Deriving greedy algorithms and Lagrangian-relaxation algorithms

Neal E. Young

February 16, 2007

set cover

standard randomized rounding
- existence proof
- method of conditional probabilities
- algorithm

iterated sampling
- existence proof
- method of conditional probabilities
- algorithm

vertex cover (duality)
- existence proof
- method of conditional probabilities
- algorithm
- implicit primal-dual algorithm

multicommodity flow
- existence proof
- algorithm for integer solution
- algorithm for fractional solution

lower bound on iterations

fast algorithm for explicitly given problems

two open questions

set cover

input: collection $s_1, s_2, \ldots, s_m$ of sets over universe U

$$\text{minimize} \sum_{i=1}^{m} x_i \text{ subject to}$$

$$\begin{aligned} (\forall e \in U) \quad \sum_{s_i \ni e} x_i &\geq 1 \\ (\forall i) \quad x_i &\in \{0, 1\} \end{aligned}$$

- Value of optimal fractional solution x^*
  is a lower bound on optimal integer solution.

a fractional set cover x^*

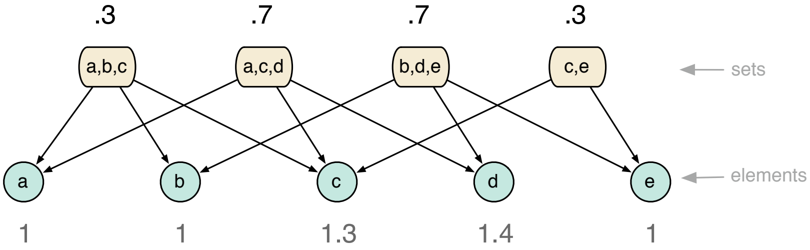

standard randomized rounding

Let x^* be an optimal fractional set cover.

Let $\lambda = \ln 2n$.

For each set $s_i \in S$ independently do:

choose s_i with probability $p_i \doteq \min\{\lambda x_i^*, 1\}$.

Theorem

With positive probability, chosen sets form a cover of size at most $2\ln(2n)\sum_i x_i^*$.

coverage

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For each set $s_i \in S$ independently do:
choose s_i with probability $p_i \doteq \min\{\lambda x_i^*, 1\}$.

Probability element e not covered:

$$\begin{aligned}
\prod_{s_i \ni e} 1 - p_i &< \prod_{s_i \ni e} \exp(-\lambda x_i^*) \\
&= \exp\left(-\lambda \sum_{s_i \ni e} x_i^*\right) \\
&\le \exp(-\lambda) \\
&= 1/2n
\end{aligned}$$

$$\Pr[\text{ exists uncovered element }] < 1/2$$

cost

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For each set $s_i \in S$ independently do:
choose s_i with probability $p_i \doteq \min\{\lambda x_i^*, 1\}$.

Expected number of sets chosen is

$$\sum_i p_i \;\leq\; \ln(2n) \sum_i x_i^*.$$

$$\Pr[\text{ more than } 2\ln(2n) \textstyle\sum_i x_i^* \text{ sets chosen }] \;\leq\; 1/2$$

proof of theorem

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For each set $s_i \in S$ independently do:
  choose s_i with probability $p_i \doteq \min\{\lambda x_i^*, 1\}$.

Theorem

With positive probability, chosen sets form a cover of size at most $2\ln(2n)\sum_i x_i^*$.

Proof.

$\Pr[\text{ exists uncovered element }] < 1/2$

$\Pr[\text{ more than } 2\ln(2n)\sum_i x_i^* \text{ sets chosen }] \le 1/2$

$\Pr[\text{ chosen sets form cover of size} \le 2\ln(2n)\sum_i x_i^* \text{ }] \ > \ 0$

□

method of conditional probabilities

converts existence proof into an efficient algorithm

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For each set $s_i \in S$ independently do:
choose s_i with probability $p_i \doteq \min\{\lambda x_i^*, 1\}$.

