Budgeted Allocations in the Full-Information Setting⋆

Aravind Srinivasan[1]

Dept. of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742.

Abstract. We build on the work of Andelman & Mansour and Azar, Birnbaum, Karlin, Mathieu & Thach Nguyen to show that the full-information (i.e., offline) budgeted-allocation problem can be approximated to within 4/3: we conduct a rounding of the natural LP relaxation, for which our algorithm matches the known lower-bound on the integrality gap.

1 Introduction

Sponsored-search auctions are a key driver of advertising, and are a topic of much current research (Lahaie, Pennock, Saberi & Vohra [10]). A fundamental problem here is *online budgeted allocation*, formulated and investigated by Mehta, Saberi, Vazirani & Vazirani [12]. Recent work has also focused on the offline version of this basic allocation problem; we improve on the known results, demonstrating a rounding approach for a natural LP relaxation that yields a 4/3-approximation, matching the known integrality gap. We also show that in the natural scenario where bidders' individual bids are much smaller than their budgets, our algorithm solves the problem almost to optimality.

Our problem is as follows. We are given a set U of *bidders* and a set V of *keywords*. Each bidder i is willing to pay an amount $b_{i,j}$ for keyword j to be allocated to them; each bidder i also has a budget B_i at which their total payment is capped. Our goal is to assign each keyword to at most one bidder, in order to maximize the total payment obtained. This models the problem of deciding which bidder (if any) gets to be listed for each keyword, in order to maximize the total revenue obtained by, say, a search engine. That is, we want to solve the following integer linear program (ILP), where $x_{i,j}$ is the indicator variable for keyword j getting assigned to bidder i: maximize $\sum_{i \in U} \min\{B_i, \ \sum_{j \in V} b_{i,j} x_{i,j}\}$, subject to $\sum_i x_{i,j} \leq 1$ for each j, and $x_{i,j} \in \{0, 1\}$ for all (i, j). (It is easy to see that the "min" term can be appropriately rewritten in order to express this as a standard ILP.)

⋆ Research supported in part by NSF ITR Award CNS-0426683 and NSF Award CNS-0626636. Part of this work was done while the author was on sabbatical at the Network Dynamics and Simulation Science Laboratory of the Virginia Bioinformatics Institute, Virginia Tech.

Known results. This NP-hard problem has been studied by Garg, Kumar & Pandit, who presented an $(1+\sqrt{5})/2 \sim 1.618$-approximation algorithm for the problem [8]. (As usual, for our maximization problem, a ρ-approximation algorithm, for $\rho \geq 1$, is a polynomial-time algorithm that always presents a solution of value at least $1/\rho$ times optimal; in the case of randomized algorithms, the *expected* solution-value should be at least $1/\rho$ times optimal.) In addition to other results, Lehmann, Lehmann & Nisan [11] have developed a greedy 2-approximation algorithm for this problem. Now, the natural LP relaxation for the problem is obtained by relaxing each $x_{i,j}$ to lie in $[0,1]$, in the above ILP. Andelman & Mansour [2] presented a rounding algorithm for this LP that achieves an approximation of $e/(e-1) \sim 1.582$; this was improved – for a more general problem – by Feige & Vondrak to $e/(e-1) - \epsilon$, for an ϵ that is about 10^{-4} [6]. More recently, Azar, Birnbaum, Karlin, Mathieu & Thach Nguyen [3] have improved the approximation ratio to 3/2. There are also two interesting special cases of the problem: the *uniform* case, where each j has a price p_j such that $b_{i,j} \in \{0, p_j\}$ for all i, and the case where all the budgets B_i are the same. Two additional results are obtained in [2]: that the integrality gap of the above LP-relaxation is at least 4/3 even for the first (i.e., uniform) special case, and that the second special case can be approximated to within 1.39. See, e.g., [12, 4, 9] for online versions of the problem.

Our results. We build on the work of [2, 3] and show how to round the LP to obtain an approximation of 4/3: note from the previous paragraph that this meets the integrality gap. Anna Karlin (personal communication, March 2008) has informed us that Chakrabarty & Goel have independently obtained this approximation ratio, as well as improved hardness-of-approximation results – a preprint of this work is available [5]. We also present two extensions in Section 3: (a) the important special case where each bidder's bids are much smaller than their budget [12, 4] can be solved near-optimally: if, for some $\epsilon \in [0,1]$, $b_{i,j} \leq \epsilon \cdot B_i$ for all (i,j), our algorithm's approximation ratio is $4/(4-\epsilon)$; and (b) suppose that for some $\lambda \geq 1$, we have for all (i,j,j') that if $b_{i,j}$ and $b_{i,j'}$ are nonzero, then $b_{i,j} \leq \lambda \cdot b_{i,j'}$. For this case, our algorithm yields a better-than-4/3 approximation if $\lambda < 2$. In particular, if $\lambda = 1$, our algorithm has an approximation ratio of $(\sqrt{2}+1)/2 \sim 1.207$.

