

CSC2420: Algorithm Design, Analysis and Theory

Fall 2022

**An introductory (i.e. foundational) level
graduate course.**

Allan Borodin

October 18, 2022

Week 5

Announcements

- Next week, Calum returns to complete the discussion of online bipartite
- Today the first assignment was due. I had a number of requests for hints. Of course, if I give anyone a substantial hint, I should give that hint to everyone. I think it is good to try to do the problems without additional hints beyond what I indicate in class or on the assignment. But there is no shame in not being able to solve a question and one learns from the attempt. So ... once the assignments are submitted, I will give some further hints and allow everyone the option of handing in a late submission or revising a submission. This additional submission is due two days later (same time).
- A hint for the $2 - \epsilon$ ratio for Graham's greedy makespan in the ROM model.
- A hint for the $1/2$ randomized priority algorithm for general $\{0,1\}$ knapsack.

Today's agenda

- We ended the October 11 class having introduced the set packing and s set packing problems. This is where we will continue.
- More specifically we will see two “natural greedy algorithms for the s -set packing problem and a not so natural greedy algorithm for the set packing problem.
- There may be a nice open problem for the s -set packing and set packing problems which I will mention.
- “Finally” for greedy algorithms, we will discuss the “weighted independent set in a matroid” problem.
- We next consider local search algorithms.
- In particular, we will make a distinction between oblivious and non-oblivious local search. **Note:** I don't like the terminology but that is what is used.

Greedy algorithms for the set packing problem

One of the new areas in theoretical computer science is algorithmic game theory and mechanism design and, in particular, auctions including what are known as *combinatorial auctions* (CAs). The underlying combinatorial problem in CAs is the set packing problem.

The set packing problem

We are given n subsets S_1, \dots, S_n from a universe U of size m . In the weighted case, each subset S_i has a weight w_i . The goal is to choose a disjoint subcollection \mathcal{S} of the subsets so as to maximize $\sum_{S_i \in \mathcal{S}} w_i$. In the s -set packing problem we have $|S_i| \leq s$ for all i .

- This is a well studied problem and by reduction from the max clique problem, there is an $m^{\frac{1}{2}-\epsilon}$ hardness of approximation assuming $NP \neq ZPP$. For s -set packing with constant $s \geq 3$, there is an $\Omega(s/\log s)$ hardness of approximation assuming $P \neq NP$.
- We will consider two “natural” greedy algorithms for the s -set packing problem and a non obvious greedy algorithm for the set packing problem. These greedy algorithms are all fixed order priority,

The first natural greedy algorithm for set packing

Greedy-by-weight ($Greedy_{wt}$)

Sort the sets so that $w_1 \geq w_2 \dots \geq w_n$.

$\mathcal{S} := \emptyset$

For $i : 1 \dots n$

 If S_i does not intersect any set in \mathcal{S} then

$\mathcal{S} := \mathcal{S} \cup S_i$.

End For

- In the unweighted case (i.e. $\forall i, w_i = 1$), this is an online algorithm.
- In the weighted case, greedy-by-weight provides an s -approximation for the s -set packing problem.
- The approximation bound can be shown by a [charging argument](#) where the weight of every set in an optimal solution is charged to the first set in the greedy solution with which it intersects.

The second natural greedy algorithm for set packing

Greedy-by-weight-per-size

Sort the sets so that $w_1/|S_1| \geq w_2/|S_2| \dots \geq w_n/|S_n|$.

$\mathcal{S} := \emptyset$

For $i : 1 \dots n$

 If S_i does not intersect any set in \mathcal{S} then

$\mathcal{S} := \mathcal{S} \cup S_i$.

End For

- In the unweighted case, greedy-by-weight-per-size is known to provide a \sqrt{m} approximation.
- In the weighted case, greedy-by-weight provides an s -approximation for the s -set packing problem.
- For both greedy algorithms, the approximation ratio for the weighted case is tight; that is, there are examples where this is essentially the approximation. These algorithms only provide an m -approximation where $m = |U|$.
- We usually assume $n \gg m$ and note that by just selecting the set of largest weight we obtain an n -approximation. So the goal is to do ^{6/54}

Improving the approximation for set packing

- In the unweighted case, greedy-by-weight-per-size can be restated as sorting so that $|S_1| \leq |S_2| \dots \leq |S_n|$ and it can be shown to provide an \sqrt{m} -approximation for set packing.
- On the other hand, greedy-by-weight-per-size does not improve the m -approximation for weighted set packing.

Greedy-by-weight-per-squareroot-size

Sort the sets so that $w_1/\sqrt{|S_1|} \geq w_2/\sqrt{|S_2|} \dots \geq w_n/\sqrt{|S_n|}$.

