An $O(\log n / \log\log n)$-approximation Algorithm for the Asymmetric Traveling Salesman Problem

and more recent developments

CATS @ UMD

April 22, 2016

The Asymmetric Traveling Salesman Problem (ATSP)

Problem (ATSP)

Given a set V if n points and a cost function $c : V \times V \to \mathbb{R}^+$, find a minimum cost tour that visits every vertex at least once.

The Asymmetric Traveling Salesman Problem (ATSP)

Problem (ATSP)

Given a set V if n points and a cost function $c : V \times V \to \mathbb{R}^+$, find a minimum cost tour that visits every vertex at least once.

- c is not necessarily symmetric.

The Asymmetric Traveling Salesman Problem (ATSP)

Problem (ATSP)

Given a set V if n points and a cost function $c : V \times V \to \mathbb{R}^+$, find a minimum cost tour that visits every vertex at least once.

- c is not necessarily symmetric.
- Every arc (u, v) in the tour can be replaced by the shortest path from u to v. Hence, we can assume c satisfies the triangle inequality.

The Asymmetric Traveling Salesman Problem (ATSP)

Problem (ATSP)

Given a set V if n points and a cost function $c : V \times V \to \mathbb{R}^+$, find a minimum cost tour that visits every vertex at least once.

- c is not necessarily symmetric.
- Every arc (u, v) in the tour can be replaced by the shortest path from u to v. Hence, we can assume c satisfies the triangle inequality.

The Asymmetric Traveling Salesman Problem (ATSP)

Problem (ATSP)

Given a set V if n points and a cost function $c : V \times V \to \mathbb{R}^+$, find a minimum cost tour that visits every vertex at least once.

- c is not necessarily symmetric.
- Every arc (u, v) in the tour can be replaced by the shortest path from u to v. Hence, we can assume c satisfies the triangle inequality.
- Integrality gap
  - Lower bound: 2.
  - Upper bound: $\log \log^{O(1)} n$.

The Traveling Salesman Problem (TSP)

i.e. Symmetric

- Metric-TSP:
  - APX-hard: 220/219.
  - 1.5-approximation by Christofides.

The Traveling Salesman Problem (TSP)

i.e. Symmetric

- Metric-TSP:
  - APX-hard: 220/219.
  - 1.5-approximation by Christofides.
- Euclidean-TSP:
  - PTAS by Arora and Mitchell.

The Traveling Salesman Problem (TSP)

i.e. Symmetric

- Metric-TSP:
  - APX-hard: 220/219.
  - 1.5-approximation by Christofides.
- Euclidean-TSP:
  - PTAS by Arora and Mitchell.
- Graph-TSP:
  - APX-hard.
  - Lower bound on integrality gap: 4/3.

The Traveling Salesman Problem (TSP)

i.e. Symmetric

- Metric-TSP:
  - APX-hard: 220/219.
  - 1.5-approximation by Christofides.
- Euclidean-TSP:
  - PTAS by Arora and Mitchell.
- Graph-TSP:
  - APX-hard.
  - Lower bound on integrality gap: 4/3.
  - Recent results breaking the 1.5 barrier.

Christofides

Algorithm

- Let T be the MST of G.
- Let O be the odd degree vertices of T. ($|O|$ is even.)
- Compute a minimum-weight perfect matching M for O.
- Combine the edges from M and T. (Every vertex has an even degree)
- Find an Eulerian circuit in $M \cup T$.
- Make the circuit Hamiltonian by skipping repeated vertices (shortcutting).

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.
  - Decompose OPT into a set of paths going from one vertex in O to the next in cyclic order.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.
  - Decompose OPT into a set of paths going from one vertex in O to the next in cyclic order.
  - Group all paths starting at a vertex with an even index, call it P_{even}. Similarly for odd indices we get P_{odd}.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.
  - Decompose OPT into a set of paths going from one vertex in O to the next in cyclic order.
  - Group all paths starting at a vertex with an even index, call it P_{even}. Similarly for odd indices we get P_{odd}.
  - $OPT = P_{\text{even}} \cup P_{\text{odd}}$. By averaging, either $c(P_{\text{even}}) \leq c(OPT)/2$ or $c(P_{\text{odd}}) \leq c(OPT)/2$.
  - Recognize that each group defines a matching on O.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.
  - Decompose OPT into a set of paths going from one vertex in O to the next in cyclic order.
  - Group all paths starting at a vertex with an even index, call it P_{even}. Similarly for odd indices we get P_{odd}.
  - $OPT = P_{\text{even}} \cup P_{\text{odd}}$. By averaging, either $c(P_{\text{even}}) \leq c(OPT)/2$ or $c(P_{\text{odd}}) \leq c(OPT)/2$.
  - Recognize that each group defines a matching on O.
  - It follows that $c(M) \leq c(OPT)/2$ as well.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.
  - Decompose OPT into a set of paths going from one vertex in O to the next in cyclic order.
  - Group all paths starting at a vertex with an even index, call it P_{even}. Similarly for odd indices we get P_{odd}.
  - $OPT = P_{\text{even}} \cup P_{\text{odd}}$. By averaging, either $c(P_{\text{even}}) \leq c(OPT)/2$ or $c(P_{\text{odd}}) \leq c(OPT)/2$.
  - Recognize that each group defines a matching on O.
  - It follows that $c(M) \leq c(OPT)/2$ as well.
- For $T \cup M$, $c(T) + c(M) \leq 1.5 \cdot c(OPT)$.