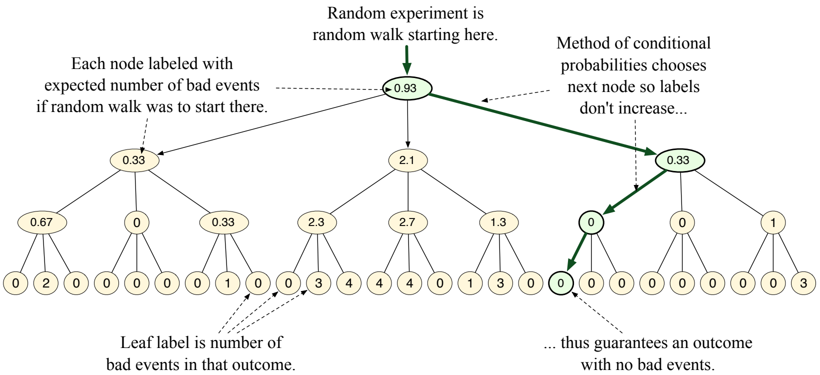

algorithm (incomplete)

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For each set $s_i \in S$ independently do:
choose s_i with probability $p_i \doteq \min\{\lambda x_i^*, 1\}$.

method of conditional probabilities

$\Downarrow$

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For $i = 1, 2, \ldots, m$ sequentially do:
include or *exclude* s_i — whichever keeps conditional probability of failure below 1.

conditional probability of failure

— coverage

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For $i = 1, 2, \ldots, m$ sequentially do:
include or *exclude* s_i — whichever keeps conditional probability of failure below 1.

Given first t choices, probability that elt e won't be covered is zero if e is already covered, and otherwise

$$\prod_{s_i \ni e, i > t} 1 - p_i .$$

Conditional probability that chosen sets will fail to cover is at most

$$\sum_{\substack{e \text{ not yet} \\ \text{covered}}} \prod_{s_i \ni e, i \geq t} 1 - p_i .$$

conditional probability of failure
— cost

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For $i = 1, 2, \ldots, m$ sequentially do:
include or *exclude* s_i — whichever keeps conditional probability of failure below 1.

Given first t choices, expected number of chosen sets is

$$\#\text{ first } t \text{ sets chosen} + \sum_{i>t} p_i.$$

Given first t choices, probability that too many sets will be chosen is at most

$$\frac{\#\text{ first } t \text{ sets chosen} + \sum_{i>t} p_i}{2 \ln 2n \sum_i x_i^*}.$$

pessimistic estimator Φ_t

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For $i = 1, 2, \ldots, m$ sequentially do:
include or *exclude* s_i — whichever keeps conditional probability of failure below 1.

Given first t choices, probability of failure is at most

$$\Phi_t \doteq \sum_{\substack{e \text{ not yet} \\ \text{covered}}} \prod_{s_i \ni e, i \geq t} 1 - p_i$$

$$+ \frac{\#\text{ first } t \text{ sets chosen} + \sum_{i>t} p_i}{2 \ln 2n \sum_i x_i^*}.$$

pessimistic estimator Φ_t

Let x^* be an optimal fractional set cover.
Let $\lambda = \ln 2n$.
For $i = 1, 2, \ldots, m$ sequentially do:
include or *exclude* s_i — whichever keeps conditional probability of failure below 1.

Given first t choices, probability of failure is at most

$$\Phi_t \doteq \sum_{\substack{e \text{ not yet} \\ \text{covered}}} \prod_{s_i \ni e, i \geq t} 1 - p_i$$

$$+ \frac{\#\text{ first } t \text{ sets chosen} + \sum_{i>t} p_i}{2 \ln 2n \sum_i x_i^*}.$$

- $\Phi_0 < 1$
- $E[\Phi_{t+1} \mid \Phi_t] \leq \Phi_t$
- If $\Phi_m < 1$, then outcome is successful.

algorithm

Let x^* be an optimal fractional set cover.

Let $\lambda = \ln 2n$.

