1 "Cut-free" Proof

Problem 1 (b-free knapsack) *Consider a set of elements E and two weights $w : E \to \mathbb{Z}_+$ and $c : E \to \mathbb{Z}_+$ and a budget $b \in \mathbb{Z}_+$. Given a feasible set $\mathcal{F} \subset 2^E$, find $\min_{X\in\mathcal{F},F\subset E} w(X \setminus F)$ such that $c(F) \le b$.*

Always remember that $\mathcal{F}$ is usually not explicitly given.

Problem 2 (Normalized knapsack) *Given the same input as Problem 1, find $\min_{X\in\mathcal{F},F\subset E} \frac{w(X\setminus F)}{B-c(F)}$ such that $c(F) \le b$.*

In [4] the normalized min-cut problem use $B = b + 1$. Here we use any integer $B > b$ and see how their method works.

Denote by τ the optimum of Problem 2. Define a new weight $w_\tau : E \to \mathbb{R}$,

$$w_\tau(e) = \begin{cases} w(e) & \text{if } w(e) < \tau \cdot c(e) \text{ (light elem)} \\ \tau \cdot c(e) & \text{otherwise (heavy elem)} \end{cases}$$

Lemma 1.1 *Let (X^N, F^N) be the optimal solution to Problem 2. Every element in F^N is heavy.*

The proof is exactly the same as [4, Lemma 1].

The following two lemmas show (a general version of) that the optimal cut C^N to normalized min-cut is exactly the minimum cut under weights w_τ.

Lemma 1.2 *For any $X \in \mathcal{F}$, $w_\tau(X) \ge \tau B$.*

Lemma 1.3 $X^N \in \arg\min_{X\in\mathcal{F}} w_\tau(X)$.

Proof:

$$\begin{aligned} w_\tau(X^N) &\le w(X^N \setminus F^N) + w_\tau(F^N) \\ &= \tau \cdot (B - c(F^N)) + \tau \cdot c(F^N) \\ &= \tau B \end{aligned}$$

Thus by Lemma 1.2, X^N gets the minimum. □

Now we show the counter part of [4, Theorem 5], which states the optimal solution to Problem 1 is a α-approximate solution to $\min_{F\in\mathcal{F}} w_\tau(F)$.

Lemma 1.4 (Lemma 4 in [4]) *Let (X^*, F^*) be the optimal solution to Problem 1. X^* is either an α-approximate solution to $\min_{X\in\mathcal{F}} w_\tau(X)$ for some $\alpha > 1$, or $w(X^* \setminus F^*) \ge \tau(\alpha B - b)$.*

Then following the argument of Corollary 1 in [4], assume that X^* is not an α-approximate solution to $\min_{X\in\mathcal{F}} w_\tau(X)$ for some $\alpha > 1$. We have

$$\frac{w(C^N \setminus F^N)}{w(C^* \setminus F^*)} \le \frac{\tau(B - c(F^N))}{\tau(\alpha B - b)} \le \frac{B}{\alpha B - b},$$

where the second inequality uses Lemma 1.4. One can see that if $\alpha > 2$, $\frac{w(C^N\setminus F^N)}{w(C^*\setminus F^*)} \le \frac{B}{\alpha B-b} < 1$ which implies (C^*, F^*) is not optimal. Thus for $\alpha > 2$, X^* must be a 2-approximate solution to $\min_{X\in\mathcal{F}} w_\tau(X)$.

Finally we get a knapsack version of Theorem 4:

Theorem 1.5 (Theorem 4 in [4]) *Let $X^{\min}$ be the optimal solution to $\min_{X\in\mathcal{F}} w_\tau(X)$. The optimal set X^* in Problem 1 is a 2-approximation to $X^{\min}$.*

Thus to obtain a FPTAS for Problem 1, one need to design a FPTAS for Problem 2 and a polynomial time alg for finding all 2-approximations to $\min_{X\in\mathcal{F}} w_\tau(X)$.

