On maximizing welfare when utility functions are subadditive

Uriel Feige
Microsoft Research and Weizmann Institute
urifeige@microsoft.com

October 31, 2005

Abstract

We consider the problem of maximizing welfare when allocating m items to n players with subadditive utility functions. Our main result is a way of rounding any fractional solution to a linear programming relaxation to this problem so as to give a feasible solution of welfare at least half that of the value of the fractional solution. This approximation ratio of $1/2$ improves over an $\Omega(1/\log m)$ ratio of Dobzinski, Nisan and Schapira [STOC 2005]. We also show an approximation ratio of $1 - 1/e$ when utility functions are fractionally subadditive. A result similar to this last result was previously obtained by Dobzinski and Schapira [Soda 2005], but via a different rounding technique that requires the use of a so called "XOS oracle".

The randomized rounding techniques that we use are *oblivious* in the sense that they only use the primal solution to the linear program relaxation, but have no access to the actual utility functions of the players. This allows us to suggest new incentive compatible mechanisms for combinatorial auctions, extending previous work of Lavi and Swamy [FOCS 2005].

1 Introduction

We consider the following problem. There are m items and n players. A feasible allocation allocates every item to at most one player. For every player P_i, her utility w_i depends only on the set S of items that she receives. Utility functions are nonnegative, monotone and subadditive. That is,

$$0 \leq w_i(S) \leq w_i(S \cup S') \leq w_i(S) + w_i(S')$$

for every i, S, S'. Let S_i be the set of items allocated to player i. The goal is to find a feasible allocation that maximizes social welfare, namely, maximizes $\sum_i w_i(S_i)$.

Dobzinski, Nisan and Schapira [4] considered the following linear programming relaxation of the problem, that we call the *welfare maximizing LP*. $x_{i,S}$ is intended to be an indicator variable that specifies whether player i gets set S.
Maximize $\sum_{i,s} x_{i,S} w_i(S)$ *subject to:*

- Item constraints: $\sum_{i,S|j \in S} x_{i,S} \leq 1$ for every item j.
- Player constraints: $\sum_S x_{i,S} \leq 1$ for every player i.
- Nonnegativity constraints: $x_{i,S} \geq 0$.

This linear program can be solved optimally in polynomial time, assuming that there is an oracle that can answer *demand queries*: which set of items would a player want to buy given a setting of prices to the individual items? See [4] for more details.

In [4] it is shown that any solution to the linear program can be rounded to give a feasible allocation of welfare at least $\Omega(1/\log m)$ of the value of the fractional solution. (In fact, more careful analysis of the rounding technique of [4] establishes an approximation ratio of $\Omega(\frac{\log \log m}{\log m})$. See Section 2.2.) Moreover, Conjecture 1 in [4] states that this is best possible.

In this paper we refute the conjecture from [4]. Our main result (described in Section 3.2) is a new randomized rounding technique that given any fractional solution to the welfare maximizing LP (whose value we denote by $w(LP)$), produces a feasible (integer) allocation of expected welfare at least $w(LP)/2$. We show that a wide variety of other rounding techniques (including all other rounding techniques described in this paper) fail to give a constant approximation ratio (see Section A.4).

Another result of this paper is that if the utility functions are further restricted to be *fractionally subadditive* (see definition is Section 1.1), then one can round the solution to the LP so as to obtain a feasible allocation of welfare at least $(1 - 1/e)w(LP)$. This last result is not new. An approximation ratio of $1 - 1/e$ was previously obtained in [5] for a class of utility functions known as XOS, and as we show (see Proposition 9), the class XOS is the same as the class of fractionally subadditive utility functions. Nevertheless, our result uses a rounding technique that is inherently different from that of [5], and may be of independent interest. See discussion below.

We note that both the $1 - 1/e$ approximation ratio for fractionally subadditive utility functions and the $1/2$ approximation ratio for subadditive utility functions are best possible, in the sense that they match the integrality ratio of the linear program for the corresponding cases, up to low order terms.

