CSC2420: Algorithm Design, Analysis and Theory Spring 2019

Allan Borodin

March 14, 2019

Week 9

Announcements:

1. One more standard type of assignment. I have so far posted 2 questions. The assignment will be due April 4.
2. A critical review of a topic. Please send me your proposal for a topic (perhaps with some relevant papers to help clarify the topic) as soon as possible. The critical review should indicate the state of the topic, what is known and what outstanding issues remain, what are the most relevant papers, what algorithms, proof techniques, algorithms are particularly novel or what is relatively standard.

   I would suggest a 5-10 page review but that is just a rough guideline. The written review and a 10-15 minute oral presentation is due April 4.

Todays agenda

Todays agenda

The lecture slides will be sketchy as I want to mention a number of topics. In addition to any posted slides, chapters 7 and 8 are being written now and you will find related material there.

- Complete the discussion of the two sided greedy algorithm for the unconstrained submodular maximization (USM) problem. The relation between the USM proof and the presentation in Poloczek et al randomized algorithm for max-sat.
- Discuss proof(s) of the KVV algorithm for bipartite matching.
- Brief introduction to primal dual algorithms and primal dual based analysis of combinatorial algorithms.
- Extensions of bipartite matching following Chris's lecture last week.
- ROM and iid input models
- Other computational models where randomization is provably necessary.

A comment on the focus of this course

Note: To give some focus to this course, I am emphasizing conceptually simple algorithms (eg greedy and online) and the role of randomization in algorithms. Online algorithms (and other algorithmic models) provide a setting where we can establish proofs (without complexity assumptions) of the power of randomization. We will also consider, streaming algorithms, sublinear time algorithms and random walk algorithms.

In the lecture for Week 7, we stated both the deterministic and randomized algorithms for the USM problem. We also sketched the proof for the deterministic $\frac{1}{3}$ bound.

For convenience we repeat these algorithms which I refer to as the deterministic and randomized double-sided greedy algorithms for the USM problem.

The deterministic 1/3 approximation for USM

$X_0 := \emptyset$; $Y_0 := U$

For $i := 1 \ldots n$

$\quad a_i := f(X_{i-1} \cup \{u_i\}) - f(X_{i-1})$; $b_i := f(Y_{i-1} \setminus \{u_i\}) - f(Y_{i-1})$

$\quad$ **If** $a_i \geq b_i$

$\quad\quad$ **then** $X_i := X_{i-1} \cup \{u_i\}$; $Y_i := Y_{i-1}$

$\quad\quad$ **else** $X_i := X_{i-1}$; $Y_i := Y_{i-1} \setminus \{u_i\}$

$\quad$ **End If**

End For

The randomized double-sided algorithm for the USM problem

The randomized 1/2 approximation for USM

$X_0 := \varnothing; Y_0 := U$

For $i := 1 \ldots n$

$a_i := f(X_{i-1} \cup \{u_i\}) - f(X_{i-1})$; $b_i := f(Y_{i-1} \setminus \{u_i\}) - f(Y_{i-1})$

$a'_i = \max\{a_i, 0\}; b'_i = \max\{b_i, 0\}$

If $a_i = b_i = 0$

then Set $X_i := X_{i-1} \cup \{u_i\}; Y_i := Y_{i-1}$

Else

With probability $\frac{a'}{a'+b'}$

Set $X_i := X_{i-1} \cup \{u_i\}; Y_i := Y_{i-1}$

With probability $\frac{b'}{a'+b'}$

Set $X_i := X_{i-1}; Y_i := Y_{i-1} \setminus \{u_i\}$

End If

End For

Outline of the randomized $\frac{1}{2}$ bound

We had the following key lemma for the deterministic case:

$$f(OPT_{i-1} - f(OPT_i) \leq [f(X_i - f(X_{i-1}] + [f(Y_i) - f(Y_{i-1}]$$

The proof of the $\frac{1}{3}$ bound then follows from this key lemma by a standard telescoping argument.