For $i = 1, 2, \ldots, m$ sequentially do:

include or *exclude* s_i — whichever makes $\Phi_i < 1$.

$$\Phi_t \doteq \Big(\sum_{\substack{e \text{ not yet} \\ \text{covered}}} \prod_{s_i \ni e, i \geq t} 1 - p_i \Big) + \frac{\#\text{ first } t \text{ sets chosen} + \sum_{i>t} p_i}{2 \ln 2n \sum_i x_i^*}.$$

Corollary

Algorithm returns a cover of size at most $2\ln(2n) \times \text{OPT}$.

sample and increment

randomized rounding via iterated sampling

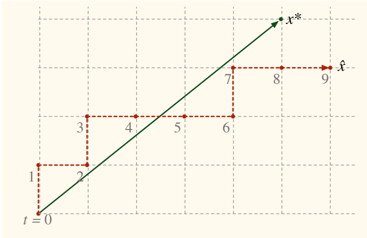

Let $x^* \geq 0$ be a fractional solution.

Let $|x^*|$ denote $\sum_i x_i^*$.

Define distribution p by $p_i \doteq x_i^* / \sum_{i'} x_{i'}^*$.

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots$ do:
  Sample random index i according to p.
  Increment $\hat{x}_i$.

Let $\hat{x}^{(t)}$ denote $\hat{x}$ after t samples.

... like weighted balls in bins.

illustration of sampling distribution

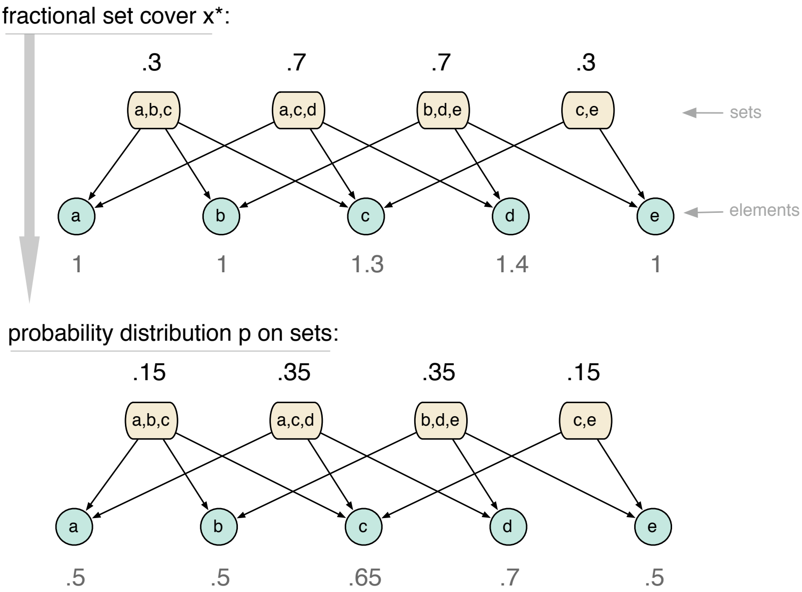

sample and increment

— for set cover

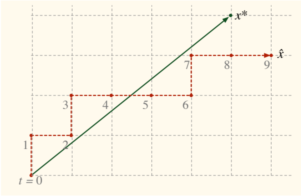

Let $x^* \geq 0$ be a fractional solution.

Let $|x^*|$ denote $\sum_i x_i^*$.

Define distribution p by $p_i \doteq x_i^*/|x^*|$.

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots$ do:
  Sample random index i according to p.
  Increment $\hat{x}_i$ — add s_i to the cover.

Let $\hat{x}^{(t)}$ denote $\hat{x}$ after t samples.

sample and increment

— for set cover

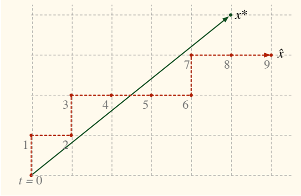

Let $x^* \geq 0$ be a fractional solution.

Let $|x^*|$ denote $\sum_i x_i^*$.

Define distribution p by $p_i \doteq x_i^*/|x^*|$.

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots$ do:
  Sample random index i according to p.
  Increment $\hat{x}_i$ — add s_i to the cover.

Let $\hat{x}^{(t)}$ denote $\hat{x}$ after t samples.