An interesting feature of our rounding techniques is that they are *oblivious* in the following sense. As input, they use only the values $x_{i,S}$ of an arbitrary feasible solution to the welfare maximizing LP. They receive absolutely no information about the actual utility functions of the players. (It may appear that the solution to the LP provides implicit information about the utility functions, but this is not the case, because this solution is only required to be (fractionally) feasible, but not optimal.) As output, our rounding techniques produce a feasible integer allocation, or rather, a distribution over feasible allocations (because oblivious rounding techniques are inherently randomized). The performance guarantee is per player. Every player is guaranteed to recover in expectation at least a certain fraction (the approximation ratio) of the utility offered to the same player under the given solution to the welfare maximizing LP.

Once a solution to the welfare maximizing LP is given, the use of oblivious rounding techniques requires no further interaction with the players. This circumvents the issue of how utility functions are represented (see discussion in Section A.2), and may also be of practical significance as it may reduce communication costs. Moreover, this is in agreement with the principle that players cannot always be trusted to report their true utility functions.

We remark that the other known rounding technique achieving an approximation ratio of $1 - 1/e$ for fractionally subadditive utility functions [5] is not oblivious. Implementing the rounding technique of [5] requires some detailed knowledge of the utility functions of the players, given in terms of a so called "XOS oracle".

The welfare maximization problem sometimes comes up in game theoretic settings. In these settings, one would like to have mechanisms that provide incentives to players to report their true utilities. Extending an approach of [8], we show how this can be done when players wish to maximize their expected benefit (namely, expected utility minus payment). Our mechanism is based on an oblivious rounding technique, and recovers at least an $\Omega(\log\log m/\log m)$ fraction of the total optimum welfare when utility functions are subadditive. This result (which partly answers an open question of [5]) appears in section A.7 in the appendix.

1.1 Classes of utility functions

In this work, we consider subadditive utility functions. However, to put this work in perspective, we discuss here some other classes of utility functions. For more details, see [9].

We denote a utility function by w, and sets of items by uppercase letters. As a convention in this work, utility functions are nonnegative and monotone. That is, $w(S) \geq 0$ for every S, and $w(S \cup T) \geq w(S)$ for every S, T. Another common convention regarding utility functions is that the utility of the empty set is 0, though this convention is not used in our work. It will be informative to consider the following classes of utility functions.

1. **Additive** (a.k.a. linear). $w(S) = \sum_{j \in S} w(j)$.

2. **Submodular**. $w(S \cup T) + w(S \cap T) \leq w(S) + w(T)$, for every S, T. A useful equivalent characterization of submodular utility functions is as those utility functions in which the marginal utility of an item decreases as sets become larger (inclusion-wise). That is, for every item j and sets $T \subset S$, $w(j \cup S) - w(S) \leq w(j \cup T) - w(T)$.

3. **Fractionally subadditive**. $w(S) \le \sum \alpha_i w(T_i)$ with $0 \le \alpha_i \le 1$ for all i, whenever the following condition holds: for every item $j \in S$, $\sum_{i|j\in T_i} \alpha_i \ge 1$. (Namely, if the sets T_i form a "fractional cover" of S, then the sum of their utilities weighted by the corresponding coefficients is at least as large as that of S.) The class of fractionally subadditive utility functions is the same as the class XOS introduced in [9]. This fact will be proved in Proposition 9.

4. **Subadditive** (a.k.a. complement free). $w(S \cup T) \le w(S) + w(T)$, for every S, T.

It can be shown that every linear utility function is submodular, every submodular utility function is fractionally subadditive, and every fractionally subadditive utility function is subadditive. To illustrate the difference between the above classes, consider a set $S = \{1, 2, 3\}$ of three items, and assume that the utility of every proper subset of it (containing either one or two items) is 1. What constraints does this place on $w(S)$? For arbitrary utility functions, the only constraint is that $w(S) \ge 1$. For subadditive utility functions, we have in addition $w(S) \le 2$, because $w(S) \le w(\{1, 2\}) + w(3)$. For fractionally subadditive utility functions we have $w(S) \le 3/2$, because $w(S) \le (w(\{1, 2\}) + w(\{2, 3\}) + w(\{1, 3\}))/2$. For submodular utility functions we have $w(S) = 1$, because $w(S)+w(1) \le w(\{1, 2\})+w(\{1, 3\})$. The function cannot be a linear utility function at all, because $w(\{1, 2\}) \ne w(1) + w(2)$.