For the randomized algorithm, we again have OPT_i, X_i, Y_i which are now, of course, random variables. We then have the following key lemma for the randomized double-sided greedy algorithm:

$$\mathbb{E}[f(OPT_{i-1}) - f(OPT_i)] \leq \frac{1}{2}\mathbb{E}[f(X_i) - f(X_{i-1})] + [f(Y_i) - f(Y_{i-1})]$$

And as before, a standard telescoping argument establishes the $\frac{1}{2}$ bound.

The correspondence between the randomized max-sat and USM algorithms

It is no coincidence that there is a strong correspondence between the proof of the randomized $\frac{3}{4}$ bound for max-sat and the $\frac{1}{2}$ bound for the USM problem. As we noted, the Buchbinder et al USM algorithm can be extended to provide a max-sat algorithm which can be seen to be equivalent to an algorithm independently given by van Zuylen as discussed in Poloczek et al whose analyis we follow.

This hold for the high level structure of the proofs based on the key lemmas to bound $\mathbb{E}[f(OPT_{i-1}) - f(OPT_i)]$ and then using a telescoping argument. Moreover, these key lemmas also follow a similar sequence of supporting lemmas.
In particular, we have $f_i + t_i \geq 0$ correspnding to $a' + b' \geq 0$ and $\mathbb{E}[w(OPT_{i-1} - w(OPT_i)] \leq \max\{0, \frac{2f_i t_i}{f_i+t_i}\}$ corresponding to $\mathbb{E}[w(OPT_{i-1} - w(OPT_i)] \leq \max\{0, \frac{a'_i b'_i}{a'_i+b'_i}\}$

Note: We will soon indicate this correspondence in chapter 8 of the text.

The KVV algorithm

Chris stated the KVV Ranking algorithm for the unweighted bipartite matching problem (BMM). The Karp, Vazirani and Vazirani paper is a seminal paper that has led to many extentions relating to online advertising (e.g. the ADWORDS problem).

The KVV algorithm is at first sight (if not, always) surprising in that the more natural randomized online greedy algorithm (i.e., match each online node to a random available offline node) has a tight asymptotic $\frac{1}{2}$ competitive ratio.

The Ranking algorithm instead initally chooses a random permutation of the offline nodes so as to provide a ranking for choosing a matching node.

Given its importance (and because there was a techincal error in the original proof), there have been a number of alternative proofs. In Chapter 8 of the text we are writing, there is a combinatiorial proof due to Birnbaum and Mathieu. A number of alternative proofs rely on primal dual analysis.

The Ranking algorithm

The Ranking algorithm

V is the set of offline nodes and U is the set of online nodes.
Choose a permutation (ranking) σ on V uniformly at random
$M := \varnothing$
$i := 1$
While $i \leq n$ **do**
    Online node u_i arrives together with its neighbourood $N(u_i)$
    Let $N_c(u_i)$ be the nodes in $N(u_i)$ that are not yet matched.
    If $N_c(u_i) \neq \varnothing$ **then**
        Let $v = argmin\{\sigma(u) : v \in N_c(u_i)\}$
        $M := M \cup \{u_i, v)\}$ $i := i + 1$
End While

We will sketch a proof in our text that KVV acheives competitve ratio $1 - \frac{1}{e}$ using the proof in Birnbaum and Mathieu.

Sketch of KVV competitive bound

Let p_t denote the probability over σ that the vertex of rank t in V is matched by Ranking. We are interested in computing the expected size of the matching returned by Ranking, which is given by $\sum_{t=1}^{n} p_t$. The analysis of Ranking will be centered around establishing the following lemma:

Lemma:

For all $t \in [n]$ we have $1 - p_t \leq (1/n) \sum_{s=1}^{t} p_s$.

