\
\text{subject to} & & & \\
x_e & = & \displaystyle\sum_i y_{i,e} & \forall e \\
y_{i,e} & \in & \{0,1\} & \forall i, e \\
y_{i,e} & : & \text{flow conservation} &
\end{array}
$$

$y_{i,e}$: Indicator of whether demand i is routed through edge e.
x_e: Load in edge e.

Cost functions

Linear functions

- Cost functions $f_e(x) = \mu \cdot x$.
- Shortest path routing is optimal.

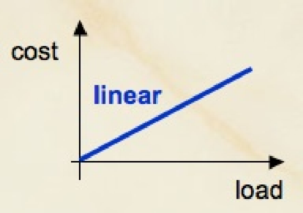

Cost functions

Subadditive functions

- Cost functions $f_e(x) + f_e(x') > f_e(x + x')$.
- Well-studied optimization problem: Buy-at-bulk.
- $f_e(x)$ uniform: $O(\log n)$-approx ratio, $\Omega(\log^{1/4} n)$ hardness.
- $f_e(x)$ nonuniform: $O(\log^3 n)$-approx ratio, $\Omega(\log^{1/2} n)$ hardness.

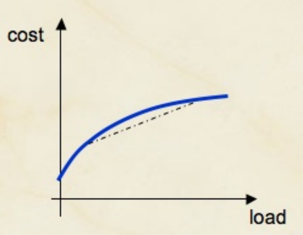

Cost functions

Other functions

- Functions that are not linear nor subadditive have been less studied.
- Energy consumption curves as function of processing speed fall in this class.

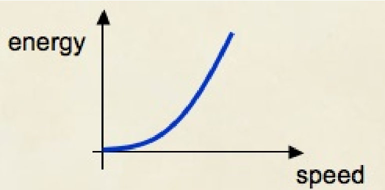

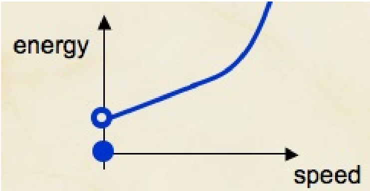

Why do we care?

Reasons for energy savings in ICT

- Growing energy bill (1.5% of US consumption in 2006).
- Pressure from governments and society to reduce carbon emissions.

Potential savings: 80% reduction at the access layer. 40% reduction in the network core.

How to save

- Speed scaling: Adapt the speed (hence energy) of network elements to the carried load.
- Globally optimize energy consumption in the network.

Contributions

Superadditive functions

- NP-hard with no finite approximation.

Polynomial functions

- $f_e(x) = \mu_e x^\alpha$: Constant-approximation.

Polynomial functions with startup

- $f_e(x) = \begin{cases} 0 & \text{if } x = 0, \\ \sigma_e + \mu_e x^\alpha & \text{if } x > 0. \end{cases}$
- $\Omega(\log^{1/4} n)$-hardness, for $\alpha > 1$.
- $O((\max_e\{\sigma_e/\mu_e\})^{1/\alpha} + 1)$-approximation.
- $O(d)$-approximation, for d demands.

Inapproximability of superadditive functions

- Reduction from edge-disjoint path (EDP):
  - Given network G and demands (s_i, t_i), decide if there are edge-disjoint routes.
  - NP-hard
- Consider monotone function $f_e(1) = 0$, $f_e(2) > 0$.
- If disjoint routes, $\sum_e f_e(x_e) = 0$, otherwise $\sum_e f_e(x_e) > 0$.
- Any bounded approximation returns 0 iff disjoint routes.

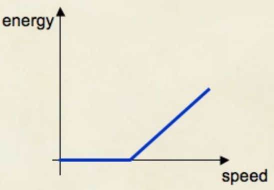

Polynomial functions

Functions of the form $f_e(x) = \mu_e x^\alpha$.

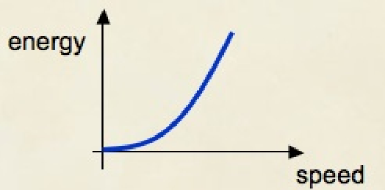

Approach

- Relax the integrality constrain $y_{i,e} \in \{0, 1\}$.
- Solve the problem.
- Apply randomized rounding.

Relaxed formulation

Convex program

$$(P_1) \qquad \min \sum_e \mu_e x_e^{\alpha}$$

subject to

$$x_e = \sum_i y_{i,e} \qquad \forall e$$

$$y_{i,e} \in [0,1] \qquad \forall i, e$$

$$y_{i,e} : \text{flow conservation}$$

- Polynomial-time solvable!