$\mathcal{S} := \emptyset$

For $i : 1 \dots n$

 If S_i does not intersect any set in \mathcal{S} then

$\mathcal{S} := \mathcal{S} \cup S_i$.

End For

Theorem: Greedy-by-weight-per-squareroot-size provides a $2\sqrt{m}$ -approximation for the set packing problem. We note that \sqrt{m} is asymptotically the best possible approximation assuming $NP \neq ZPP$.

Another way to obtain an $O(\sqrt{m})$ approximation

There is another way to obtain the same asymptotic improvement for the weighted set packing problem. Namely, we can use the idea of partial enumeration greedy; that is somehow combining some kind of brute force (or naive) approach with a greedy algorithm.

Partial Enumeration with Greedy-by-weight ($PGreedy_k$)

Let Max_k be the best solution possible when restricting solutions to those containing at most k sets. Let G be the solution obtained by $Greedy_{wt}$ applied to sets of cardinality at most $\sqrt{m/k}$. Set $PGreedy_k$ to be the best of Max_k and G .

- Theorem: $PGreedy_k$ achieves a $2\sqrt{m/k}$ -approximation for the weighted set packing problem (on a universe of size m)
- In particular, for $k = 1$, we obtain a $2\sqrt{m}$ approximation and this can be improved by an arbitrary constant factor \sqrt{k} at the cost of the brute force search for the best solution of cardinality k ; that is, at the cost of say n^k .

Combinatorial auctions

For those who are interested, a combinatorial auction is one where n agents (buyers) want one of possibly many subsets of items S_1, S_2, \dots, S_k . Each agent has a private value $w(S_i)$ for any desired set. An auctioneer wants to allocate at most one subset to each agent so as to maximize the total value of disjoint sets that are allocated. This is called the “social welfare” objective.

We assume “free disposal” (as in the display ads problem) so that $w(S) \leq w(T)$ if $S \subset T$.

If every agent was truthful about their valuations, this would be the set packing problem. But agents are self interested and may not truthfully report the values for the sets they desire if they think it will be helpful. That is, they could over state or understate their valuations.

In order to incentivize agents to be truthful, the auctioneer decides on prices for each agent.

Combinatorial auctions continued

If we could compute an optimal solution we would have a way to specify a price for the subset to be allocated to each agent so that the agent would be truthful. But set packing is an NP hard approximation problem as we have indicated. This result (called VCG auctions) does not hold for approximation algorithms.

Perhaps the most prominent question in CAs, is how how much truthfulness can hurt the approximation ratio for a given CA if we insist upon efficient polynomial allocations or insist upon conceptually simple allocations and pricing.

Although we need to be more precise, truthful auctions using priority algorithm allocations for say an s -CA will result in asymptotically the worst possible approximation (i.e., $\Omega(\min\{n,m\})$).

Max weighted independent set in a matroid.

We conclude the current discussion of greedy and priority algorithms by mentioning a classic result about greedy algorithms and matroids.

Let U be a set of elements and \mathcal{I} be a collection of subsets of U . (U, \mathcal{I}) is a matroid if the following hold:

- (Hereditary property) If $I \in \mathcal{I}$ and $I' \subset I$, then $I' \in \mathcal{I}$.
- (Exchange property) If $I', I \in \mathcal{I}$ and $|I'| < |I|$, then $\exists u \in I \setminus I'$ such that $I' \cup \{u\} \in \mathcal{I}$.

An *hereditary set system* (U, \mathcal{I}) is any set system satisfying the hereditary property so that a matroid is an hereditary set system that also satisfies the exchange property.

The sets $I \in \mathcal{I}$ are referred to as the *independent sets*.

The objective of “Max weighted independent set in a matroid” is to select an independent set $I \subseteq U$ so as to maximize $\sum_{u \in I} w(u)$ where $w(u)$ is the weight of element u .

Alternative definitions of a matroid

We note that there are alternative equivalent definitions. In particular, an alternative to the exchange property is that every maximal independent set is called a *basis* and has the same size. This maximum size is called the *rank* of the matroid.

There are alternative statements of the exchange property.

(Basis Exchange property) If $I \neq I^*$ are bases (i.e. maximal independent sets in I), then $\forall x \in I^* \setminus I, \exists y \in I \setminus I^*$ such that $(I^* \setminus \{x\}) \cup \{y\}$ is also a basis.

Matroids and the greedy algorithm

Any acyclic subset of edges in a graph is an independent set in a matroid. If the graph is connected then a maximal acyclic subset is a spanning tree. This is called a graphic matroid.