FPTAS for Problem 2 in [4] (The name "FPTAS" here is not precise since we do not have a approximation scheme but an enumeration algorithm. But I will use this term anyway.) In their settings, $\mathcal{F}$ is the collection of all cuts in some graph. Let OPT^N be the optimum of Problem 2. We can assume that there is no $X \in \mathcal{F}$ s.t. $c(X) \le b$ since this is polynomially detectable (through min-cut on $c(\cdot)$) and the optimum is 0. Thus we have $\frac{1}{b+1} \le \mathrm{OPT}^N \le |E| \cdot \max_e w(e)$. Then we enumerate $\frac{(1+\varepsilon)^i}{b+1}$ where $i \in \left\{0, 1, \dots, \left\lfloor \log_{1+\varepsilon}(|E| w_{\max}(b+1)) \right\rfloor\right\}$. There is a feasible i s.t. $(1-\varepsilon)\,\mathrm{OPT}^N \le \frac{(1+\varepsilon)^i}{b+1} \le \mathrm{OPT}^N$ since $\frac{(1+\varepsilon)^i}{b+1} \le \mathrm{OPT}^N \le \frac{(1+\varepsilon)^{i+1}}{b+1}$ holds for some i.

Note that this enumeration scheme also holds for arbitrary $\mathcal{F}$ if we have a non-zero lowerbound on OPT^N.

Conjecture 1.6 *Let (C, F) be the optimal solution to connectivity interdiction. The optimum cut C can be computed in polynomial time. In other words, connectivity interdiction is almost as easy as knapsack.*

2 Connections

For unit costs, connectivity interdiction with budget $b = k-1$ is the same problem as finding the minimum weighted edge set whose removal breaks k-edge connectivity.

It turns out that Problem 2 is just a necessary ingredient for MWU. Authors of [4] $\subset$ authors of [1].

How to derive normalized min cut for connectivity interdiction?

$$
\begin{aligned}
\max \quad & z & & \\
\text{s.t.} \quad & \sum_e y_e c(e) \le B & & \text{(budget for } F\text{)} \\
& \sum_{e \in T} x_e \ge 1 & \forall T \quad & (x \text{ is a cut}) \\
\sum_e \min(0, x_e - y_e) w(e) \ge z & & & \\
& y_e, x_e \in \{0,1\} & \forall e &
\end{aligned}
$$

we can assume that $y_e \le x_e$.

$$
\begin{aligned}
\min \quad & \sum_e (x_e - y_e) w(e) & & \\
\text{s.t.} \quad & \sum_{e \in T} x_e \ge 1 & \forall T \quad & (x \text{ is a cut}) \\
& \sum_e y_e c(e) \le B & & \text{(budget for } F\text{)} \\
& x_e \ge y_e & \forall e \quad & (F \subset C) \\
& y_e, x_e \in \{0,1\} & \forall e &
\end{aligned}
$$

Now this LP looks similar to the normalized min-cut problem.
A further reformulation (the new x is $x - y$) gives us the following,

$$
\begin{aligned}
\min \quad & \sum_e x_e w(e) \\
s.t. \quad & \sum_{e \in T} x_e + y_e \geq 1 \qquad \forall T \quad (x + y \text{ is a cut}) \\
& \sum_e y_e c(e) \leq b \qquad (\text{budget for } F) \\
& y_e, x_e \in \{0, 1\} \quad \forall e
\end{aligned} \tag{1}
$$

Note that now this is almost a positive covering LP. Let $L(\lambda) = \min\{w(C \setminus F) - \lambda(b - c(F))|\forall \text{cut } C \ \forall F \subset C\}$. Consider the Lagrangian dual,

$$\max_{\lambda \geq 0} L(\lambda) = \max_{\lambda \geq 0} \min\{w(C \setminus F) - \lambda(b - c(F))|\forall \text{cut } C \ \forall F \subset C\}$$

We have shown that the budget B in normalized min-cut does not really matter as long as $B > b$. Note that $L(\lambda)$ and the normalized min-cut look similar to the principal sequence of partitions of a graph and the graph strength problem.

2.1 graph strength

Given a graph $G = (V, E)$ and a cost function $c : V \to \mathbb{Z}_+$, the strength $\sigma(G)$ is defined as $\sigma(G) = \min_\Pi \frac{c(\delta(\Pi))}{|\Pi|-1}$, where Π is any partition of V, $|\Pi|$ is the number of parts in the partition and $\delta(\Pi)$ is the set of edges between parts. Note that an alternative formulation of strength (using graphic matroid rank) is $\sigma(G) = \min_{F \subset E} \frac{|E-F|}{r(E)-r(F)}$, which in general is the fractional optimum of matroid base packing.