2 The algorithm and analysis

We will round the natural LP-relaxation mentioned in Section 1. Our algorithm is randomized, and can be derandomized using the method of conditional probabilities.

Observe that for the original (integral) problem, setting

$$b_{i,j} := \min\{b_{i,j}, B_i\} \tag{1}$$

keeps the problem unchanged. Thus, we will assume

$$\forall (i,j),\ b_{i,j} \leq B_i. \tag{2}$$

Notation. When we refer to the *load* on a bidder i w.r.t. some (fractional) allocation x, we mean the sum $\sum_j b_{i,j} x_{i,j}$; note that we do not truncate at B_i in this definition.

Suppose we are given some feasible fractional allocation x; of course, the case of interest is where this is an optimal solution to the LP, but we do not require it. It is also immediate that the following assumption is without loss of generality:

$$\text{if } b_{i,j} = 0, \text{ then } x_{i,j} = 0. \tag{3}$$

As in [3], we may assume that the bipartite graph (with (U, V) as the partition) induced by those $x_{i,j}$ that lie in $(0, 1)$, is a *forest* F. This can be effected by an efficient algorithm, such that the resulting fractional objective-function value equals that of the original value that we started with [3]. This forest F is the structure that we start with; we show how to round those $x_{i,j}$ in F. We are motivated by the approaches of [1, 13, 7]; however, our method is different, especially in step **(P2)** below. Each iteration is described next.

2.1 Iteration s, $s \geq 1$

Remove all (i, j) that have already been rounded to 0 or 1; let F be the current forest consisting of those $x_{i,j}$ that lie in $(0, 1)$. Choose any **maximal** path $P = (w_0, w_1, \ldots, w_k)$ in F; we will now probabilistically round at least one of the edges in P to 0 or 1. For notational simplicity, let the current x value of edge $e_t = (w_{t-1}, w_t)$ in P be denoted y_t; note that all the y_t lie in $(0, 1)$. We will next choose values $z_1, z_2, \ldots, z_k$ probabilistically, and update the x value of each edge $e_t = (w_{t-1}, w_t)$ to $y_t + z_t$. Suppose we initialize some value for z_1, and that we have chosen the increments $z_1, z_2, \ldots, z_t$, for some $t \geq 1$. Then, the value z_{t+1} (corresponding to edge $e_{t+1} = (w_t, w_{t+1})$) is chosen as follows:

(P1) If $w_t \in V$ (i.e., is a keyword), then $z_{t+1} = -z_t$ (i.e., we retain the total assignment value of w_t);

(P2) if $w_t \in U$ (i.e., is a bidder), then we choose z_{t+1} so that the load on w_t remains unchanged (recall that in computing the load, we do not truncate at B_{w_t}); i.e., we set $z_{t+1} = -b_{w_t, w_{t-1}} z_t / b_{w_t, w_{t+1}}$, which ensures that the incremental load $b_{w_t, w_{t-1}} z_t + b_{w_t, w_{t+1}} z_{t+1}$ is zero. (Since $x_{w_t, w_{t+1}}$ is nonzero by the definition of F, $b_{w_t, w_{t+1}}$ is also nonzero by (3); therefore, dividing by $b_{w_t, w_{t+1}}$ is admissible.)

Observe that the vector $z = (z_1, z_2, \ldots, z_k)$ is completely determined by z_1, the path P, and the matrix of bids b; more precisely, there exist reals $c_1, c_2, \ldots, c_k$ that depend only on the path P and the matrix b, such that

$$\forall t, \; z_t = c_t z_1. \tag{4}$$

We will denote this resultant vector z by $f(z_1)$.

Now let μ be the smallest positive value such that if we set $z_1 := \mu$, then all the x values (after incrementing by the vector z as mentioned above) stay in $[0, 1]$, and at least one of them becomes 0 or 1. Similarly, let γ be the smallest positive value such that if we set $z_1 := -\gamma$, then this "rounding-progress" property holds. (It is easy to see that μ and γ are strictly positive, since all the y_i lie in $(0, 1)$.) We now choose the vector z as follows:

(R1) with probability $\gamma/(\mu + \gamma)$, let $z = f(\mu)$;
(R2) with the complementary probability of $\mu/(\mu + \gamma)$, let $z = f(-\gamma)$.

2.2 Analysis

If $Z = (Z_1, Z_2, \ldots, Z_k)$ denotes the random vector z chosen in steps (R1) and (R2), the choice of probabilities in (R1) and (R2) ensures that $\mathbf{E}[Z_1] = 0$. So, we have from (4) that

$$\forall t,\ \mathbf{E}[Z_t] = 0. \tag{5}$$

The algorithm clearly rounds at least one edge permanently in each iteration (and removes all such edges from the forest F), and therefore terminates in polynomial time. We now analyze the expected revenue obtained from each bidder i, and prove that it is not too small.