From the lemma, the following inequalities can be derived:

$$\sum_{t=1}^{n} p_t = p_1 + \sum_{t=2}^{n} p_t \geq \frac{1}{n} + \left(1 - \frac{1}{n}\right) \sum_{t=1}^{n} \left(\frac{n}{n+1}\right)^{t-1} \geq \frac{1}{n} + \left(1 - \frac{1}{n}\right) \frac{1 -}{1 -}$$

$$= \frac{1}{n} + \left(n - \frac{1}{n}\right) \left[1 - \left(\frac{n}{n+1}\right)^{n}\right] \geq n \left[1 - \left(1 - \frac{1}{n+1}\right)^{n}\right],$$

Completing the proof sketch of the KVV competitive bound

To prove (and motivate) the lemma, we proceed as follows:
Let A_t denote the set of permutations such that a vertex of rank t is matched by Ranking. Let $S_{[n]}$ denote the set of all permutations $V \to V$ and define $B_t = S_{[n]} \setminus A_t$; that is, B_t is the set of permutations such that a vertex of rank t is not matched by Ranking.

We then construct an injection of the form $[n] \times B_t \to \bigcup_{i=1}^{t} A_i$. This will prove the lemma.

Duality: See Vazirani and Shmoys/Williamson texts, and Williamson article

- For a primal maximization (resp. minimization) LP in standard form, the dual LP is a minimization (resp. maximization) LP in standard form.
- Specifically, if the primal $\mathcal{P}$ is:
  - Minimize $\mathbf{c} \cdot \mathbf{x}$
  - subject to $A_{m \times n} \cdot \mathbf{x} \geq \mathbf{b}$
  - $\mathbf{x} \geq 0$
- then the dual LP $\mathcal{D}$ with dual variables $\mathbf{y}$ is:
  - Maximize $\mathbf{b} \cdot \mathbf{y}$
  - subject to $A^{tr}_{n \times m} \cdot \mathbf{y} \leq \mathbf{c}$
  - $\mathbf{y} \geq 0$
- Note that the dual (resp. primal) variables are in correspondence to primal (resp. dual) constraints.
- If we consider the dual $\mathcal{D}$ as the primal then its dual is the original primal $\mathcal{P}$. That is, the dual of the dual is the primal.

An example: set cover

As already noted, the vertex cover problem is a special case of the set cover problem in which the elements are the edges and the vertices are the sets, each set (ie vertex v) consisting of the edges adjacent to v.

The set cover problem as an IP/LP

minimize $\sum_j w_j x_j$
subject to $\sum_{j:e_i \in S_j} x_j \geq 1$ for all i; that is, $e_i \in U$
$x_j \in \{0, 1\}$ (resp. $x_j \geq 0$)

The dual LP

maximize $\sum_i y_i$
subject to $\sum_{i:e_i \in S_j} y_i \leq w_j$ for all j
$y_i \geq 0$

If all the parameters in a standard form minimization (resp. maximization) problem are non negative, then the problem is called a covering (resp. packing) problem. Note that the set cover problem is a covering problem and its dual is a packing problem.

Duality Theory Overview

- An essential aspect of duality is that a finite optimal value to either the primal or the dual determines an optimal value to both.
- The relation between these two can sometimes be easy to interpret. However, the interpretation of the dual may not always be intuitively meaningful.
- Still, duality is very useful because the duality principle states that optimization problems may be viewed from either of two perspectives and this might be useful as the solution of the dual might be much easier to calculate than the solution of the primal.
- In some cases, the dual might provide additional insight as to how to round the LP solution to an integral solution.
- Moreover, the relation between the primal $\mathcal{P}$ and the dual $\mathcal{D}$ will lead to primal-Dual algorithms and to the so-called dual fitting analysis.
- In what follows we will initially assume the primal is a minimization problem to simplify the exposition.

Strong and Weak Duality

Strong Duality

If x^* and y^* are (finite) optimal primal and resp. dual solutions, then $\mathcal{D}(\mathbf{y}^*) = \mathcal{P}(\mathbf{x}^*)$.

