

## Approximation Algorithms

Complexity class review:

P – easy

NP – easy to check “yes”

co-NP – easy to check “no”

NP problems are technically decision problems

- but we can also say some optimization problems are NP hard, if it is NP hard to find out if the optimum is less than some threshold

For NP-hard optimization problems, finding the optimal solution is hard, BUT we may be able to find a sub-optimal solution easily.

- we can have both constant-factor approximations and functions of the input (like  $\log(I)$ ) – N.B. We don't want anything much worse than this)

Some example problems:

Vertex cover

-Approximation: Find a maximal matching in  $G \rightarrow$  factor 2 approximation (any cover has to choose at least 2 endpoints of those edges, this picks both.

Note: We are interested in 3 questions when we find an approximation:

Can the guarantee for this approximation be improved?

$\rightarrow$  find a tight example

Can the same algorithm give a better guarantee?

$\rightarrow$  use linear programming/analysis

$\rightarrow$  find the integrality gap

Is there another lower-bounding scheme which is better?

$\rightarrow$  show finding a better approximation is NP hard

\*\* Many approximations don't prove these things \*\*

Set Cover

-Given a universe  $U$ , a set of subsets, and a cost function, find a min cost subcollection covering  $U$

$\rightarrow$  see slides for greedy algorithm

$\rightarrow$  this is an  $H_n$  factor approximation algorithm (note  $H_n$  diverges  $1+1/2+1/3+\dots$ )

$\rightarrow$  A tight example for this:

$\rightarrow$  singleton sets which each cost 1, and the set of all of them which costs  $1+\epsilon$   
(see slide for image)

$\rightarrow$  there is another algorithm with a better bound

$\rightarrow$  see also: the knapsack problem

$\rightarrow$  greedy doesn't do it for that problem – can perform arbitrarily badly.

$\rightarrow$  a  $\log(n)$  approximation is about all we're willing to tolerate.

$\rightarrow$  note that we can relax some assumptions and do well probabilistically

Steiner Trees

-Have a set of required vertices, and a set of optional steiner vertices  $\rightarrow$  find a tree on a graph including all the required vertices and some of the steiner vertices

$\rightarrow$  We can take a steiner tree and convert to a spanning tree.

## TSP (Traveling Salesman Problem)

- Cannot be approximated within  $\alpha(n)$  unless  $P = NP$
- if we restrict it to a metric graph (think of  $R^2$  or  $R^3$ )
  - then we can get a 2-approximation (see slides for algorithm)
- combining MST with min-cost perfect matching can yield a 1.5-approximation

## Linear Programs and Primal/Dual Problems

- Set cover ILP formation
  - we can relax the problem (let go of integrality) and do rounding
- if  $f$  is the frequency of the most frequent element, and we pick all  $s$  where  $x_s \geq 1/f$ 
  - this is an  $f$ -approximation algorithm.
- we can relax the primal and dual of LP → gives us a primal/dual schema
  - we get slackness variables, and intrinsic lower bounds from slackness.

## **Blum et al., “Approximation Algorithms for Orienteering and Discounted-Reward TSP”**

- consider TSP – this problem is NP-complete
- robot navigation is similar, but
  - it can't go everywhere
    - many problems don't require completeness, but do require some notion of timeliness
- introduce the idea of prizes (at a vertex), quotas (amount of prize desired), penalties (cost of moving between vertices) and budgets (amount of penalty that can be paid).
- For quota-based tsp, there is a 2-approximation
  - quota-tsp is the special case where the prizes are all one
    - k-tsp the quota is  $k$  (this is tied to the min-cost tree – k-MST)
  - all have a version where a budget is given and we're trying to get the maximum reward

## ORIENTEERING

- a tour with maximum reward with length less than some budget  $B$
- unrooted orienteering – can start from anywhere
  - a 2-approximation exists, using PCTSP
  - no previous algorithm for rooted orienteering
- Discounted-Reward TSP
  - have an undirected weighted graph
  - the prize is discounted by  $g_{tv}$  if  $v$  is visited at  $t_v$  the total prize is the sum of the prizes times  $g_{tv}$ 
    - getting to high-reward locations early is good
  - discounting
    - encourages early reward collection
      - good if the environment might change
    - optimal strategy is a mapping from states to actions
    - compare to MDP, which visits states many times, and maximizes the reward in a stochastic setting, where DRTSP visits states once, and is deterministic (no

uncertainty)