Let $L_i^{(s)}$ denote the load on bidder i at the end of iteration s; the values $L_i^{(0)}$ refer to the initial values obtained by running the subroutine of [3] that obtains the forest F. Property (P2) shows that as long as i has degree at least two in the forest F, $L_i^{(s)}$ stays at its initial value $L_i^{(0)}$ with probability 1. (Recall that whenever we refer to F etc., we always refer to its subgraph containing those edges with x values in $(0, 1)$; edges that get rounded to 0 or 1 are removed from F.) In particular, if i never had degree one at the end of any iteration, then its final load equals $L_i^{(0)}$ with probability one, so the expected approximation ratio for i is one. So, suppose the degree of i came down to one at the end of some iteration s. Let the corresponding unique neighbor of i be j, let $\beta = b_{i,j}$, and suppose, at the end of iteration s, the total already-rounded load on i and the value of $x_{i,j}$ are $\alpha \geq 0$ and $p \in (0, 1)$ respectively. Note that j, α, β as well as p are all random variables, and that $L_i^{(s)} = \alpha + \beta p$; so,

$$\Pr[\alpha + \beta p = L_i^{(0)}] = 1.$$

Fix any j, α, β and p that satisfy $\alpha + \beta p = L_i^{(0)}$; all calculations from now on will be conditional on this fixed choice, and on all random choices made up to the end of iteration s. Property (5) and induction on the iterations show that the final load on i (which is now a random variable that is a function of the random choices made from iteration $s + 1$ onward) is:

$$\alpha,\ \text{with probability } 1 - p; \text{ and } \alpha + \beta,\ \text{with probability } p. \tag{6}$$

Let $B = B_i$ for brevity. Thus, the final expected revenue from i is $(1 - p) \cdot \min\{\alpha, B\} + p \cdot \min\{\alpha + \beta, B\}$; the revenue obtained from i in the LP solution is

$\min\{\alpha+\beta p, B\}$. So, by the linearity of expectation, the expected approximation ratio is the maximum possible value of

$$\frac{\min\{\alpha+\beta p, B\}}{(1-p)\cdot\min\{\alpha, B\}+p\cdot\min\{\alpha+\beta, B\}}.$$

It is easily seen that this ratio is 1 if $\alpha > B$ or if $\alpha+\beta < B$. Also note from (2) that $\beta \leq B$. Thus, we want the minimum possible value of the reciprocal of the approximation ratio:

$$r = \frac{(1-p)\alpha + pB}{\min\{\alpha+\beta p, B\}}, \tag{7}$$

subject to the constraints

$$p \in [0,1];\ \ \alpha, \beta \leq B;\ \ \alpha+\beta \geq B. \tag{8}$$

(Of course, we assume the denominator of (7) is nonzero. In the case where it is zero, it is easy to see that so is the numerator, in which case it follows trivially that $(1-p)\alpha + pB \geq (3/4)\cdot\min\{\alpha+\beta p, B\}$.)

We consider two cases, based on which term in the denominator of r is smaller:

Case I: $\alpha+\beta p \leq B$. Here, we want to minimize

$$r = \frac{(1-p)\alpha + pB}{\alpha+\beta p}. \tag{9}$$

Keeping all other variables fixed and viewing α as a variable, r is minimized when α takes one of its extreme values, since r is a non-negative rational function of α. From our constraints, we have $B-\beta \leq \alpha \leq B-\beta p$. Thus, r is minimized at one of these two extreme values of α. If $\alpha+\beta = B$, then $r = 1$. Suppose $\alpha = B-\beta p$. Then,

$$r = \frac{(1-p)\alpha + pB}{B}. \tag{10}$$

Since

$$\alpha = B-\beta p \geq B(1-p), \tag{11}$$

we have

$$r = \frac{(1-p)\alpha + pB}{B} \geq (1-p)^2 + p,$$

which attains a minimum value of $3/4$ when $p = 1/2$.

Case II: $\alpha+\beta p \geq B$. We once again fix all other variables and vary α; the extreme values for α now are $\alpha = B - Bp$ (with $\beta = B$) and $\alpha = B$. In the former case, the argument of Case I shows that $r \geq 3/4$; in the latter case, r is easily seen to be 1.

This completes the proof that our expected approximation ratio is at most $4/3$. Also, it is easy to derandomize the algorithm by picking one of the two possible updates in each iteration using the method of conditional probabilities; we will describe this in the full version. Thus we have the following theorem:

Theorem 1. *Given any feasible fractional solution to the LP-relaxation of the offline budgeted-allocation problem with the truncations (1) done without loss of generality, it can be rounded to a feasible integer solution with at least* 3/4*-th the value of the fractional solution in deterministic polynomial time. Therefore, the offline budgeted-allocation problem can be approximated to within* 4/3 *in deterministic polynomial time.*

3 Extensions

The following two extensions hold.