The class of utility functions based on *set cover* problems serves as a useful example to show the distinction between subadditive and fractionally subadditive utility function. Let $T_1, \dots, T_k$ be some ground sets whose union contains all items. Then $w(S) = \min t$ such that there are t sets satisfying $S \subset \bigcup_{j=1}^{t} T_{i_j}$ is a subadditive utility function, but in general is not fractionally subadditive.

1.2 Integrality gap for the subadditive case

It is shown in [4] that it is impossible to get an approximation ratio strictly better than 1/2 with only polynomial amount of communication with the players. In our context, it may be more informative to view this as an integrality gap for the LP, or as a hardness of approximation result.

Proposition 1 *For every $\epsilon > 0$, it is NP-hard to approximate the maximum welfare (when players have subadditive utility functions) within a factor of $1/2 + \epsilon$.*

Proof: It is known [1] that for every $\epsilon > 0$ there is an $\alpha > 0$ such that it is NP-hard to distinguish between "yes instances" in which a graph has an independent set of size αn and "no instances" in which every independent set is of size at most $\epsilon\alpha n$. Let the edges of an input graph be the items, let the number of players be αn, and let $w(S) = w_i(S) = 2$ if there is some vertex such that S contains all edges incident with it, and $w(S) = 1$ otherwise. This utility function w is subadditive. On yes instances the maximum welfare is $2\alpha n$ (by giving each player the edges incident with some vertex of a maximum independent set), and on no instances it is at most $(1 + \epsilon)\alpha n$. □

We remark that it is known (and was rediscovered multiple times) that without subadditivity the maximum welfare cannot be approximated even within factors close to $\sqrt{m}$ (essentially by the same proof as above, but setting $w(S) = 0$ in the "otherwise" case).

Observe that for a clique on $2n$ vertices and utility functions as in Proposition 1, the optimal allocation has welfare $n+1$, whereas the LP has a feasible fractional solution of value $2n$ (e.g., by having all $x_{i,v} = 1/2n$, where v is a shorthand notation for the set of edges incident with vertex v). This establishes an integrality gap of $1/2 + 1/2n$ for the LP.

2 Basic oblivious rounding techniques

This section contains known results [4], but our presentation is based on oblivious rounding techniques, and hence will lead more naturally to our new results.

2.1 One step randomized rounding

Perhaps the simplest randomized rounding scheme for the LP is as follows. For each item j we have the constraint $\sum_{i,S|j\in S} x_{i,S} \leq 1$. In the randomized rounding procedure, allocate item j to at most one player, by selecting player i with probability $\sum_{S|j\in S} x_{i,S}$. This gives a feasible allocation. A player gets an item with probability equal to the fractional value of the item assigned to the player under the solution to the LP. When utility functions are additive, then in expectation a player's utility will be similar to its contribution to $w(LP)$. However, when utility functions are not additive, this is far from true.

Consider for example the case were $n = \sqrt{m}$. For every player i independently, define a random partition of all items into n sets $T_1^i, \ldots, T_n^i$. The utility functions are $w_i(S) = \max_{j=1}^n |S \cap T_j^i|$. This utility function is fractionally subadditive. A feasible fractional solution to the LP assigns to player i each of the corresponding sets T_j^i with probability $1/n$. Hence the fractional utility to player i is $\frac{1}{n}\sum_{j=1}^n |T_j^i| = m/n = n$, and the total fractional welfare is $n^2 = m$. However, the randomized rounding procedure described above is unlikely to ever allocate more than $O(\log n)$ items from the same set T_j^i to player i, and hence the total welfare will be $O(n \log n)$, which is a factor of $\Omega(\sqrt{m}/\log m)$ worse than the fractional solution.