- For any element e, with each sample, $\Pr[e \text{ is covered}] = \sum_{s_i \ni e} x_i^*/|x^*| \geq 1/|x^*|$.

existence proof for set cover

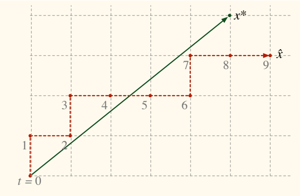

Theorem

With positive probability, after $T = \lceil \ln(n)|x^*| \rceil$ samples, $\hat{x}^{(T)}$ is a set cover.

Proof.

For any element e:

- With each sample,
  $\Pr[e \text{ is covered}] = \sum_{s_i \ni e} x_i^* / |x^*| \geq 1/|x^*|$.
- After T samples,
  $\Pr[e \text{ is not covered}] \leq (1 - 1/|x^*|)^T < 1/n$.

So, expected number of uncovered elements is less than 1. □

Corollary

There exists a set cover of size at most $\lceil \ln(n)|x^*| \rceil$.

method of conditional probabilities

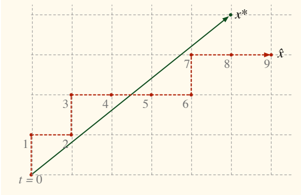

Let $x^* \geq 0$ be a fractional solution.

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots, T$ do:
  Increment $\hat{x}_i$, where i is chosen to keep
  expected number of not-covered elements below 1.

Return $\hat{x}^{(T)}$.

Given first t samples, expected number of not-covered elements is at most

$$\Phi_t \doteq \sum_{\substack{e \text{ not yet} \\ \text{covered}}} (1 - 1/|x^*|)^{T-t}.$$

algorithm

the greedy set-cover algorithm

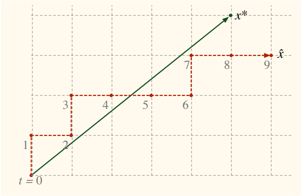

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots, T$ do:

Increment $\hat{x}_i$, where i is chosen to minimize the number of not-yet-covered elements.

Return $\hat{x}^{(T)}$.

Corollary

The greedy algorithm returns a cover of size at most $\lceil \ln(n) \min_{x^*} |x^*| \rceil$.

algorithm
the greedy set-cover algorithm

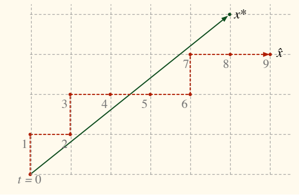

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots, T$ do:

Increment $\hat{x}_i$, where i is chosen to minimize the number of not-yet-covered elements.

Return $\hat{x}^{(T)}$.

Corollary

The greedy algorithm returns a cover of size at most $\lceil \ln(n) \min_{x^*} |x^*| \rceil$.

Can also derive Chvatal's weighted set cover algorithm and show $H(\max_s |s|)$-approximation.

vertex cover

attack via its dual — maximum matching

vertex cover: minimize $\sum_v y_v$ s.t. $(\forall e \in E) \quad \sum_{v \in e} y_v \geq 1$

matching: maximize $\sum_e x_e$ s.t. $(\forall v \in V) \quad \sum_{e \in v} x_e \leq 1$

Let x^* be a fractional matching.

Define probability distribution p on edges by $p_e \doteq x^*_e / |x^*|$.

Let $\hat{x} \leftarrow \mathbf{0}$. Say vertex v is *matched* when $\sum_{e \ni v} \hat{x}_e = 1$.

Repeat until each edge has a matched vertex:

> Sample an edge e from distribution p.
> If e has no matched vertex, then increment $\hat{x}_e$.

Return $\hat{x}$.

Let x^* be a fractional matching.
Define probability distribution p on edges by $p_e \doteq x^*_e / |x^*|$.
Let $\hat{x} \leftarrow \mathbf{0}$. Say vertex v is *matched* when $\sum_{e \ni v} \hat{x}_e = 1$.
Repeat until each edge has a matched vertex:
  Sample an edge e from distribution p.
  If e has no matched vertex, then increment $\hat{x}_e$.
Return $\hat{x}$.