Christofides

Analysis

Let OPT be the optimal tour for G.

- Removing one edge from OPT yields a spanning tree. Hence, $c(T) \leq c(OPT)$.
- Number the vertices of O in cyclic order around OPT.
  - Decompose OPT into a set of paths going from one vertex in O to the next in cyclic order.
  - Group all paths starting at a vertex with an even index, call it P_{even}. Similarly for odd indices we get P_{odd}.
  - $OPT = P_{\text{even}} \cup P_{\text{odd}}$. By averaging, either $c(P_{\text{even}}) \leq c(OPT)/2$ or $c(P_{\text{odd}}) \leq c(OPT)/2$.
  - Recognize that each group defines a matching on O.
  - It follows that $c(M) \leq c(OPT)/2$ as well.
- For $T \cup M$, $c(T) + c(M) \leq 1.5 \cdot c(OPT)$.
- Shortcutting cannot increase the cost.

The Held-Karp Relaxation

Define $\delta^+(U) = \{a = (u, v) \in \overrightarrow{E} \mid u \in U, v \notin U\}$, and $\delta^-(U) = \delta^+(V \setminus U)$.

The Held-Karp Relaxation

Define $\delta^+(U) = \{a = (u, v) \in \overrightarrow{E} \mid u \in U, v \notin U\}$, and $\delta^-(U) = \delta^+(V \setminus U)$.

$$
\begin{array}{lll}
\text{minimize} & \displaystyle\sum_a c(a)x_a & \\
\text{subject to} & x(\delta^+(U)) \geq 1 & \forall U \subset V, \\
& x(\delta^+(v)) = x(\delta^-(v)) = 1 & \forall v \in V, \\
& x_a \geq 0 & \forall a.
\end{array}
\tag{1}
$$

The Held-Karp Relaxation

Define $\delta^+(U) = \{a = (u, v) \in \overrightarrow{E} \mid u \in U, v \notin U\}$, and $\delta^-(U) = \delta^+(V \setminus U)$.

$$\begin{array}{lll} \text{minimize} & \sum_a c(a)x_a & \\ \text{subject to} & x(\delta^+(U)) \geq 1 & \forall U \subset V, \\ & x(\delta^+(v)) = x(\delta^-(v)) = 1 & \forall v \in V, \\ & x_a \geq 0 & \forall a. \end{array} \tag{1}$$

Remark: the second set of constraints imply that $x(\delta^+(U)) = x(\delta^-(U))\ \forall U \subset V$.

The Held-Karp Relaxation over Spanning Trees

- Let x^* denote an optimum solution to the Held-Karp LP. Thus, $c(x^*) = \mathrm{OPT}_{\mathrm{HK}}$.

The Held-Karp Relaxation over Spanning Trees

Symmetrization

- Let x^* denote an optimum solution to the Held-Karp LP. Thus, $c(x^*) = \mathrm{OPT}_{\mathrm{HK}}$.
- Define $z^*_{\{u,v\}} = (1 - \frac{1}{n})(x^*_{uv} + x^*_{vu})$.

The Held-Karp Relaxation over Spanning Trees

Symmetrization

- Let x^* denote an optimum solution to the Held-Karp LP. Thus, $c(x^*) = \mathrm{OPT}_{\mathrm{HK}}$.
- Define $z^*_{\{u,v\}} = (1 - \frac{1}{n})(x^*_{uv} + x^*_{vu})$.
- Also define the cost of an edge $\{u, v\}$ as $\min\{c(u, v), c(v, u)\}$.

$\mathrm{OPT}_{\mathrm{HK}}$ in the Spanning Tree Polytope

Lemma

z^* belongs to the relative interior of the spanning tree polytope P.