2.2 Two step randomized rounding

We present here an oblivious two step randomized rounding technique. It is a straightforward variation of the rounding technique of [4] (which was not oblivious).

1. Each player chooses at most one set of items, where player i chooses set S with probability $x_{i,S}$. In expectation, the welfare of a player does not change. However, her chosen set might intersect with sets chosen by other players. Hence, the solution might not be feasible. We call this a tentative allocation.

2. For each item j, if it is allocated to several players under the tentative allocation, choose uniformly at random which of these players gets the item j. This results in a feasible solution, called the final allocation.

We now analyse the quality of the final solution. Standard probabilistic arguments (see Proposition 10 in Section A.3 in the appendix) imply that with high probability (say, probability $1 - 1/m$), no item belongs to more than $k = O(\log m/\log\log m)$ players in the tentative allocation. (In [4] only a weaker bound of $O(\log m)$ bound is claimed, but the basic idea is the same.) Hence when computing the final allocation to a player (the second step of the rounding) every item of the tentative allocation is included independently with probability at least $1/k$. Now is the point when we use subadditivity of the utility functions, namely, Proposition 2, which together with monotonicity implies the desired bound. (Monotonicity is appealed to because the probability per item in Proposition 2 is exactly $1/k$, whereas we are interested in the case where the probability is at least $1/k$.)

Proposition 2 *Let $k \geq 1$ be integer and let w be an arbitrary subadditive utility function. For a set S, pick a random subset $S' \subset S$ by picking each item of S independently at random with probability $1/k$. Then $E[w(S')] \geq w(S)/k$.*

Proof: Color independently at random each item of S with one of k colors. This gives k mutually disjoint subsets $S_1, \ldots, S_k$, where every such subset is distributed exactly like S'. By subadditivity, $\sum_i w(S_i) \geq w(S)$. Now the proposition follows from the linearity of the expectation. □

Summing up, the two step randomized rounding procedure gives the following guarantee to every player i. The expected utility of her tentative set is exactly $\sum_S x_{i,S} w_i(S)$, namely, her contribution to the fractional solution of the LP. Thereafter, with overwhelming probability (say, $1 - 1/m$), no bad event happens, in the sense that no item is in more than k tentative sets. Conditioned on no bad event happening, the expected utility of her final set is at least a $1/k$ fraction of the utility of her tentative set, by Proposition 2. By linearity of expectation, it follows that the expected welfare of the allocation delivered by the two step rounding technique is at least a $(1 - 1/m)/k = \Omega(\log \log m / \log m)$ fraction of the value of the fractional solution to the welfare maximizing LP.

2.3 Fractionally subadditive utility functions

For fractionally subadditive utility functions we can use a strengthening of Proposition 2. The difference between the two propositions is the removal of the requirement for statistical independence among items.

Proposition 3 *Let $k \geq 1$ be integer and let w be an arbitrary fractionally subadditive utility function. For a set S, consider a distribution over subsets $S' \subset S$ such that each item of S is included in S' with probability at least $1/k$. Then $E[w(S')] \geq w(S)/k$.*

Proof: Let p_i be the probability that set S_i is chosen. Then $\sum p_i = 1$, and $k \sum p_i S_i$ fractionally covers S. Hence also $\sum \min[1, kp_i] S_i$ fractionally cover S, and by fractional subadditivity, $w(S) \leq \sum \min[1, kp_i] w(S_i) \leq k \sum p_i w(S_i) = kE[w(S')]$, as desired. □

An alternative proof of Proposition 3 follows from the equivalence between fractionally subadditive and XOS utility function (Proposition 9), but is omitted here.

Consider now the two step rounding procedure of Section 2.2. From the point of view of player i, step 1 of the other players can be viewed as being part of step 2, as follows.

1. Player i chooses a tentative set S_i.
2. (a) All other players choose their tentative sets.

(b) Item $j \in S_i$ is allocated to player i with probability $1/(n_j + 1)$, where n_j is the number of other players who have item j in their tentative sets.