Theorem

The expected size of the matching returned by sample-and-increment is at least $|x^*|/2$.

Proof.

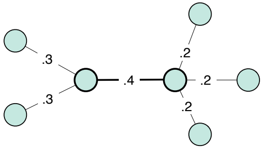

For any edge e,

$$\Pr[e \text{ chosen}] = \frac{p_e}{\sum_{e' : e \cap e' \neq \emptyset} p_{e'}} = \frac{x^*_e}{\sum_{e' : e \cap e' \neq \emptyset} x^*_{e'}} \geq \frac{x^*_e}{2}.$$

Theorem follows by linearity of expectation. □

matching

method of conditional probabilities

Let x^* be a fractional matching.
Define probability distribution p on edges by $p_e \doteq x^*_e / |x^*|$.
Let $\hat{x} \leftarrow \mathbf{0}$. Say vertex v is *matched* when $\sum_{e \ni v} \hat{x}_e = 1$.
Repeat until each edge has a matched vertex:
  Sample an edge e from distribution p.
  If e has no matched vertex, then increment $\hat{x}_e$.
Return $\hat{x}$.

Given the solution $\hat{x}^{(t)}$ after t samples,
the expected size $|\hat{x}^{(T)}|$ of the final matching is at least

$$\Phi_t \doteq |\hat{x}^{(t)}| + \sum_{\substack{e \text{ not yet} \\ \text{blocked}}} x^*_e/2.$$

Choosing an unblocked edge (u, v) and incrementing $\hat{x}_{(u,v)}$
increases Φ by at least

$$1 - \sum_{e \ni u} x^*_e/2 - \sum_{e \ni v} x^*_e/2$$

$$\geq 1 - 1/2 - 1/2 = 0.$$

algorithm

Let $\hat{x} \leftarrow \mathbf{0}$. Say vertex v is *matched* when $\sum_{e \ni v} \hat{x}_e = 1$.

Repeat until each edge has a matched vertex:

  Choose an edge e with no matched vertex. Increment $\hat{x}_e$.

Return $\hat{x}$.

Corollary

The algorithm returns a matching of size at least $(1/2) \max_{x^*} |x^*|$.

duality

primal: $\max c \cdot x : Ax \leq b$

dual: $\min b \cdot y : A^t y \geq c$

weak duality: x, y feasible $\Rightarrow c \cdot x \leq b \cdot y$, because

$$c^t x \;\leq\; (y^t A)x \;=\; y^t (Ax) \;\leq\; y^t b.$$

strong duality: *Every linear inequality*
that is valid for all feasible primal solutions x
can be expressed via weak duality.

duality

- Analysis of algorithm shows $|\hat{x}| \geq |x^*|/2$

  for any feasible solution x^*.

- Analysis *must be expressible via weak duality*.

weak duality relation for matching x / vertex cover y:

$$|x| = \sum_e x_e \leq \sum_e x_e \sum_{v \in e} y_v = \sum_v y_v \sum_{e \ni v} x_e \leq \sum_v y_v = |y|.$$

- Find implicit dual solution by looking for coefficients of x_e^* in the inequalities in the analysis.

dual solution implicit in analysis

look for coefficients of x_e^* in inequalities used in proof

$$\phi_t \doteq |\hat{x}^{(t)}| + \sum_{\substack{e \text{ not yet} \\ \text{blocked}}} x_e^*/2$$

Analysis showed $|x^*|/2 = \Phi_0 \le \Phi_T = |\hat{x}|$.