Note: Before it was known that solving LPs was in polynomial time, it was observed that strong duality proves that LP (as a decision problem) is in $\mathbf{NP} \cap \mathbf{co} - \mathbf{NP}$ which strongly suggested that LP was not NP-complete.

Weak Duality for a Minimization Problem

If $\mathbf{x}$ and $\mathbf{y}$ are primal and resp. dual solutions, then $\mathcal{D}(\mathbf{y}) \leq \mathcal{P}(\mathbf{x})$.

- Duality can be motivated by asking how one can verify that the minimum in the primal is at least some value z. To get witnesses, one can explore non-negative scaling factors (i.e. the dual variables) that can be used as multipliers in the constraints. The multipliers, however, must not violate the objective (i.e cause any multiplies of a primal variable to exceed the coefficient in the objective) we are trying to bound.

Motivating duality

Consider the motivating example in V. Vazirani's text:

Primal

minimize $7x_1 + x_2 + 5x_3$

subject to

- (1) $x_1 - x_2 + 3x_3 \geq 10$
- (2) $5x_1 + 2x_2 - x_3 \geq 6$
- $x_1, x_2, x_3 \geq 0$

Dual

maximize $10y_1 + 6y_2$

subject to

$$y_1 + 5y_2 \leq 7$$
$$-y_1 + 2y_2 \leq 1$$
$$3y_1 - y_2 \leq 5$$
$$y_1, y_2 \geq 0$$

Adding (1) and (2) and comparing the coefficient for each x_i, we have:

$7x_1 + x_2 + 5x_3 \geq (x_1 - x_2 + 3x_3) + (5x_1 + 2x_2 - x_3) \geq 10 + 6 = 16$

Better yet,

$7x_1 + x_2 + 5x_3 \geq 2(x_1 - x_2 + 3x_3) + (5x_1 + 2x_2 - x_3) \geq 26$

For an upper bound, setting $(x_1, x_2, x_3) = (7/4, 0, 11/4)$

$7x_1 + x_2 + 5x_3 = 7 \cdot (7/4) + 1 \cdot 0 + 5 \cdot (11/4) = 26$

This proves that the optimal value for the primal and dual (with solution $(y_1, y_2) = (2, 1)$ must be 26.

Easy to prove weak duality

The proof for weak duality

$\mathbf{b} \cdot \mathbf{y} = \sum_{j=1}^{m} b_j y_j$
$\leq \sum_{j=1}^{m} (\sum_{i=1}^{n} A_{ji} x_i) y_j$
$\leq \sum_{i=1}^{n} \sum_{j=1}^{m} (A_{ji} y_j) x_i$
$\leq \sum_{i=1}^{n} c_i x_i = \mathbf{c} \cdot \mathbf{x}$

Max flow-min Cut in terms of duality

- While the max flow problem can be naturally formulated as a LP, the natural formulation for min cut is as an IP. However, for this IP, it can be shown that the *extreme point solutions* (i.e. the vertices of the polyhedron defined by the constraints) are all integral $\{0,1\}$ in each coordinate. Moreover, there is a precise sense in which max flow and min cut can be viewed as dual problems. This is described nicely in Vazarani (section 12.2).
- In order to formulate max flow in standard LP form we reformulate the problem so that all flows (i.e. the LP variables) are non-negative. And to state the objective as a simple linear function (of the flows) we add an edge of infinite capacity from the terminal t to the source s and hence define a circulation problem.

The max flow LP

maximize $f_{t,s}$

subject to $f_{i,j} \leq c_{i,j}$ for all $(i,j) \in E$

$\sum_{j:(j,i)\in E} f_{j,i} - \sum_{j:(i,j)\in E} f_{i,j} \leq 0$ for all $i \in V$

$f_{i,j} \geq 0$ for all $(i,j) \in E$

Max flow-min cut duality continued

For the primal edge capacity constraints, introduce dual ("distance") variables $d_{i,j}$ and for the vertex flow conservation constraints, introduce dual ("potential") variables p_i.