An independent set of vectors in a vector space is a matroid. ' Whitney [1935] defined this elegant abstraction that applies to many other systems. Generalizations to intersections of matroids and more general independence systems are also known.

The minimum spanning tree (MST) problem can be seen as a maximization problem (negate the signs of the edge weights or subtract each edge weight from the maximum edge weight).

Kruskal's MST algorithm can be seen as the natural greedy algorithm for computing a maximum weight independent set in a graphic matroid. The proof of the optimality of Kruskal's algorithm can be carried over to show how the natural greedy algorithm obtains an optimal solution for computing an optimal independent set **in any matroid**.

The Rado-Edmonds theorem

What is this natural greedy algorithm for computing a max weight independent set in a matroid?

The Rado-Edmonds theorem

What is this natural greedy algorithm for computing a max weight independent set in a matroid?

Standard Greedy for a Maximization Problem

Let $S := \emptyset$

While $\exists u : S \cup \{u\}$ is independent

$S := S \cup \{u\}$ where u is an element of largest weight
such that $S \cup \{u\}$ is independent

End While

Rado: For independence in a matroid M , *Greedy* is an optimal algorithm for maximizing the weight of basis (i.e. maximizing a linear function subject to a matroid constraint).

Remarkably, there is what can be seen as a converse to this fact.

The Rado-Edmonds theorem continued

Edmonds: Consider any hereditary set system. If the natural greedy algorithm produces an optimal solution for every linear function of elements in an independent set, then the set system is a greedy algorithm.

Does this define greedy algorithms?

The Rado-Edmonds theorem continued

Edmonds: Consider any hereditary set system. If the natural greedy algorithm produces an optimal solution for every linear function of elements in an independent set, then the set system is a greedy algorithm.

Does this define greedy algorithms? No. It doesn't say anything about approximation algorithms.

Is the optimality of greedy algorithms limited to independent sets in matroids?

The Rado-Edmonds theorem continued

Edmonds: Consider any hereditary set system. If the natural greedy algorithm produces an optimal solution for every linear function of elements in an independent set, then the set system is a greedy algorithm.

Does this define greedy algorithms? No. It doesn't say anything about approximation algorithms.

Is the optimality of greedy algorithms limited to independent sets in matroids?

No. Consider the greedy algorithm for unweighted interval scheduling on m identical machines.

Local search: the “other” conceptually simplest paradigm

Along with greedy and greedy-like algorithms, *local search* is (for me) one of the two conceptually simplest search/optimization paradigms. Like greedy algorithms, there are many variations of this paradigm.

The vanilla local search paradigm

“Initialize” S

While there is a “better” solution S'
in the “local neighbourhood” $Nbhd(S)$

$S := S'$

End While

If and when the algorithm terminates, the algorithm has computed a *local optimum*.

Local search as a well defined algorithm

To make local search a precise algorithmic model, we have to say:

- 1 How are we allowed to choose an initial solution?
- 2 What constitutes a reasonable definition of a **local neighbourhood**?
- 3 What do we mean by “better”?

Answering these questions (especially as to defining a local neighbourhood) will often be quite problem specific.

Towards a more precise definition for local search

- We clearly want the initial solution to be efficiently computed and to be precise we can (for example) say that the initial solution is a random solution, or a greedy solution or adversarially chosen. Of course, in practice we can use any efficiently computed solution.
- We want the local neighbourhood $Nbhd(S)$ to be such that we can efficiently search for a “better” solution (if one exists).
 - 1 In many problems, a solution S is a subset of the input items or equivalently a $\{0,1\}$ vector, and in this case we often define the $Nbhd(S) = \{S' \mid d_H(S, S') \leq k\}$ for some “small” k where $d_H(S, S')$ is the Hamming distance.
 - 2 More generally whenever a solution is a vector over a small domain D , we can use Hamming distance to define a local neighbourhood. Hamming distance k implies that $Nbhd(S)$ can be searched in at most time $|D|^k$.
 - 3 We can view Ford Fulkerson flow algorithms as local search algorithms where the (possibly exponential size but efficiently search-able) neighbourhood of a flow solution S are flows obtained by adding an **augmenting path** flow.

What does “better” solution mean? Oblivious and non-oblivious local search

- For a search problem, we would generally have a non-feasible initial solution and “better” can then mean “closer” to being feasible.
- For an optimization problem it usually means being an improved solution which respect to the given objective. For reasons I cannot understand, this has been termed *oblivious local search*. I think it should be called greedy local search.
- For some applications, it turns out that rather than searching to improve the given objective function, we search for a solution in the local neighbourhood that improves a related *potential function* and this has been termed *non-oblivious local search*.
- In searching for an improved solution, we may want an arbitrary improved solution, a random improved solution, or the best improved solution in the local neighbourhood.
- For efficiency we sometimes insist that there is a “sufficiently better” improvement rather than just better.