$\mathrm{OPT}_{\mathrm{HK}}$ in the Spanning Tree Polytope

Lemma

z^* belongs to the relative interior of the spanning tree polytope P.

- It follows that z^* may be expressed as a convex combination of spanning trees, with strictly positive coefficients (marginal probabilities).

$\mathrm{OPT}_{\mathrm{HK}}$ in the Spanning Tree Polytope

Lemma

z^* belongs to the relative interior of the spanning tree polytope P.

- It follows that z^* may be expressed as a convex combination of spanning trees, with strictly positive coefficients (marginal probabilities).
- Next step: round z^* to a spanning tree.

Maximum Entropy Distribution

Let $\mathcal{T}$ be the collection of all spanning trees of G.

Maximum Entropy Distribution

Let $\mathcal{T}$ be the collection of all spanning trees of G.
Define the maximum entropy distribution p^* w.r.t z by the following convex program:

$$\begin{aligned} \text{infimum} \quad & \sum_{T \in \mathcal{T}} p(T) \log p(T) \\ \text{subject to} \quad & \sum_{T \ni e} p(T) = z_e \qquad \forall e \subset E, \\ & p(T) \geq 0 \qquad \forall T \in \mathcal{T}. \end{aligned} \tag{2}$$

Maximum Entropy Distribution

Let $\mathcal{T}$ be the collection of all spanning trees of G.
Define the maximum entropy distribution p^* w.r.t z by the following convex program (CP):

$$\begin{aligned} \text{infimum} \quad & \sum_{T \in \mathcal{T}} p(T) \log p(T) \\ \text{subject to} \quad & \sum_{T \ni e} p(T) = z_e \qquad \forall e \subset E, \\ & p(T) \geq 0 \qquad\qquad \forall T \in \mathcal{T}. \end{aligned} \tag{2}$$

Remark: the constraints imply that $\sum_{T \in \mathcal{T}} p(T) = 1$.

The Lagrange Dual

- For every edge $e \in E$, associate a Lagrange multiplier δ_e to the constraint for z_e.

The Lagrange Dual

- For every edge $e \in E$, associate a Lagrange multiplier δ_e to the constraint for z_e.
- Letting $\delta(T) = \sum_{e \in T} \delta_e$, it follows that

$$p(T) = e^{\delta(T)-1}.$$

The Lagrange Dual

- For every edge $e \in E$, associate a Lagrange multiplier δ_e to the constraint for z_e.
- Letting $\delta(T) = \sum_{e \in T} \delta_e$, it follows that

$$p(T) = e^{\delta(T)-1}.$$

- Apply a change of variables $\gamma_e = \delta_e - \frac{1}{n-1}$.

The Lagrange Dual

- For every edge $e \in E$, associate a Lagrange multiplier δ_e to the constraint for z_e.
- Letting $\delta(T) = \sum_{e \in T} \delta_e$, it follows that

$$p(T) = e^{\delta(T)-1}.$$

- Apply a change of variables $\gamma_e = \delta_e - \frac{1}{n-1}$.
- z being in the relative interior of P is a Slater condition, so strong duality holds and the Lagrange dual value equals $\mathrm{OPT}_{\mathrm{CP}}$.

Sampling Spanning Trees

Theorem

Given a vector z in the relative interior of the spanning tree polytope P on G, there exist γ_e^* for all $e \in E$ such that if we sample a spanning tree T of G according to $p^*(T) = e^{\gamma^*(T)}$, $Pr[e \in T] = z_e$ for every $e \in E$.

Sampling Spanning Trees

Theorem

Given a vector z in the relative interior of the spanning tree polytope P on G, there exist γ_e^* for all $e \in E$ such that if we sample a spanning tree T of G according to $p^*(T) = e^{\gamma^*(T)}$, $Pr[e \in T] = z_e$ for every $e \in E$.

- Suffices to find $\tilde{\gamma}_e$ while allowing $\tilde{z}_e \leq (1+\epsilon)z_e$ for $\epsilon = 0.2$.

Sampling Spanning Trees

Theorem

Given a vector z in the relative interior of the spanning tree polytope P on G, there exist γ_e^* for all $e \in E$ such that if we sample a spanning tree T of G according to $p^*(T) = e^{\gamma^*(T)}$, $Pr[e \in T] = z_e$ for every $e \in E$.

- Suffices to find $\tilde{\gamma}_e$ while allowing $\tilde{z}_e \leq (1+\epsilon)z_e$ for $\epsilon = 0.2$.
- The tree sampled is λ-random for $\lambda_e = e^{\lambda_e}$. This yields efficient sampling procedures and sharp concentration bounds using the tools developed for λ-random trees.