The fractional min cut dual

minimize $\sum_{(i,j)\in E} c_{i,j} d_{i,j}$

subject to $d_{i,j} - p_i + p_j \geq 0$

$p_s - p_t \geq 1$

$d_{i,j} \geq 0; p_i \geq 0$

- Now consider the IP restriction : $d_{i,j}, p_i \in \{0,1\}$ and let $\{(d^*_{i,j}, p^*_i)\}$ be an intergal optimum.
- The $\{0,1\}$ restriction and second constraint forces $p^*_s = 1; p^*_t = 0$.
- The IP optimum then defines a cut (S,T) with $S = \{i | p^*_i = 1\}$ and $T = \{i | p^*_i = 0\}$.
- Suppose (i,j) is in the cut, then $p^*_i = 1, p^*_j = 0$ which by the first constraint forces $d_{i,j} = 1$.
- The optimal $\{0,1\}$ IP solution (of the dual) defines a a min cut.

Solving the f-frequency set cover by a primal dual algorithm

- In the f-frequency set cover problem, each element is contained in at most f sets.
- Clearly, the vertex cover problem is an instance of the 2-frequency set cover.
- As in the vertex cover LP rounding, we can similarly solve the f-frequency cover problem by obtaining an optimal solution $\{x_j^*\}$ to the (primal) LP and then rounding to obtain $\bar{x}_j = 1$ iff $x_j^* \geq \frac{1}{f}$. This is, as noted before, a conceptually simple method but requires solving the LP.
- We know that for a minimization problem, any dual solution is a lower bound on any primal solution. One possible goal in a primal dual method for a minimization problem will be to maintain a fractional feasible dual solution and continue to try improve the dual solution. As dual constraints become tight we then set the corresponding primal variables.

Primal dual for f-frequency set cover continued

Suggestive lemma

Claim: Let $\{y_i^*\}$ be an optimal solution to the dual LP and let $\mathcal{C}' = \{S_j | \sum_{e_i \in S_j} y_i^* = w_j\}$. Then $\mathcal{C}'$ is a cover.

This suggests the following algorithm:

Primal dual algorithm for set cover

Set $y_i = 0$ for all i

$\mathcal{C}' := \varnothing$

While there exists an e_i not covered by $\mathcal{C}'$

  Increase the dual variables y_i until there is some $j : \sum_{\{k: e_i \in S_j\}} y_i = w_j$

  $\mathcal{C}' := \mathcal{C}' \cup \{S_j\}$

  Freeze the y_i associated with the newly covered e_i

End While

Theorem: Approximation bound for primal dual algorithm

The cover formed by tight constraints in the dual solution provides an f approximation for the f-frequency set cover problem

Comments on the primal dual algorithm

- What is being shown is that the integral primal solution is within a factor of f of the dual solution which implies that the primal dual algorithm is an f-approximation algorithm for the f-frequency set cover problem.
- In fact, what is being shown is that the integraility gap of this IP/LP formulation for f-frequency set cover problem is at most f.
- In terms of implementation we would calculate the minimum ϵ needed to make some constraint tight so as to chose which primal variable to set. This ϵ could be 0 if a previous iteration had more than one constraint that becomes tight simultaneously. This ϵ would then be subtracted from w_j for j such that $e_i \in S_j$.

Using dual fitting to prove the approximation ratio of the greedy set cover algorithm

We have already seen the following natural greedy algorithm for the weighted set cover problem:

The greedy set cover algorithm

$\mathcal{C}' := \varnothing$

While there are uncovered elements

Choose S_j such that $\frac{w_j}{|\tilde{S}_j|}$ is a minimum where

$\tilde{S}_j$ is the subset of S_j containing the currently uncovered elements

$\mathcal{C}' := \mathcal{C}' \cup S_j$

End While

We wish to prove the following theorem (Lovasz[1975], Chvatal [1979]):

Approximation ratio for greedy set cover

The approximation algorithm for the greedy algorithm is H_d where d is the maximum size of any set S_j.