Exact Max- k -Sat

- **Given:** An exact k -CNF formula

$$F = C_1 \wedge C_2 \wedge \dots \wedge C_m,$$

where $C_i = (\ell_i^1 \vee \ell_i^2 \dots \vee \ell_i^k)$ and $\ell_i^j \in \{x_k, \bar{x}_k \mid 1 \leq k \leq n\}$.

- In the **weighted** version, each C_i has a weight w_i .
- **Goal:** Find a truth assignment τ so as to maximize

$$W(\tau) = w(F \mid \tau),$$

the weighted sum of satisfied clauses w.r.t the truth assignment τ .

- It is NP hard to achieve an approximation better than $\frac{7}{8}$ for exact Max-3-Sat and hence that hard for the non exact version of Max- k -Sat for $k \geq 3$.
- Max-2-Sat can be approximated to within a factor $\approx .87856$.

The natural oblivious local search

- A natural oblivious local search algorithm uses a Hamming distance d neighbourhood:

$$N_d(\tau) = \{ \tau' : \tau \text{ and } \tau' \text{ differ on at most } d \text{ variables} \}$$

Oblivious local search for Exact Max- k -Sat

Choose any initial truth assignment τ

WHILE there exists $\hat{\tau} \in N_d(\tau)$ such that $W(\hat{\tau}) > W(\tau)$

$\tau := \hat{\tau}$

END WHILE

How good is this oblivious local search algorithm?

- Note: Following the standard convention for Max-Sat, I am using approximation ratios < 1 .
- It can be shown that for $d = 1$, the approximation ratio for Exact-Max-2-Sat is $\frac{2}{3}$.
- In fact, for every exact 2-Sat formula, the algorithm finds an assignment τ such that $W(\tau) \geq \frac{2}{3} \sum_{i=1}^m w_i$, the weight of all clauses, and we say that the “totality ratio” is at least $\frac{2}{3}$.
- More generally for Exact Max- k -Sat the ratio is $\frac{k}{k+1}$. This ratio is essentially a tight ratio for any $d = o(n)$.
- This is in contrast to an online greedy algorithm derived from a naive randomized algorithm that achieves totality ratio $(2^k - 1)/2^k$.
- “In practice”, the local search algorithm often performs better than the naive greedy and one could always start with the greedy algorithm and then apply local search.

Analysis of the oblivious local search for Exact Max-2-Sat

- Let τ be a local optimum and let
 - ▶ S_0 be those clauses that are not satisfied by τ
 - ▶ S_1 be those clauses that are satisfied by exactly one literal by τ
 - ▶ S_2 be those clauses that are satisfied by two literals by τ

- Let $W(S_i)$ be the corresponding weight.

Analysis of oblivious Exact-Max-2-Sat local search continued

- We will say that a clause involves a variable x_j if either x_j or \bar{x}_j occurs in the clause. Then for each j , let
- A_j be those clauses in S_0 involving the variable x_j .
- B_j be those clauses C in S_1 involving the variable x_j such that it is the literal x_j or \bar{x}_j that is satisfied in C by τ .
- C_j be those clauses in S_2 involving the variable x_j .
- Let $W(A_j), W(B_j), W(C_j)$ be the corresponding weights.

Analysis of the oblivious local search (continued)

- Summing over all variables x_j , we get
- $2W(S_0) = \sum_j W(A_j)$ noting that each clause in S_0 gets counted twice.
- $W(S_1) = \sum_j W(B_j)$
- Given that τ is a local optimum, for every j , we have

$$W(A_j) \leq W(B_j)$$

or else flipping the truth value of x_j would improve the weight of the clauses being satisfied.

- Hence (by summing over all j),

$$2W_0 \leq W_1.$$

Finishing the analysis

- It follows then that the ratio of clause weights not satisfied to the sum of all clause weights is

$$\frac{W(S_0)}{W(S_0)+W(S_1)+W(S_2)} \leq \frac{W(S_0)}{3W(S_0)+W(S_2)} \leq \frac{W(S_0)}{3W(S_0)}$$

- It is not easy to verify but there are examples showing that this $\frac{2}{3}$ bound is essentially tight for any N_d neighbourhood for $d = o(n)$.
- It is also claimed that the bound is at best $\frac{4}{5}$ whenever $d < n/2$. For $d = n/2$, the algorithm would be optimal.
- In the weighted case, we have to worry about the number of iterations. And here we can speed up the termination by insisting that any improvement has to be sufficiently better.