Sampling Spanning Trees

Theorem

Given a vector z in the relative interior of the spanning tree polytope P on G, there exist γ_e^* for all $e \in E$ such that if we sample a spanning tree T of G according to $p^*(T) = e^{\gamma^*(T)}$, $Pr[e \in T] = z_e$ for every $e \in E$.

- Suffices to find $\tilde{\gamma}_e$ while allowing $\tilde{z}_e \leq (1+\epsilon)z_e$ for $\epsilon = 0.2$.
- The tree sampled is λ-random for $\lambda_e = e^{\lambda_e}$. This yields efficient sampling procedures and sharp concentration bounds using the tools developed for λ-random trees.
- Namely, the events $[e \in T]$ are negatively correlated.

Concentration Bounds

Theorem

For each edge e, let X_e be an indicator random variable associateed with the event $[e \in T]$, where T is a λ-random tree. Also, for any subset C of the edges of G, define $X(C) = \sum_{e \in C} X_e$. Then we have

$$Pr[X(C) \geq (1+\delta)\mathbb{E}[X(C)]] \leq (\frac{e^{\delta}}{(1+\delta)^{1+\delta}})^{\mathbb{E}[X(C)]}.$$

Concentration Bounds

Theorem

For each edge e, let X_e be an indicator random variable associateed with the event $[e \in T]$, where T is a λ-random tree. Also, for any subset C of the edges of G, define $X(C) = \sum_{e \in C} X_e$. Then we have

$$Pr[X(C) \geq (1+\delta)\mathbb{E}[X(C)]] \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mathbb{E}[X(C)]}.$$

By the negative correlation of the events $[e \in T]$, this follows directly from the result of Panconesi and Srinivasan.

The Thinness Property

α-thin tree

Definition

We say that a tree T is α-thin if for each set $U \subset V$,

$$|T \cap \delta(U)| \leq \alpha \cdot z^*(\delta(U)).$$

Also we say that T is (α, s)-thin if it is α-thin and moreover,

$$c(T) \leq s \cdot \mathrm{OPT}_{\mathrm{HK}}.$$

Sampling Thin Trees

w.r.t. a single cut

Lemma

If T is a spanning tree sampled from distribution $\tilde{p}(.)$ for $\epsilon = 0.2$ in a graph G with $n \geq 5$ vertices then for any set $U \subset V$,

$$Pr[|T \cap \delta(U)| > \beta \cdot z^*(\delta(U))] \leq n^{-2.5z^*(\delta(U))},$$

where $\beta = 4\frac{\log n}{\log \log n}$.

Sampling Thin Trees

w.r.t. a single cut

Lemma

If T is a spanning tree sampled from distribution $\tilde{p}(.)$ for $\epsilon = 0.2$ in a graph G with $n \geq 5$ vertices then for any set $U \subset V$,

$$Pr[|T \cap \delta(U)| > \beta \cdot z^*(\delta(U))] \leq n^{-2.5z^*(\delta(U))},$$

where $\beta = 4\frac{\log n}{\log \log n}$.

The proof follows from the concentration bound with $1 + \delta = \beta\frac{z^*(\delta(U))}{\tilde{z}(\delta(U))} \geq \frac{\beta}{1+\epsilon}$.

Sampling Thin Trees

Theorem

Let $n \geq 5$ and $\epsilon = 0.2$. Let $T_1, \ldots, T_{\lceil 2 \log n \rceil}$ be $\lceil 2 \log n \rceil$ independent samples from $\tilde{p}(.)$. Let T^* be the tree among these samples that minimizes $c(T_j)$. Then T^* is $(4 \log n / \log \log n, 2)$-thin with high probability.

Sampling Thin Trees

Theorem

Let $n \geq 5$ and $\epsilon = 0.2$. Let $T_1, \ldots, T_{\lceil 2 \log n \rceil}$ be $\lceil 2 \log n \rceil$ independent samples from $\tilde{p}(.)$. Let T^* be the tree among these samples that minimizes $c(T_j)$. Then T^* is $(4 \log n / \log \log n, 2)$-thin with high probability.

The proof follows by a union bound over all individual cuts using a result of Karger showing that there are at most n^{2k} cuts of size at most k times the minimum cut value for any half-integer $k \geq 1$.

From a Thin Trees to an Eulerian Walk

Theorem

Assume that we are given an (α, s)-thin spanning tree T^* w.r.t the LP relaxation x^*. Then we can find a Hamiltonian cycle of cost no more than $(2\alpha + s)c(x^*) = (2\alpha + s)OPT_{HK}$ in polynomial time.

Thank You

Questions?

akader@cs.umd.edu