The dual fitting analysis

The greedy set cover algorithm setting prices for each element

$\mathcal{C}' := \varnothing$

While there are uncovered elements

Choose S_j such that $\frac{w_j}{|\tilde{S}_j|}$ is a minimum where

$\tilde{S}_j$ is the subset of S_j containing the currently uncovered elements

%Charge each element e in $\tilde{S}_j$ the average cost $price(e) = \frac{w_j}{|\tilde{S}_j|}$

% This charging is just for the purpose of analysis

$\mathcal{C}' := \mathcal{C}' \cup S_j$

End While

- We can account for the cost of the solution by the costs imposed on the elements; namely, $\{price(e)\}$. That is, the cost of the greedy solution is $\sum_e price(e)$.

Dual fitting analysis continued

- The goal of the dual fitting analysis is to show that $y_e = price(e)/H_d$ is a feasible dual and hence any primal solution must have cost at least $\sum_e price(e)/H_d$.
- Consider any set $S = S_j$ in $\mathcal{C}$ having say $k \leq d$ elements. Let $e_1, \ldots, e_k$ be the elements of S in the order covered by the greedy algorithm (breaking ties arbitrarily). Consider the iteration is which e_i is first covered. At this iteration $\tilde{S}$ must have at least $k - i + 1$ uncovered elements and hence S could cover cover e_i at the average cost of $\frac{w_j}{k-i+1}$. Since the greedy algorithm chooses the most cost efficient set, $price(e_i) \leq \frac{w_j}{k-i+1}$.
- Summing over all elements in S_j, we have $\sum_{e_i \in S_j} y_{e_i} = \sum_{e_i \in S_j} price(e_i)/H_d \leq \sum_{e_i \in S_j} \frac{w_j}{k-i+1} \frac{1}{H_d} = w_j \frac{H_k}{H_d} \leq w_j$. Hence $\{y_e\}$ is a feasible dual.

More comments on primal dual algorithms

- We have just seen an example of a basic form of the primal dual method for a minimization problem. Namely, we start with an infeasible integral primal solution and feasible (fractional) dual. (For a covering primal problem and dual packing problem, the initial dual solution can be the all zero solution.) Unsatisfied primal constraints suggest which dual constraints might be tightened and when one or more dual constraints become tight this determines which primal variable(s) to set.
- Some primal dual algorithms extend this basic form by using a second (reverse delete) stage to achieve minimality.
- **NOTE** In the primal dual method we are not solving any LPs. Primal dual algorithms are viewed as "combinatorial algorithms" and in some cases they might even suggest an explicit greedy algorithm.

Dual fitting applied to a maximization problem

Krysta [2005] applies dual fitting approach to a maximization problem, namely to analyze (in my terminology) fixed order prioriity algorithms (such as the Lehman et al [1999] greedy $2\sqrt{m}$ approximate set packing algorithm) for generalizations of the weighted set packing problem (which can be used to formulate many natural integer packing problems).

Generalized Set Packing

As in weighted set packing, we have a collection of sets $S \in \mathcal{S}$ over some universe U. Each set has a weight w_S. Now we allow sets to be multi-sets and let $q(u, S)$ to be the number of copies of $u \in S$. Furthermore, we also allow each element $u \in U$ to have some maximum number b_u of copies that can occur in a feasible solution (in contrast to the basic set packing problem where $b_u = 1$ for all $u \in U$).
The goal is to select a subcollection $\mathcal{C}$ of sets satisfying the feasibility constraints on the $\{b_u\}$ so as to maximize the sum of the weights of the sets in $\mathcal{C}$.