- Note that sometimes we can get better bounds in a stochastic setting
- How to compare deterministic/stochastic algorithms?
  - Sometimes deterministic is an upper bound for the stochastic version
    - But not always!
  - Sometimes there are problems with modeling deterministically
  - we can use stochastic algorithms in the stochastic case
- Overall Strategy
  - Approximate the optimum difference between the length of a prize collecting path and the length of the shortest path between endpoints
  - This paper gives
    - an algorithm that gives a constant-factor approximation for the difference
    - a formula for the approximation (this is good – not everyone does)
  - Reduces to another problem
- Path excess
  - distance along P from s-t MINUS the minimum s-t distance
  - min excess path of total prize  $\pi$  is also min-cost path of prize  $\pi$
  - an s-t path approximating optimal excess E by alpha will have length
    - $d(s,t) + \alpha * E \leq \alpha(d(s,t) + E)$
  - Approximate min excess by alpha will also approximate min cost path by alpha
  - Results:

Problem	Approximation Factor	Source
kTSP	$a_{cc}$	Prior work (best is $a_{cc} = 2$ )
Min-excess	$a_{ep} = 3/2a_{cc} + 1$	This paper
Orienteering	$1 + \text{ceiling}(a_{ep})$	This paper
Discounted-reward TSP	$e(a_{ep} + 1)$ (roughly)	This paper

- will show that min-excess solves orienteering, discounted-reward TSP
- main result: showing the factor of min-excess algorithm

+++ SEE SLIDES FOR ALGORITHMS +++

- these have been implemented, you can get the code
  - Time:  $\log(\pi) * V * C_{minexcess} =? \log(\pi) * V^3$

\*Deterministic Discounted Reward is similar to orienteering\*

The reward function is sub-modular, and MDP requires modular rewards, so MDP algorithms don't work for this problem directly

We *could* use an MDP, but it would use exponentially many states.

Open question: could we compress states down (for some problems) and use an MDP?

## Krause and Guestrin, “Near-optimal Observation Selection Using Submodular Functions”

- Motivation
  - need to monitor spatial phenomena
  - this paper is a high-level overview
- Problem
  - we have a lake, and we need to pick locations to monitor in order to get as much information as possible.
    - “information” can be RMSE or entropy
  - We take advantage of submodularity → diminishing returns
    - putting a sensor where there are already many sensors doesn't gain as much as putting it where there aren't many
  - More formally:
    - we have a set  $V = \{1, 2, \dots, n\}$
    - a function  $F: 2^V \rightarrow \mathbb{R}$
    - assume  $F(\text{empty set}) = 0$
    - Assume a black box that can evaluate  $F$  for any input  $A$
  - Submodularity means that if we have 2 sets and  $A$  is a subset of  $B$ , then adding an element to  $A$  gains you more than adding an element to  $B$ .
    - set cover is submodular
  - Ex: mutual information:
    - when adding measurements, how much we gain on existing variables (this is information gain on the variables we care about)
    - mutual information is submodular
  - Submodular functions are closed under linear combinations
  - A greedy algorithm for placing sensors (SEE SLIDE FOR ALGORITHM)
    - intuition: greedily place sensors in best location
      - $F(A_{\text{greedy}}) \geq (1 - 1/e)F(A_{\text{opt}})$  → better than a 2-approximation
      - guarantees best possible approximation unless  $P = NP$
  - Note that this greedy algorithm is empirically quite close to the optimal: why?
    - Lazy greedy algorithm → keep a list from previous iterations and reevaluate  $S_i$  only for the first element, since previous values are an upper bound
    - Lazy is much better than naïve greedy
    - Data-dependent bounds
      - suppose you have a candidate solution and an optimal solution → if you look at what you'd gain by adding others in, you get a data-dependent bound
  - the pSPIEL algorithm
    - padded Sensor Locations at Informative and cost-Efficient Locations
    - have sensors and a difficulty measure for communication along edges
    - For two sets of measurements far enough apart, measurements are independent
      - $F(A_1 \cup A_2) \geq F(A_1) + g \cdot F(A_2)$
    - thus
      - split the problem into regions
      - run the greedy algorithm on the regions

- can get to within a constant factor of both communication cost and objective function
  - the trick → decomposition
    - submodularity and locality
- Padded decomposition scales polynomially in large graphs (good!)
- Note: problems with complicated objective functions are HARD
- when would we have padded decomposition in robotics
  - robot placement (but not a whole lot else, because we want things to move)
- The algorithm is randomized and poly-time
  - can we combine ideas from TSP and this sensor placement stuff and get reasonable results?
    - Maybe, but the graph would have to be metric.