Approximation algorithm

Raghavan-Thompson randomized rounding

- Solve the convex program (P_1), and obtain optimal fractional solution (x^*, y^*).
- For each demand i, iteratively decompose the unit flow of i into paths according to the $y^*_{i,e}$ values:
  - Assign to path p weight $w_p = \min_{e \in p} y^*_{i,e}$.
  - For each $e \in p$, decrement $y^*_{i,e}$ by w_p.
- For each demand i, several flow paths $p_{i,1}, p_{i,2}, p_{i,3}, \ldots$ are obtained such that $\sum_j w_{p_{i,j}} = 1$.
- For each demand i, randomly choose one of the paths using the weights as probabilities. Let the output be $(\hat{x}, \hat{y})$.

Integrality gap of direct approach

- $f_e(x) = x^{\alpha}$ and only one demand (s, t).
- Optimal fraction solution is $x_1 = x_2 = \cdots = x_m = 1/m$.
- Optimal integral solution is $x_1 = 1, x_2 = \cdots = x_m = 0$.

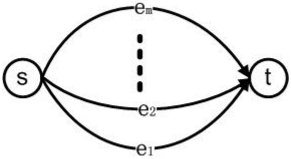

The integrality gap is

$$\frac{1}{m \cdot f_e(1/m)} = \frac{1}{m \cdot (1/m)^{\alpha}} = m^{\alpha-1}$$

Modified cost function

- The previous gap comes from the behavior of $f_e(x)$ in $[0, 1]$.
- We can use function $\mu_e x$, in this interval, since they agree in 0 and 1.
- We define a new cost function: $g_e(x) = \mu_e \max\{x_e^\alpha, x_e\}$

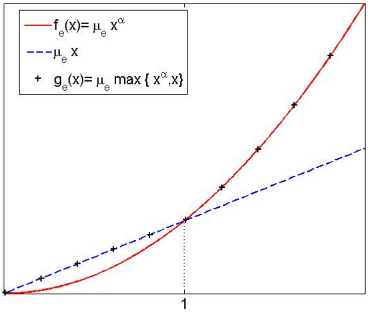

New formulation

Convex program

$$(P_1') \qquad \min \sum_e g_e(x_e)$$

subject to

$$\begin{aligned} x_e &= \sum_i y_{i,e} & \forall e \\ y_{i,e} &\in [0,1] & \forall i, e \\ y_{i,e} &: \text{ flow conservation} \end{aligned}$$

- $\sum_e g_e(x_e)$ and $\sum_e f_e(x_e)$ have the same integral optimal.

Constant approximation

Lemma

Randomized rounding the optimal fractional solution x_e^* guarantees that $E[f_e(\hat{x}_e)] \leq \gamma g_e(x_e^*)$ for all e.

Consider two cases:

- If $x_e^* \leq 1$, partition the possible values of $\hat{x}_e$ into the intervals $[0, 1), [1, 2), \ldots [2^j, 2^{j+1}), \ldots,$ and apply a Chernoff bound to $\Pr[\hat{x}_e \geq 2^j]$.
- If $x_e^* > 1$, partition the possible values of $\hat{x}_e$ into the intervals $[0, x_e^*), [x_e^*, 2x_e^*), \ldots [2^j x_e^*, 2^{j+1} x_e^*), \ldots,$ and apply a Chernoff bound to $\Pr[\hat{x}_e \geq 2^j x_e^*]$.

Constant approximation

Theorem

Randomized rounding provides a γ-approximation in the expected value of the total energy cost. For any constant c, randomized rounding provides a $c\gamma$-approximation with probability at least $1 - 1/c$.

This result also applies when demands are uniform and can be generalized to the case when demands are non-uniform.

Theorem

For nonuniform demands, randomized rounding can be used to achieve a $O(\log^{\alpha-1} D)$-approximation, where $D = \max_i d_i$.

Polynomial functions with startup: Hardness

$$f_e(x) = \begin{cases} 0 & \text{if } x = 0, \\ \sigma_e + \mu_e x^\alpha & \text{if } x > 0. \end{cases}$$

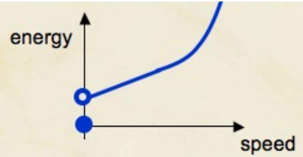

Theorem

For $\alpha > 1$ there is a uniform $f_e(x)$ such that no algorithm has $\log^{1/4} n$ approximation ratio unless $NP \subseteq ZPTIME(n^{polylogn})$.

The proof constructs an instance of the routing problem that encodes a 3CNF(5) formula. A $\log^{1/4} n$ approximation algorithm could be used to probabilistically decide satisfiability of a subset of clauses, contradicting the PCP theorem.