In expectation, in step 1 player i looses nothing compared to the fractional solution. Now consider steps 2(a) and 2(b) combined. The expected value of n_j is at most 1, due to the item constraints of the LP. It is not hard to see that this implies that the expected value of $1/(n_j + 1)$ is at least $1/2$. It now follows from Proposition 3 that for fractionally subadditive utility functions, the rounded solution is expected to recover at least half the value of the fractional solution.

3 Improved approximation ratios

In this section we show how to improve over the approximation ratios presented in Section 2. First we give an overview of our approach.

Recall the two step randomized rounding technique. In the first step, each player is assigned at most one tentative set. In the second step we resolve contention: if several players have the same item j in their tentative set (in which case, we view them as players *competing* for j), then one of the competing players is chosen uniformly at random and gets item j.

The key to our improved approximation ratios is a change in the second step. Rather that allocating item j uniformly at random, we attempt to allocate it to the player who will derive the highest marginal utility from item j.

The above principle is not new. It was used by Fleischer, Goemans, Mirrokni and Sviridenko [6] in a setting in which utility functions are additive (and then the implementation of this principle is straightforward). For XOS utility functions (maximum of additive utility functions), Dobzinski and Schapira [5] showed that if one can determine for every player which additive utility function maximizes the utility of its tentative set, then one can use this information to obtain an approximation ratio of $1-1/e$. Moreover, this is best possible, in the sense that the LP relaxation has an integrality gap tending to $1-1/e$. As we noted in Section 1.1, XOS utility functions are essentially the same as fractionally subadditive utility functions, and hence the results of [5] give an approximation ratio of $1-1/e$ for fractionally subadditive utility functions.

In contrast to [5] (and to [6]), we present rounding techniques that are *oblivious*. That is, our goal is to give item j to the player that would derive maximum marginal utility from it, but we wish to achieve this goal without knowing anything about the utility functions of the players. Of course, this cannot be done. Nevertheless, we design randomized oblivious rounding techniques that achieve the best possible approximation ratios (in the sense that they match the integrality gap of the LP). For fractionally subadditive utility functions, the new aspect of our results is the fact that the rounding techniques are oblivious. For subadditive utility functions (our main result), an even more important aspect is the dramatic improvement in approximation ratio, matching the NP-hardness result (Proposition 1) and the integrality bound of the welfare maximizing LP in this case.

3.1 Fractionally subadditive utility functions

In Section 2.3 we showed a factor 1/2 approximation for the case of fractionally subadditive utility functions. In this Section we show an improved rounding procedure with approximation ratio $1-1/e$. A different way of obtaining a similar approximation ratio was previously shown in [5].

We first show that the two step randomized rounding procedure of Section 2.2 does not produce an approximation ratio better than 1/2. Consider the case that there are two players and a set S of m items. Player P_1 has an additive utility function, with each item having value 1. Player P_2 has a utility function that is the maximum of $\sqrt{m}$ additive utility functions. The ith additive utility function gives value $N \gg \sqrt{m}$ for each item in the set S_i, where S_i contains those items numbered $(i-1)\sqrt{m}+1$ to $i\sqrt{m}$.

An optimal solution to the LP sets $x_{1,S} = 1 - 1/\sqrt{m}$, and $x_{2,S_i} = 1/\sqrt{m}$ for every i. All other variables are 0. The value of the LP is $(1-1/\sqrt{m})m + \sqrt{m}\frac{1}{\sqrt{m}}\sqrt{m}N \simeq N\sqrt{m}$. However, the two step rounding procedure will produce a solution of value roughly $N\sqrt{m}/2$.

As a precursor to our improved rounding technique for fractionally subadditive utility functions, we consider first the special case where there are only two players. For this we suggest the following two-player rounding procedure.