The principal sequence of partitions of G is a piecewise linear concave curve $L(\lambda) = \min_\Pi c(\delta(\Pi)) - \lambda|\Pi|$. ($L(\lambda) = \min_{F \in E} c(E \setminus F) - \lambda(r(E) - r(F) + 1)$) Cunningham used principal partition to computed graph strength [3]. There is a list of good properties mentioned in [2, Section 6](implicated stated in [3]).

1. We can assume G is connected and deal with the smallest strength component. One can see this by fractional base packing on the direct sum of matroids. Note that on disconnected graphs we should use the edge set definition instead of partitions.

2. $L(\lambda)$ is piecewise linear concave since it is the lower envelope of some line arrangement.

3. For each line segment on $L(\lambda)$ there is a corresponding partition Π. If λ^* is a breakpoint on $L(\lambda)$, then there are two optimal solution (say partitions P_1 and P_2, assume $|P_1| \leq |P_2|$) to $\min_\Pi c(\delta(\Pi)) - \lambda^*|\Pi|$. Then P_2 is a refinement of P_1.

   Proof (sketch): Suppose that P_2 is not a refinement of P_1. We claim that the meet of P_1 and P_2 achieves a objective value at least no larger than P_1 or P_2 does. The correspondence between graphic matroid rank function and partitions of V gives us a reformulation $L(\lambda^*) = \min_{F \subset E} c(E - F) - \lambda^*(r(E) - r(F) + 1)$. Here F is the set of edges in each part of Π. Let $g(F) = c(E - F) + \lambda^* r(F) - \lambda^* n$. Then the claim is equivalent to the fact that for two optimal solutions F_1, F_2 to $L(\lambda^*)$, $g(F_1 \cap F_2) \leq g(F_1) = g(F_2) \leq g(F_1 \cup F_2)$, which can be seen by the submodularity of g and the optimality of F_1, F_2. □

   The number of breakpoints on $L(\lambda)$ is at most $n - 1$.

4. Let λ^* be a breakpoint on $L(\lambda)$ induced by edge set F. The next breakpoint is induced by the edge set $F' \subset F$ and F' is the solution to strength problem on the smallest strength component of F. λ^* is the strength of the smallest strength component in F. These claims can be seen by the following arguments. From the previous bullet we have $\min_{\Delta F} c(E - F + \Delta F) - \lambda^*(r(E) - r(F - \Delta F) + 1) = L(\lambda^*)$. Consider the largest λ^* which allows $\Delta F = \emptyset$ to be an optimal solution. Such λ^* would be the next breakpoint. For any ΔF, $c(E - F + \Delta F) - \lambda^*(r(E) - r(F - \Delta F) + 1) \geq c(E - F) - \lambda^*(r(E) - r(F) + 1)$. Thus we have $\lambda^* \leq \frac{c(\Delta F)}{r(F) - r(F - \Delta F)}$.

5. Consider $\lambda \in [0, \epsilon]$ for some small enough ϵ. The Lagrangian dual $\min_F c(E \setminus F) - \lambda(r(E) - r(F) + 1)$ gets the optimum at $F = E$. That is $c(E \setminus F') - \lambda(r(E) - r(F') + 1) > -\lambda$ for all $F' \subsetneq E$. We are interested in the upperbound ϵ of λ such that the optimal F is a proper subset of E when $\lambda > \epsilon$. Therefore, the upperbound is $\epsilon = \min_{F \subsetneq E} \frac{c(E \setminus F)}{r(E) - r(F)}$, which is exactly the strength.