Using the proof to improve the algorithm

Aside: Using adversarial examples and viewing algorithms as a *game* against an adversary is an idea that is now very active in “adversarial learning”.

- We can learn something from this proof to improve the performance.
- Note that we are not using anything about $W(S_2)$.
- If we could guarantee that $W(S_0)$ was at most $W(S_2)$ then the ratio of clause weights not satisfied to all clause weights would be $\frac{1}{4}$.
- **Claim:** We can do this by enlarging the neighbourhood to include $\tau' = \text{the complement of } \tau$.

The non-oblivious local search

- We consider the idea that satisfied clauses in S_2 are more valuable than satisfied clauses in S_1 (because they are able to withstand any single variable change).
- The idea then is to weight S_2 clauses more heavily.
- Specifically, in each iteration we attempt to find a $\tau' \in N_1(\tau)$ that improves the **potential function**

$$\frac{3}{2}W(S_1) + 2W(S_2)$$

instead of the oblivious $W(S_1) + W(S_2)$.

- More generally, for all k , there is a setting of scaling coefficients c_1, \dots, c_k , such that the non-oblivious local search using the potential function $c_1W(S_1) + c_2W(S_2) + \dots + c_kW(S_k)$ results in approximation ratio $\frac{2^k-1}{2^k}$ for exact Max- k -Sat.

Sketch of $\frac{3}{4}$ totality bound for the non oblivious local search for Exact Max-2-Sat

- Renaming variables, we can assume that τ is the all true assignment.
- Let $P_{i,j}$ be the weight of all clauses in S_i containing x_j .
- Let $N_{i,j}$ be the weight of all clauses in S_i containing \bar{x}_j .
- Here is the key observation for a local optimum τ wrt the stated potential:

.in

$$-\frac{1}{2}P_{2,j} - \frac{3}{2}P_{1,j} + \frac{1}{2}N_{1,j} + \frac{3}{2}N_{0,j} \leq 0$$

- Summing over variables $P_1 = N_1 = W(S_1)$, $P_2 = 2W(S_2)$ and $N_0 = 2W(S_0)$ and using the above inequality we obtain $3W(S_0) \leq W(S_1) + W(S_2)$

Some experimental results concerning Max-Sat

- Of course, one wonders whether or not a worse case approximation will actually have a benefit in “practice”.
- “In practice”, local search becomes more of a “heuristic” where one uses various approaches to escape (in a principled way) local optima and then continuing the local search procedure. Perhaps the two most commonly used versions are [Tabu Search](#) and [Simulated Annealing](#).
- Later, we will also discuss methods based on online algorithm and “random walks” and other randomized methods (and their derandomizations).
- We view these algorithmic ideas as starting points.
- But for what it is worth, here are some 2010 experimental results both for artificially constructed instances and well as for one of the many benchmark test sets for Max-Sat.
- Recent experimental results by Poloczek and Williamson show that various ways to use greedy and local search algorithms can compete (wrt. various test sets) with “state of the art” simulated annealing algorithms and walk-sat algorithms while using much less time.

Experiment for unweighted Max-3-Sat




Fig. 1. Average performance when executing on random instances of exact MAX-3-SAT.

[From Pankratov and Borodin]

Experiment for Benchmark Max-Sat

	OLS	NOLS	TS	NOLS+TS	SA	MWS
OLS	0	457	741	744	730	567
NOLS	160	0	720	750	705	504
TS	0	21	0	246	316	205
NOLS+TS	8	0	152	0	259	179
SA	30	50	189	219	0	185
MWS	205	261	453	478	455	0

Table 2. MAX-SAT 2007 benchmark results. Total number of instances is 815. The tallies in the table show for how many instances a technique from the column improves over the corresponding technique from the row.

[From Pankratov and Borodin]

More experiments for benchmark Max-Sat

Table 2. The Performance of Local Search Methods

	NOLS+TS		2Pass+NOLS		SA		WalkSat	
	% sat	∅ time	% sat	∅ time	% sat	∅ time	% sat	∅ time
SC-APP	90.53	93.59s	99.54	45.14s	99.77	104.88s	96.50	2.16s
MS-APP	83.60	120.14s	98.24	82.68s	99.39	120.36s	89.90	0.48s
SC-CRAFTED	92.56	61.07s	99.07	22.65s	99.72	70.07s	98.37	0.66s
MS-CRAFTED	84.18	0.65s	83.47	0.01s	85.12	0.47s	82.56	0.06s
SC-RANDOM	97.68	41.51s	99.25	40.68s	99.81	52.14s	98.77	0.94s
MS-RANDOM	88.24	0.49s	88.18	0.00s	88.96	0.02s	87.35	0.06s

Figure: Table from Poloczek and Williamson 2017

Note: *2Pass* is a deterministic “2-pass online algorithm” that is derived from a randomized online algorithm that we will discuss soon.