3.1 The case of bids being small w.r.t. budgets

Here we consider the case where for some $\epsilon \in [0, 1]$, we have for all i, j that $b_{i,j} \leq \epsilon B_i$. The only modification needed to the analysis of Section 2.2 is that (11) now becomes "$\alpha = B - \beta p \geq B(1 - \epsilon p)$", and that the function to minimize is $(1-p) \cdot (1-\epsilon p) + p$ instead of $(1-p)^2 + p$. This is again minimized at $p = 1/2$, giving $r \geq 1 - \epsilon/4$. Thus, the approximation ratio in this case is at most $4/(4-\epsilon)$.

3.2 The case of similar bids for any given bidder

We now study the case where for each i, all its nonzero bids $b_{i,j}$ are within some factor λ of each other, where $1 \leq \lambda \leq 2$. Note that different bidders can still have widely-differing bid values.

Consider the analysis of Section 2.2. In the trivial case where $\alpha = 0$, it easily follows from (6) that the approximation ratio for machine i is 1. So suppose $\alpha > 0$; then the additional constraint that

$$\beta \leq \alpha\lambda \tag{12}$$

must hold, by our assumption about the bid-values.

By a tedious proof along the lines of Section 2.2, it can be shown that we get a better-than-4/3 approximation if $\lambda < 2$. We will present the calculation-details in the full version. For now, we just focus on the case where $\lambda = 1$. Recall that we aim to minimize r from (7), subject to (8) and the constraint (12), i.e., $\alpha \geq \beta$. Let us first argue that if the minimum value of r is smaller than 1, then $\alpha = \beta$ at any minimizing point. To see this, assume for a contradiction that there is a minimizing pair (α, β) with $\alpha > \beta$, and observe that we may make the following three sets of assumptions w.l.o.g.: (i) if $\alpha = 0$ or $\alpha + \beta = B$, then $r = 1$: so, we may assume that $\alpha > 0$ and $\alpha + \beta > B$; (ii) if $\beta = B$, then $\alpha \geq \beta = B = \beta$ and we are done, so we can assume $\beta < B$; (iii) if $p = 0$, then $r = 1$, so we can take $p > 0$. Now, if we perturb as $\alpha := \alpha - \delta$ and $\beta := \beta + \delta/p$ for some tiny positive δ, then we stay in the feasible region and get a smaller value for r from (7), a contradiction. So, we can take $\alpha = \beta$, and have from (8) that $\alpha = \beta \geq B/2$.

We repeat the case analysis of Section 2.2. In Case I, the extreme value $\alpha = B/2$ gives $r = 1$. The other extreme value is $\alpha = B - \beta p = B - \alpha p$, i.e.,

$\alpha = B/(1+p)$. So, the r of (10) becomes $(1-p)/(1+p)+p$, whose minimum value is $2(\sqrt{2}-1)$. Similarly in Case II. Thus, $r \geq 2(\sqrt{2}-1)$, and taking the reciprocal, we see that the approximation ratio is $(\sqrt{2}+1)/2 \sim 1.207$.

Acknowledgment. I thank Anna Karlin for informing me about related work, and the APPROX 2008 referees for their helpful suggestions.

References

[1] Alexander A. Ageev and Maxim Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.*, 8(3):307–328, 2004.

[2] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In *SWAT*, pages 26–38, 2004.

[3] Yossi Azar, Benjamin Birnbaum, Anna R. Karlin, Claire Mathieu, and C. Thach Nguyen. Improved approximation algorithms for budgeted allocations. In *ICALP*, 2008.

[4] Niv Buchbinder, Kamal Jain, and Joseph Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*, pages 253–264, 2007.

[5] Deeparnab Chakrabarty and Gagan Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. Manuscript.

[6] Uriel Feige and Jan Vondrák. Approximation algorithms for allocation problems: Improving the factor of 1 - 1/e. In *FOCS*, pages 667–676, 2006.

[7] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, 2006.

[8] Rahul Garg, Vijay Kumar, and Vinayaka Pandit. Approximation algorithms for budget-constrained auctions. In *RANDOM-APPROX*, pages 102–113, 2001.

[9] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to Adwords. In *SODA*, pages 982–991, 2008.

[10] Sébastein Lahaie, David M. Pennock, Amin Saberi, and Rakesh V. Vohra. Sponsored search auctions. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 28, pages 699–716. Cambridge University Press, 2007.

[11] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.

[12] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.

[13] Aravind Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *FOCS*, pages 588–597, 2001.