The natural IP and LP relaxation

The natural IP/LP

max $\sum_{S \in \mathcal{S}} w_S x_S$

- subject to $\sum_{S: u \in S} q(u,S) x_S \leq b_u \quad \forall u \in U$
- $x_S \in \{0,1\}$

In the LP relaxation, the {0,1} constraint becomes $0 \leq x_S \leq 1$}

NOTE: Unlike set cover, for set packing the condition $x_S \leq 1$ is necessary

The minimization dual

min $\sum_{u \in U} b_u y_u + \sum_{S \in \mathcal{S}} z_S$

- subject to $z_S + \sum_{u \in S} q(u,S) y_u \geq w_S \quad \forall S \in \mathcal{S}$
- $z_S, y_u \geq 0$

NOTE: The dual variable z_S corresponds to the constraint $x_S \leq 1$

The secretary problem as an LP

We recall the classical secrtary problem (defined in Lecture 2) which is to maximize the probability of choosing the best candidate from N candidates that arrive in random order. Bucnbinder, Kain and Singh [2010] show how to view the classical secretary problem (and many generalization) as an LP maximization problem with the following benefits:

1. Finding an optimal mechanism reduces to solving a specific linear program
2. Proving that $\frac{1}{e}$ is the best bound possible reduces to finding a solution to the dual of the LP.
3. This approach facilitates the analysis of many generalizations of the secretary problem (i.e. by adding additional constraints or modifying the objective function).
4. One of the generalizations is to obtain a *truthful* mechanism whereby agents (i.e. candidates) have no incentive to seek a particular place in the ordering (and hence making a random order more meaningful).

The LP for the classical secretary problem

The primal LP $\mathcal{P}$

max $\frac{1}{n} \sum_{i=1}^{N} i \cdot p_i$

- subject to: $i \cdot p_i \leq 1 - \sum_{j=1}^{i-1} p_j \quad 1 \leq i \leq N$
- $p_i \geq 0$

The dual LP $\mathcal{D}$

min $\sum_{i}^{N} x_i$

- subject to: $\sum_{j=i+1}^{N} x_j + i \cdot x_i \geq \frac{i}{N} \quad 1 \leq i \leq N$
- $x_i \geq 0$

Sketch of LP characterization

To prove that this LP captures the secrtary problem one needs to prove:

- If M is any mechanism and p_i^M is the probability that M selects the candidate in position i. Then $\{p_i^M\}$ is a feasible solution for the primal $\mathcal{P}$ and Prob[M selects best candidate] $\leq$ the objective value of $\mathcal{P}$
- Let $\{p_i\}$ be any feasible solution of $\mathcal{P}$. Then the following mechanism M obtains the objective function of $\mathcal{P}$:
  Select candidate i with probability $\frac{i \cdot p_i}{(1 - \sum_{j<i} p_j)}$ if the first $i-1$ candidates have not been selected and i is best so far.

Furthermore, to prove an upper bound (namely $\frac{1}{e} + o(1)$) on the best performance (i.e. best probability), it suffices to construct a feasible solution $\{x_i\}$ for the dual $\mathcal{D}$ with dual objective value $\frac{1}{e}$.

- Setting $x_i = 0$ for $1 \leq i \leq N/e$ and $x_i = \frac{1}{N}(1 - \sum_{j=i}^{N} \frac{1}{j})$ for $n/e < i \leq N$ is a feasible dual solution with value $\frac{1}{e}$.

More comments on primal dual algorithms

- We have just seen an example of a basic form of the primal dual method for a minimization problem. Namely, we start with an infeasible integral primal solution and feasible (fractional) dual. (For a covering primal problem and dual packing problem, the initial dual solution can be the all zero solution.) Unsatisfied primal constraints suggest which dual constraints might be tightened and when one or more dual constraints become tight this determines which primal variable(s) to set.
- Some primal dual algorithms extend this basic form by using a second (reverse delete) stage to achieve minimality.
- **NOTE** In the primal dual method we are not solving any LPs. Primal dual algorithms are viewed as "combinatorial algorithms" and in some cases they might even suggest an explicit greedy algorithm.

Primal dual analysis for the Ranking algorithm

A number of papers use primal dual analysis for BMM and related problems.