Oblivious and non-oblivious local search for $k + 1$ claw free graphs

- We again consider the **weighted max (independent) vertex set** in a $k + 1$ claw free graph. (Recall the argument generalizing the approximation ratio for the k set packing problem.)
- The standard greedy algorithm and the 1-swap oblivious local search both achieve a $\frac{1}{k}$ approximation for the WMIS in $k + 1$ claw free graphs. Here we define an “ ℓ -swap” oblivious local search by using neighborhoods defined by bringing in a set S of up to ℓ vertices and removing all vertices adjacent to S .

NOTE: We will continue to use fractional approximation ratios for the maximization problems being considered this week.

- For the **unweighted MIS**, Halldórsson shows that a 2-swap oblivious local search will yield a $\frac{2}{k+1}$ approximation.

Berman's [2000] non-oblivious local search

- For the **weighted MIS**, the “ ℓ -swap” oblivious local search results (essentially) in an $\frac{1}{k}$ locality gap for any constant ℓ .
- Chandra and Halldósson [1999] show that by first using a standard greedy algorithm to initialize a solution and then using a “greedy” k -swap oblivious local search, the approximation ratio improves to $\frac{3}{2k}$.
- Can we use non-oblivious local search to improve the locality gap? Once again given two solutions V_1 and V_2 having the same weight, when is one better than the other?

Berman's [2000] non-oblivious local search

- For the **weighted MIS**, the “ ℓ -swap” oblivious local search results (essentially) in an $\frac{1}{k}$ locality gap for any constant ℓ .
- Chandra and Halldósson [1999] show that by first using a standard greedy algorithm to initialize a solution and then using a “greedy” k -swap oblivious local search, the approximation ratio improves to $\frac{3}{2k}$.
- Can we use non-oblivious local search to improve the locality gap? Once again given two solutions V_1 and V_2 having the same weight, when is one better than the other?
- Intuitively, if one vertex set V_1 is small but vertices in V_1 have large weights that is better than a solution with many small weight vertices.

Berman's [2000] non-oblivious local search

- For the **weighted MIS**, the “ ℓ -swap” oblivious local search results (essentially) in an $\frac{1}{k}$ locality gap for any constant ℓ .
- Chandra and Halldósson [1999] show that by first using a standard greedy algorithm to initialize a solution and then using a “greedy” k -swap oblivious local search, the approximation ratio improves to $\frac{3}{2k}$.
- Can we use non-oblivious local search to improve the locality gap? Once again given two solutions V_1 and V_2 having the same weight, when is one better than the other?
- Intuitively, if one vertex set V_1 is small but vertices in V_1 have large weights that is better than a solution with many small weight vertices.
- Berman chooses the potential function $g(S) = \sum_{v \in S} w(v)^2$. Ignoring some small ϵ 's, his k -swap non-oblivious local search achieves a locality gap of $\frac{2}{k+1}$ for WMIS on $k + 1$ claw-free graphs.

The (metric) facility location and k -median problems

- Two extensively studied problems in operations research and CS algorithm design are the related uncapacitated facility location problem (UFL) and the k -median problem. In what follows we restrict attention to the (usual) metric case of these problems defined as follows:

Definition of UFL

Input: (F, C, d, f) where F is a set of facilities, C is a set of clients or cities, d is a metric distance function over $F \cup C$, and f is an opening cost function for facilities.

Output: A subset of facilities F' minimizing $\sum_{i \in F'} f_i + \sum_{j \in C} d(j, F')$ where f_i is the opening cost of facility i and $d(j, F') = \min_{i \in F'} d(j, i)$.

- In the capacitated version, facilities have capacities and cities can have demands (rather than unit demand). The constraint is that a facility can not have more assigned demand than its capacity so it is not possible to always assign a city to its closest facility.

UFL and k -median problems continued

Definition of k -median problem

Input: (F, C, d, k) where F, C, d are as in UFL and k is the number of facilities that can be opened.

Output: A subset of facilities F' with $|F'| = k$ minimizing $\sum_{j \in C} d(j, F')$

- These problems are clearly well motivated. Moreover, they have been the impetus for the development of many new algorithmic ideas which we will hopefully at least touch upon throughout the course.
- There are many variants of these problems and in many papers the problems are defined so that $F = C$; that is, any city can be a facility. If a solution can be found when F and C are disjoint then there is a solution for the case of $F = C$.