Want to recast as weak duality relation for some $\hat{y}$:

$$|x^*| = \sum_e x_e^* \le \sum_e x_e^* \sum_{v \in e} \hat{y}_v = \sum_v \hat{y}_v \sum_{e \ni v} x_e^* \le \sum_v \hat{y}_v = |y|.$$

Let $e_t = (u_t, v_t)$ be the edge chosen in the tth iteration.
Recall $\Phi_T \ge \Phi_0$ proved via $\sum_{t=0}^{T-1} \Phi_{t+1} - \Phi_t \ge 0$, via

$$\sum_t \Big(1 - \sum_{e \ni u_t} x_e^*/2 - \sum_{e \ni v_t} x_e^*/2\Big) \ \ge\ 0.$$

Rewrite to isolate coefficients of each x_e^*:

$$|\hat{x}| \ \ge\ \frac{1}{2} \sum_e x_e^* \sum_{v \in e} \sum_t [v \in e_t].$$

Suggests taking $\hat{y}_v \doteq \sum_t [v \in e_t]$, i.e. $\hat{y}_v = 1$ for matched vertices.

implicit primal-dual algorithm

Let $\hat{x} \leftarrow \mathbf{0}$. Say vertex v is *matched* when $\sum_{e \ni v} \hat{x}_e = 1$.

Let $\hat{y} \leftarrow \mathbf{0}$.

Repeat until each edge has a matched vertex:

> Choose an edge e with no matched vertex. Increment $\hat{x}_e$.
> For each $v \in e$, increment $\hat{y}_v$.

Return $\hat{x}$.

Corollary

The algorithm returns a feasible vertex cover $\hat{y}$, with $|\hat{y}| \leq 2|\hat{x}|$. Thus, the algorithm is a 2-approximation algorithm for VERTEX COVER.

maximum multicommodity flow

input: directed graph $G = (V, E)$, collection P of paths

$$\text{maximize} \sum_{p \in P} x_p \quad \text{s.t.} \quad (\forall e \in E) \sum_{p \ni e} x_p \le C$$

Let x^* be a fractional solution.

Define distribution q on paths by $q_p \doteq x^*_p / |x^*|$.

Let $\hat{x} \leftarrow \mathbf{0}$.

For $t = 1, 2, 3, \ldots$ do:

Sample random path p from distribution q; increment $\hat{x}_p$.

existence proof

Let x^* be a fractional solution.
Define distribution q on paths by $q_p \doteq x^*_p / |x^*|$.
Let $\hat{x} \leftarrow \mathbf{0}$.
For $t = 1, 2, 3, \ldots$ do:
Sample random path p from distribution q; increment $\hat{x}_p$.

Theorem

For $T = \lfloor |x^*| \rfloor$ and any $\varepsilon \in [0,1]$, the expected number of edges on which $\hat{x}^{(T)}$ induces flow greater than $(1+\varepsilon)C$ is at most

$$m \exp(-\varepsilon^2 C/3).$$

Proof.

Note expected flow on any edge is at most $TC/|x^*| \leq C$.
Apply Chernoff. □

Corollary

For $\varepsilon \doteq \sqrt{3\ln(m)/C}$, if $\varepsilon \leq 1$, there exists an integer flow of size at least $\lfloor |x^*| \rfloor$ that induces flow at most $(1+\varepsilon)C$ on each edge.

algorithm for integer multicommodity flow

after applying the method of conditional probabilities

Let $\hat{x} \leftarrow \mathbf{0}$. Let $\varepsilon \leftarrow \sqrt{3\ln(m)/C}$.

Repeat until $\hat{x}$ induces flow of $(1+\varepsilon)C$ on some edge:

> Let $\hat{x}(e)$ denote $\sum_{p \ni e} \hat{x}_p$, the flow on edge e.
> Choose path p to minimize $\sum_{e \in p}(1+\varepsilon)^{\hat{x}(e)}$.
> Increment $\hat{x}_p$.

Return $\hat{x}$.

Corollary

For $\varepsilon \doteq \sqrt{3\ln(m)/C}$, if $\varepsilon \leq 1$, the algorithm returns an integer flow of size at least $\lfloor \max_{x^*} |x^*| \rfloor$ that induces flow at most $(1+\varepsilon)C$ on each edge.

algorithm for fractional multicommodity flow

Additional input: ε. Idea: round to units of size $O(\varepsilon^2/\ln(m))$.

Let $\hat{x} \leftarrow \mathbf{0}$.

Choose λ so $\lambda C = 3\ln(m)/\varepsilon^2$.