Warning: Although I prefer to use U (or R for right side) for the online vertices and V (or L) for the offline vertices, this is not standard and various papers reverse the role of U and V (R and L). In what follows, I will use L (for left) for the offline nodes and R (right) for the online nodes.

Here is the standard primal IP for matching in an arbitrary graph $G = (V, E)$:

Maximize $\sum_{(i,j)\in E} x_{ij}$

subject to $\sum_{j:(i,j)\in E} x_{ij} \leq 1 \quad \forall i \in V$

$x_{ij} \in \{0,1\} \quad \forall (i,j) \in E$

The LP relaxation and its dual

Primal LP
Maximize $\sum_{(i,j)\in E} x_{ij}$
subject to $\sum_{j:(i,j)\in E} x_{ij} \leq 1 \quad \forall i \in V$

$$x_{ij} \geq 0 \quad \forall (i,j) \in E$$

Dual LP stated for a bipartite graph
Minimize $\sum_{i\in L} \alpha_i + \sum_{j\in R} \beta_j$
subject to $\alpha_i + \beta_j \geq 1 \quad \forall (i,j) \in E$

$$\alpha_i, \beta_j \geq 0 \quad \forall i, j$$

Sketch of this primal dual proof of the KVV algorithm

This is the proof in Devanur, Jain and R. Kleinberg

First we need to slightly restate the KVV algorithm as follows so that rather choose a random permutation of the offline nodes, we choose a random $Y_i \in [0,1]$ for each offline node $i \in L$. Then in matching an online vertex we choose the available offline node (if any exists) have the smallest value Y_i. (This type of restatement is the extension to offline vertex weighted online matching.)
In order to establish a c competitive ratio, it will suffice to find a dual soution D such that $P \geq c \cdot D$ where P is the primal solution. We note that P and D will be random variables but that the inequality $P \geq c \cdot D$ will hold for every instantaition of the algorithm.

We also need to relate the algorithm to the dual variables.

It is also hel;pful that we know the ratio c that we are aiming for, namely $c = 1 - \frac{1}{e}$.

Vertex weighted bipartite matching

- Aggarwal et al [2011] consider a vertex weighted version of the online bipartite matching problem. Namely, the vertices $v \in V$ all have a known weight w_v and the goal is now to maximize the weighted sum of matched vertices in V when again vertices in U arrive online.
- This problem can be shown to subsume the adwords problem when all bids $b_{q,i} = b_i$ from an advertiser are the same.
- It is easy to see that Ranking can be arbitrarily bad when there are arbitrary differences in the weight. Greedy (taking the maximum weight match) can be good in such cases. Can two such algorithms be somehow combined? Surprisingly, Aggarwal et al are able to achieve the same 1-1/e bound for this class of vertex weighted bipartite matching.

The vertex weighted online algorithm

The perturbed greedy algorithm

For each $v \in V$, pick r_v randomly in $[0, 1]$
Let $f(x) = 1 - e^{1-x}$
When $u \in U$ arrives, match u to the unmatched v (if any) having the highest value of $w_v * f(x_v)$. Break ties consistently.

In the unweighted case when all w_v are identical this is the Ranking algorithm.

Moving away from the adversarial approach

While competitive online analysis was initially motivated by arguing for worst case (i.e., adversarial inputs) analysis (as in Sleator and Tarjan's seminal article regarding online list accessing and paging), the interest in bipartite matching (and extensions) might arguably be said to have led to a renaissance in the study of online algorithms in distributional models.

In particular, the most studied distributional input setting is the i.i.d. setting which has been prominent in various BMM problem extensions (e.g. ADWORDS) and also in the more basic BMM problem. In particular, starting with a 2009 paper by Feldman et al, there has been a succession of papers showing algorithms that beat the ratio $1 - 1/e$ in the known and unknown i.i.d. settings.

Another well studied distributional setting is the *random order model* (ROM) setting.