UFL and k -median problems continued

- It is known (Guha and Khuller) that UFL is hard to approximate to within a factor better than 1.463 assuming NP is not a subset of $DTIME(n^{\log \log n})$ and the k -median problem is hard to approximate to within a factor better than $1 + 1/e \approx 1.736$ (Jain, Mahdian, Saberi).
- The UFL problem is better understood than k -median. After a long sequence of improved approximation results the current best polynomial time approximation is 1.488 (Li, 2011).
- For k -median, until recently, the best approximation was by a local search algorithm. Using a p -flip (of facilities) neighbourhood, Arya et al (2001) obtain a $3 + 2/p$ approximation which yields a $3 + \epsilon$ approximation running in time $O(n^{2/\epsilon})$.
- Li and Svensson (2013) have obtained a $(1 + \sqrt{3} + \epsilon)$ approximation running in time $O(n^{1/\epsilon^2})$. Surprisingly, they show that an α approximate “pseudo solution” using $k + c$ facilities can be converted to an $\alpha + \epsilon$ approximate solution running in $n^{O(c/\epsilon)}$ times the complexity of the pseudo solution.

Some additional comments on local search

- An interesting (but probably difficult) open problem is to use a non oblivious local search for either the UFL or k -median problems.
- But suffice it to say now that local search is the basis for many practical algorithms, especially when the idea is extended by allowing some well motivated ways to escape local optima (e.g. simulated annealing, tabu search) and combined with other paradigms.
- Although local search with all its variants is viewed as a great “practical” approach for many problems, local search is not often theoretically analyzed. It is not surprising then that there hasn’t been much interest in formalizing the method and establishing limits.
- LP is itself often solved by some variant of the simplex method, which can also be thought of as a local search algorithm, moving from one vertex of the LP polytope to an adjacent vertex.
 - ▶ No such method is known to run in polynomial time in the worst case.

Submodular functions

Let U be a universe. In what follows, we will only be interested in set functions that satisfy $f(S) \geq 0$ for all $S \subseteq U$. We will also assume functions are *normalized* in that $F(\emptyset) = 0$. These assumptions are not that essential but they are standard and without these assumptions statements and proofs become somewhat more complex.

- A *sublinear set function* satisfies the property that $f(S \cup T) \leq f(S) + f(T)$ for all subsets S, T of U .
- When $f(S \cup T) + f(S \cap T) = f(S) + f(T)$, the function is a linear (also called modular) function.
- A *submodular set function* $f : U \rightarrow \mathbb{R}$ satisfies the following property: $f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$
- It follows that modular set functions are submodular and submodular functions are sublinear.
- Submodular functions can be monotone or non-monotone. A monotone submodular function also satisfies the property that $f(S) \leq f(T)$ whenever $S \subseteq T$.

An alternative characterization and examples of submodular functions

Submodular functions satisfy and can also be defined as those satisfying a *decreasing marginal gains* property. Namely, For $S \subset T$, $f(T \cup \{x\}) - f(T) \leq f(S \cup \{x\}) - f(S)$. That is, adding additional elements has decreasing (more precisely, non increasing) marginal gain for larger sets.

Most applications of submodular functions are for monotone submodular functions. For example, in practice, when we are obtaining results from a search engine, as we obtain more and more results, we tend to obtain less additional value.

Modular functions are monotone.

The rank function of a matroid is a monotone submodular function.

The two most common examples of non-monotone submodular functions are max-cut and max-di-cut (i.e., max directed cut)

Monotone submodular function maximization

- The monotone problem is only interesting when the submodular maximization is subject to some constraint.
- Probably the simplest and most widely used constraint is a cardinality constraint; namely, to maximize $f(S)$ subject to $|S| \leq k$ for some k and since f is monotone this is the same as the constraint $f(S) = k$.
- Following Cornuéjols, Fisher and Nemhauser [1977] (who study a specific submodular function), Nemhauser, Wolsey and Fisher [1978] show that the standard greedy algorithm achieves a $1 - \frac{1}{e}$ approximation for the cardinality constrained monotone problem. More precisely, for all k , the standard greedy is a $1 - (1 - \frac{1}{k})^k$ approximation for a cardinality k constraint.

Standard greedy for submodular functions wrt cardinality constraint

$S := \emptyset$

While $|S| < k$

Let u maximize $f(S \cup \{u\}) - f(S)$

$S := S \cup \{u\}$

End While

Proof: greedy approx for monotone submodular maximization subject to cardinality constraint

We want to prove the $1 - (1 - \frac{1}{k})^k$ approximation bound.