2.2 principal sequence of partitions for cut interdiction

Now we focus on $L(\lambda) = \min\{w(C \setminus F) - \lambda(b - c(F)) | \forall \text{cut } C \; \forall F \subset C\}$. We can still assume that G is connected and see that $L(\lambda)$ is pwl concave (1 and 2 still hold). Let λ^* be a breakpoint on L. Suppose that there are two optimal solutions (C_1, F_1) and (C_2, F_2) at λ^*. For fixed C ($C_1 = C_2$), the same argument for principal partition still works. However, the difficult part is that C might not be the same. So it's unlikely that 3 and 4 hold. For cut interdiction problem, 5 shows connections between normalized mincut and the original interdiction problem. Recall that we observe the denominator in normalized min-cut can be relaxed (that is, we can use $\frac{w(C \setminus F)}{B - c(F)}$ for any $B > b$, instead of restricting to $B = b + 1$) and the analysis still works. Now following the previous argument for 5, we assume $\lambda \in [0, \epsilon]$ for small enough positive ϵ. For any C, we have $F = C$ since $w(C \setminus F)$ is dominating. For the remaining term $-\lambda(b - c(F))$ we are selecting a cut F with smallest cose with respect to c. Note that we can assume that any cut in G has larger cost than b since otherwise the optimum is simply 0. Let B be the minimum cost of cuts in G. We have $-\lambda(b - B) \leq w(C \setminus F) - \lambda(b - c(F))$ for any cut C and $F \subsetneq C$. Thus the upperbound is $\epsilon = \min \frac{w(C \setminus F)}{B - c(F)}$.

2.3 differences

Consider $L(\lambda)$ for cut problem. One can see that the optimal λ is clearly 0 since $L(\lambda)$ is pwl concave and the slope is negative at $\lambda = 0$. What we are really interested in is the first segment on L. At the left end, $L(0)$ is exactly the weight of minimum cut. (the complementary slackness condition is satisfied.) At the right end, as we have shown in the previous paragraph, λ equals to the value of the strength (which is the optimum of the linear relaxation of the cut IP). However, for cut interdiction problems $L(0)$ is not the optimum.

2.4 integrality gap

I guess the 2-approximate min-cut enumeration algorithm implies an integrality gap of 2 for cut interdiction problem.

First consider the dual of linear relaxation of Equation 1.

$$\begin{aligned} \max \quad & \sum_T z_T - b\lambda \\ s.t. \quad & \sum_{T \ni e} z_T \le w(e) \quad \forall e \in E \\ & \sum_{T \ni e} z_T \le c(e)\lambda \quad \forall e \in E \\ & z_T, \lambda \ge 0 \end{aligned} \tag{2}$$

We want to prove something like tree packing theorem for Equation 2.

Conjecture 2.1 *The optimum of Equation 2 is* $\min\left\{\frac{w(C\setminus F)}{B-c(F)} | \forall cut\ C, c(F) \le b\right\}$*, where B is the cost of mincut in G and b is the budget.*

I believe the previous conjecture is not likely to be true.

Weight truncation Assuming we know the optimal λ to the LP dual, Equation 2 in fact gives the idea of weight truncation. The capacity of each edge e in the "tree packing" is $\min\{c(e)\lambda, w(e)\}$. Therefore, the optimum of Equation 2 is $\Lambda^{fr}_{w_\tau} - b\lambda$, where $\Lambda^{fr}_{w_\tau}$ is the fractional mincut on G with weights w_τ.

The optimal λ Denote by λ^* the optimal λ that maximizes $L(\lambda)$. From the previous argument on the first segment of $L(\lambda)$ we know that $\lambda^* \ge \min \frac{w(C\setminus F)}{B-c(F)}$. Now assume $\lambda^* > \min_{c(F)\le b} \frac{w(C\setminus F)}{b-c(F)}$. We have $\min w(C \setminus F) - \lambda^*(b - c(F)) < w(C \setminus F) - \min_{c(F)\le b} \frac{w(C\setminus F)}{b-c(F)}(b - c(F)) = 0$ since the optimum must be achieved by F such that $0 \le b - c(F)$(the slope). The negative optimum of $L(\lambda)$ contradicts the fact that $L(0) = 0$ and L is concave. Hence, the optimal solution λ^* is in the range $[\min \frac{w(C\setminus F)}{B-c(F)}, \min_{c(F)\le b} \frac{w(C\setminus F)}{b-c(F)}]$.

It would be nice if we can prove that any breakpoint is of the form $\min \frac{w(C\setminus F)}{b'-c(F)}$ for some $b' \in [b, B]$. However, this seems incorrect. Let $\{(C_0, F_0), \dots, (C_h, F_h)\}$ be the sequence of solutions for each segment on $L(\lambda)$ and let $\lambda_1 < \cdots < \lambda_h$ be the sequence of breakpoints. (λ_i is the intersection of the corresponding segments of (C_{i-1}, F_{i-1}) and (C_i, F_i).)