$O(d)$-approximation

Let us use the same functions $g_e(x)$ as before and consider the following convex program.

Convex program

$$(P_2) \qquad \min \sum_e \sigma_e z_e + g_e(x_e)$$

subject to

$$\begin{aligned} x_e &= \sum_i y_{i,e} && \forall e \\ z_e &= \max_i \{y_{i,e}\} && \forall e \\ y_{i,e} &\in [0,1] && \forall i, e \\ y_{i,e} &: \text{flow conservation} \end{aligned}$$

$O(d)$-approximation

Theorem

Randomized rounding provides a $O(d)$-approximation in the expected value of the total energy cost.

Follows from

- $E[f_e(\hat{x}_e)] \leq \gamma g_e(x_e^*)$
- $E[\hat{z}_e] = Pr(\hat{z}_e = 1) = 1 - Pr(\hat{y}_{i,e} = 0 \text{ for all } i)$
  $= 1 - \prod_i (1 - y_{i,e}^*) \leq \sum_i y_{i,e}^* \leq dz_e^*$.

For non-uniform demands, $O(d + \log^{\alpha-1} D)$ approximation ratio can be achieved, where $D = \max_i d_i$.

Approximation with fitting function

Let $s_e = (\sigma_e/(\alpha - 1)\mu_e)^{1/\alpha}$. Define

$$\beta_e = \begin{cases} \sigma_e + \mu_e & \text{if } s_e < 1, \\ \alpha\mu_e(\sigma_e/(\alpha-1)\mu_e)^{1-1/\alpha} & \text{if } s_e \geq 1. \end{cases}$$

$$g_e(x) = \begin{cases} \beta_e x & \text{if } x \in [0, \max(1, s_e)), \\ \sigma_e + \mu_e x^\alpha & \text{if } x \geq \max(1, s_e). \end{cases}$$

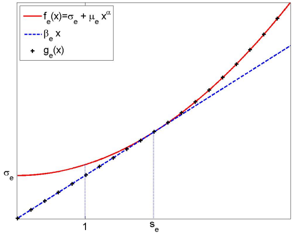

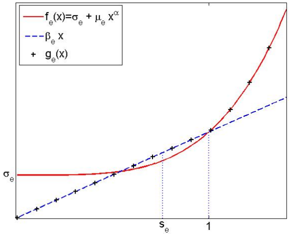

Formulation

Convex programming

$$(P_1') \qquad \min \sum_e g_e(x_e)$$

subject to

$$x_e = \sum_i y_{i,e} \quad \forall e$$

$$y_{i,e} \in [0,1] \quad \forall i, e$$

$$y_{i,e} : \text{flow conservation}$$

Approximation

Theorem

Randomized rounding to the fractional solution obtained from the convex program (P'_1), guarantees $E[f_e(\hat{x}_e)] \leq O(1 + (\sigma_e/\mu_e)^{1/\alpha}) \cdot g_e(x^*_e)$, for each link e. Hence, $O((\max_e\{\sigma_e/\mu_e\})^{1/\alpha} + 1)$-approximation.

The proof considers three cases

- $x^*_e \leq 1$. Implies $E[f_e(\hat{x}_e)] \leq O(1 + (\sigma_e/\mu_e)^{1/\alpha}) \cdot g_e(x^*_e)$.
- $x^*_e \geq \max(1, s_e)$. Implies $E[f_e(\hat{x}_e)] \leq \gamma \cdot g_e(x^*_e)$.
- $x^*_e \in (1, s_e)$. Implies $E[f_e(\hat{x}_e)] \leq O(1 + (\sigma_e/\mu_e)^{1/\alpha}) \cdot g_e(x^*_e)$.

For non-uniform demands a similar result with an additional $O(\log^{\alpha-1} D)$ factor, where $D = \max_i d_i$, can be shown.

Conclusions

- We have explored a min-cost integer routing problem for cost functions that model energy consumption.
- Positive and negative results have been shown for polynomial functions with and without startup cost.
- While pure polynomial functions can be efficiently routed (constant approximation), reaching the polylogarithmic lower bound seems to be hard when there are startup costs.
- Simulation results (omitted) show that the proposed algorithms work well in practice.

Future work

- For polynomial functions with startup cost, we would like to get approximation ratios independent of the functions and the demands, and as close as possible to the polylogarithmic bound.
- Other classes of functions can be explored.
- Additional restrictions in the solutions can be imposed (discrete speeds, capacity limits).
- Other models of energy savings can be explored: E.g., powering down the links.

Thanks!

Thank you!!