Repeat until $\hat{x}$ induces flow of $(1+\varepsilon)\lambda C$ on some edge:

- Let $\hat{x}(e)$ denote $\sum_{p \ni e} \hat{x}_p$, the flow on edge e.
- Choose path p to minimize $\sum_{e \in p} (1+\varepsilon)^{\hat{x}(e)}$.
- Increment $\hat{x}_p$.

Return $\hat{x}/\lambda$.

Corollary

Given $\varepsilon \in [0,1]$, the algorithm returns a flow of size at least $\max_{x^*} |x^*|$ that induces flow at most $(1 + O(\varepsilon))C$ on each edge.

General alg. requires $3m\ln(m)/\varepsilon^2$ shortest-path computations.

a lower bound on number of iterations

critical dependence on $1/\varepsilon^2$ is inherent?

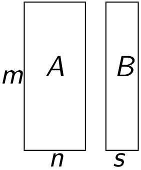

Define $V(A) \doteq \max\{|x| : Ax \leq \mathbf{1}\}$.

Theorem

Let $n \in \mathbb{N}$, $m = n^2$, and $\varepsilon > 0$ such that $\varepsilon^{-2} \leq n^{1-\Omega(1)}$.

Choose $A \in \{0,1\}^{m \times n}$ uniformly at random.

With probability $1 - o(1)$, for $s \leq \ln(m)/\varepsilon^2$, every $m \times s$ submatrix B of A satisfies

$$V(B) < (1 - \Omega(\varepsilon))V(A).$$

Proof.

Discrepancy argument based on "tightness" of Chernoff bound. □

a lower bound on number of iterations

$\Omega(\log(m)/\varepsilon^2)$ iterations are necessary

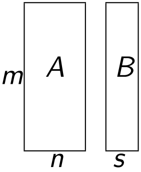

Corollary

Let $n \in \mathbb{N}$, $m = n^2$, *and* $\varepsilon > 0$ *such that* $\varepsilon^{-2} \leq n^{1-\Omega(1)}$.

Choose $A \in \{0,1\}^{m \times n}$ *uniformly at random.*

Then with probability $1 - o(1)$, *for the fractional packing problem of computing* $V(A)$, *any* $(1-\varepsilon)$*-approximate solution* $\hat{x}$ *has* $\Omega(\log(m)/\varepsilon^2)$ *non-zero entries* $\hat{x}_i$.

fast algorithm for explicitly given problems

reducing significance of $1/\varepsilon^2$

Theorem

A $(1 \pm \varepsilon)$*-approximate primal-dual pair for the linear program* $\max\{c \cdot x : Ax \geq b, x \geq \mathbf{0}\}$ *can be computed in expected time*

$$O(\#\textit{non-zeroes} + n \log(n)/\varepsilon^2)$$

where $n =$*(#constraints) + (#variables).*

Proof.

Clever use of duality, randomization, algorithmic engineering. □

(Strengthens and generalizes result by Grigoriadis and Khachiyan.)

two open questions

- SET COVER with *demands and multiplicity constraints* is

  $$\min\{c \cdot x : Ax \geq b, x \leq \mathbf{1}\}$$

  where A is $\{0,1\}$.

  The greedy algorithm is an $\ln(n)$-approximation algorithm.

  Is there a corresponding rounding scheme?

- For FACILITY LOCATION, the sample-and-increment rounding scheme gives a solution of expected cost at most

  $$\text{assignment-cost}(\textsc{opt}) + \ln(n) \times \text{facility-cost}(\textsc{opt}).$$

  Is there a corresponding greedy algorithm?

set cover

standard randomized rounding
- existence proof
- method of conditional probabilities
- algorithm

iterated sampling
- existence proof
- method of conditional probabilities
- algorithm

vertex cover (duality)
- existence proof
- method of conditional probabilities
- algorithm
- implicit primal-dual algorithm

multicommodity flow
- existence proof
- algorithm for integer solution
- algorithm for fractional solution

lower bound on iterations

fast algorithm for explicitly given problems

two open questions