Lemma 2.2 $\lambda_i = \min \frac{w(C\setminus F) - w(C_{i-1}\setminus F_{i-1})}{c(F_{i-1}) - c(F)}$*, where the minimum is taken over all cut C and $F \subset C$ such that both the numerator and denominator are positive.*

The proof is using the argument for showing $\lambda_1 = \min \frac{w(C\setminus F)}{B-c(F)}$ and induction. λ_i looks similar to normalized mincut but is related to the slope and vertical intercept of a previous segment.

Conjecture 2.3 *Equation 1 has an integrality gap of 4.*

However, Conjecture 2.3 is wrong. The integrality gap is unbounded. Consider a cycle C_n of n vertices with two special edges e_1, e_2. Let L be a large number.

$$w(e) = \begin{cases} 1 & e = e_1 \\ L & e = e_2 \\ 2 & \text{else} \end{cases}, \quad c(e) = \begin{cases} L & e = e_1 \\ 1 & \text{else} \end{cases}, \quad b = 2 - \epsilon$$

For IP, it is clear that $F = \{e_2\}, C \setminus F = \{e_1\}$ and the optimum is 1
For LP, we assign $x = 0$ and $y_e = \frac{1}{n-2}$ for every edge except e_1. The optimum is 0.
What is the gap if we only relax λ in the Lagrangian dual?

3 Random Stuff

3.1 remove box constraints

Given a positive covering LP,

$$\begin{aligned} LP1 = \min \quad & \sum_e w(e)x_e \\ s.t. \quad & \sum_{e\in T} c(e)x_e \geq k \quad \forall T \\ & c(e) \geq x_e \geq 0 \quad \forall e, \end{aligned}$$

we want to remove constraints $c(e) \geq x_e$. Consider the following LP,

$$\begin{aligned} LP2 = \min \quad & \sum_e w(e)x_e \\ s.t. \quad & \sum_{e\in T} c(e)x_e \geq k \qquad \forall T \\ & \sum_{e\in T\setminus f} c(e)x_e \geq k - c(f) \quad \forall T\ \forall f \in T \\ & x_e \geq 0 \qquad \forall e, \end{aligned}$$

These two LPs have the same optimum. Any feasible solution to LP1 is feasible in LP2. Thus OPT($LP1$) ≥ OPT($LP2$). Next we show that any x_e in the optimum solution to LP2 is always in $[0, c(e)]$. Let x^* be the optimum and suppose that $c(f) < x_f \in x^*$. Consider all constraints $\sum_{e\in T\setminus f} c(e)x_e \geq k - c(f)$ on $T \ni f$. For any such constraint, we have $\sum_{e\in T} c(e)x_e > k$ since we assume $x_f > c(f)$, which means we can decrease x_f without violating any constraint. Thus it contradicts the assumption that x^* is optimal. Then we can add redundant constraints $x_e \leq c(e)\ \forall e$ to LP2 and see that LP1 and LP2 have the same optimum.

This applies to [1] but cannot get an improvement on their algorithm.(MWU does not care the number of constraints.) So does this trick apply to connectivity interdiction?

$$\min_{\text{cut } C, f\in C} \frac{\sum_{e\in C\setminus\{f\}} w(e)x_e}{k - c(f)}$$

References

[1] Parinya Chalermsook, Chien-Chung Huang, Danupon Nanongkai, Thatchaphol Saranurak, Pattara Sukprasert, and Sorrachai Yingchareonthawornchai. Approximating k-Edge-Connected Spanning Subgraphs via a Near-Linear Time LP Solver. *LIPIcs, Volume 229, ICALP 2022*, 229:37:1–37:20, 2022.

[2] Chandra Chekuri, Kent Quanrud, and Chao Xu. LP Relaxation and Tree Packing for Minimum k-Cut. *SIAM Journal on Discrete Mathematics*, 34(2):1334–1353, January 2020.

[3] William H. Cunningham. Optimal attack and reinforcement of a network. *Journal of the ACM*, 32(3):549–561, July 1985.

[4] Chien-Chung Huang, Nidia Obscura Acosta, and Sorrachai Yingchareonthawornchai. An FPTAS for Connectivity Interdiction. In Jens Vygen and Jarosław Byrka, editors, *Integer Programming and Combinatorial Optimization*, volume 14679, pages 210–223, Cham, 2024. Springer Nature Switzerland.