Let S_i be the set after i iterations of the standard greedy algorithm and let $S^* = \{x_1, \dots, x_k\}$ be an optimal set so that $OPT = f(S^*)$. For any set S and element x , let $f_S(x) = f(S \cup \{x\}) - f(S)$ be the marginal gain by adding x to S . The proof uses the following sequence of inequalities:

$$\begin{aligned} f(S^*) &\leq f(S_i \cup S^*) \text{ by monotonicity} \\ &\leq f(S_i) + (f(S_i \cup \{x_1\}) - f(S_i)) + (f(S_i \cup \{x_1, x_2\}) - f(S_i \cup \{x_1\})) + \dots \\ &\quad \text{(by submodularity; question 5(a) on assignment)} \\ &\leq f(S_i) + f_{S_i}(x_1) + f_{S_i}(x_2) + \dots + f_{S_i}(x_k) \\ &\quad \text{(again by submodularity)} \\ &\leq f(S_i) + k \cdot (f(S_{i+1}) - f(S_i)) \text{ by the greedy assumption} \end{aligned}$$

Equivalently, $f(S_{i+1}) \geq f(S_i) + \frac{1}{k}(f(OPT) - f(S_i))$

The proof is completed by showing $f(S_i) \geq (1 - (1 - \frac{1}{k})^i) \cdot OPT$ by induction on i .

Generalizing to a matroid constraint

- Nemhauser and Wolsey [1978] showed that the $1 - \frac{1}{e}$ approximation is optimal in the sense that an exponential number of value oracle queries would be needed to beat the bound for the cardinality constraint.
- Furthermore, Feige [1998] shows it is NP hard to beat this bound even for the explicitly represented maximum k -coverage problem.
- Following their first paper, Fisher, Nemhauser and Wolsey [1978] extended the cardinality constraint to a **matroid** constraint.
- Fisher, Nemhauser and Wolsey show that both the standard greedy algorithm and a 1-exchange local search algorithm (that will follow) achieve a $\frac{1}{2}$ approximation for maximizing a monotone submodular function subject to an arbitrary matroid constraint.
- They also showed that this bound was tight for the greedy and 1-exchange local search algorithms.

Monotone submodular maximization subject to a matroid constraint

We need some additional facts about matroids and submodular functions.

- Brualdi [1969] Let O and S be two independent sets in a matroid of the same size (in particular they could be two bases). Then there is a bijection π between $O \setminus S$ and $S \setminus O$ such that for all $x \in O$, $(S \setminus \{\pi(x)\}) \cup x$ is independent.
- We have the following facts for a submodular function f on a ground set U :
 - 1 Let $C = \{c_1, \dots, c_\ell\} \subseteq U \setminus S$. Then

$$\sum_{i=1}^{\ell} [f(S + c_i) - f(S)] \geq f(S \cup C) - f(S)$$

- 2 Let $\{t_1, \dots, t_\ell\}$ be elements of S . Then

$$\sum_{i=1}^{\ell} [f(S) - f(S \setminus \{t_i\})] \leq f(S)$$

The 1-exchange local search algorithm

We can start with any basis S (eg using the natural greedy algorithm). Then we keep trying to find an element of $x \notin S$ such that $(S \setminus \{\pi(x)\}) \cup \{x\} > f(S)$. Here π is the bijection as in Brualdi's result.

The previous local search algorithm provides a $\frac{1}{2}$ -approximation for maximizing a monotone submodular function.

Now let S be a local optimum and O an optimal solution. By local optimality, for all $x \in O \setminus S$, we have

$$f(S) \geq f((S \setminus \{\pi(x)\}) \cup \{x\})$$

Subtracting $f(S \setminus \{\pi(x)\})$ from both sides, we have

$$f(S) - f(S \setminus \{\pi(x)\}) \geq f((S \setminus \{\pi(x)\}) \cup \{x\}) - f(S \setminus \{\pi(x)\})$$

From submodularity,

$$f((S \setminus \{\pi(x)\}) \cup \{x\}) - f(S \setminus \{\pi(x)\}) \geq f(S \cup \{x\}) - f(S)$$

Thus for all $x \in O \setminus S$

$$f((S \setminus \{\pi(x)\}) \cup \{x\}) \geq f(S \cup \{x\}) - f(S)$$

Completing the local search approximation

Summing over all such x yields

$$\sum_{x \in O \setminus S} [f(S) - f(S \setminus \{\pi(x)\})] \geq \sum_{x \in O \setminus S} [f(S \cup \{x\}) - f(S)]$$

Applying the first fact on slide 28 to the right hand side of this inequality and the second fact to the left hand side, we get

$$f(S) \geq f(S \cup (O \setminus S)) - f(S) = f(O \cup S) - f(S) \geq f(O) - f(S)$$

which gives the desired $\frac{1}{2}$ -approximation.