1. Each player chooses at most one set of items, where player i chooses her *tentative* set S with probability $x_{i,S}$.
2. Let S_i denote the tentative set chosen by player i, for $i \in \{1,2\}$. For every item j independently do the following.
   (a) If $j \in S_1 \setminus S_2$, allocate j to player 1.
   (b) If $j \in S_2 \setminus S_1$, allocate j to player 1.
   (c) If $j \in S_1 \cap S_2$, then allocate j to player 1 with probability $\frac{\sum_{S|j\in S} x_{2,S}}{\sum_{S|j\in S} x_{1,S} + \sum_{S|j\in S} x_{2,S}}$ and to player 2 with probability $\frac{\sum_{S|j\in S} x_{1,S}}{\sum_{S|j\in S} x_{1,S} + \sum_{S|j\in S} x_{2,S}}$.
   (d) If $j \notin S_1 \cup S_2$, allocate j arbitrarily (this will not be used in our analysis of the approximation ratio).

Proposition 4 *For every player i, if her utility function is fractionally subadditive, then the expected utility of the random set allocated to the player under the above two-player rounding technique is at least three quarters of the utility allocated to her by the fractional solution of the linear program, namely, at least $\frac{3}{4}\sum_S x_{i,S} w_i(S)$.*

Proof: By symmetry, it suffices to prove the proposition with respect to player 1. The expected utility of the random tentative set S_1 that player 1 receives in step 1 is the same as the utility allocated to her by the LP. However, some items of S_1 might be given to player 2, if these items happen also to be in S_2, and moreover, step 2(c) allocates them to player 2. Hence an item $j \in S_1$ is given to player 1 with probability

$$1 - (\sum_{S|j\in S} x_{2,S}) \frac{\sum_{S|j\in S} x_{1,S}}{\sum_{S|j\in S} x_{1,S} + \sum_{S|j\in S} x_{2,S}} \geq 3/4$$

where the inequality follows from the fact that $\sum_{S|j\in S} x_{2,S} + \sum_{S|j\in S} x_{2,S} \leq 1$ (the item constraints). Now the proof follows from Proposition 3. $\square$

We now consider the case when the number of players is $n > 2$ (due to space limitations the proofs are moved to the appendix). In Section A.5 we present an oblivious rounding procedure with approximation ratio no worse than $1 - (1 - 1/n)^n$. Its running time is polynomial in m (the number of items) whenever the number of players n is constant. When n is not a constant, then $1-(1-1/n)^n$ tends to $1-1/e$ (from above). In Section A.6 we present an oblivious rounding technique that achieves an approximation ratio no worse than $1 - 1/e$, and whose running time is polynomial regardless of the number of players.

3.2 Subadditive utility functions

In Section A.4 in the appendix we present an instance based on the set cover utility function (which is subadditive but not fractionally subadditive) for which the two step randomized rounding procedure of section 2.2 does not produce a constant approximation ratio (and neither do all other rounding techniques described in this paper).

We now present a rounding procedure for the LP that has an approximation ratio of 1/2 when utility functions are subadditive. It will involve an object called the *guiding graph* which is not of polynomial size. Later we shall show how the rounding procedure can be implemented in expected polynomial time.

1. **The input.** The input to our rounding procedure is an arbitrary (not necessarily optimal) fractional (primal) solution to the LP. Neither the utility functions of the players nor the value of the solution (denoted by $w(LP)$) are needed as part of the input.

2. **The guiding graph.** Consider an arbitrary regular bipartite graph G of degree n and girth g, where g is sufficiently large compared to the number of players n and the number of items m. The graph G is called the *guiding graph*.

3. **Edge coloring.** The edges of every n-regular bipartite graph can be partitioned into n matchings (an edge coloring with n colors). Partition the edges of G into n matchings. Player i controls all edges of matching M_i.

4. **Random edge labelling.** For every $i \in \{1, \ldots, n\}$ and every edge $(u, v) \in M_i$, label the edge independently at random. The label of the edge is a set $S_{(u,v)}$ of items, where set S is chosen as the label with probability $x_{i,S}$.

5. **The item subgraphs.** We derive from G in combination with the edge labelling m edge induced subgraphs, one for every item. Subgraph G_j is obtained by keeping in G those edges (u, v) whose label satisfies $j \in S_{(u,v)}$, and removing all other edges.

6. **Tree property.** As we shall explain later, we may assume that every connected component of G_j is a tree.

7. **Edge orientation.** For every subgraph G_j and for every connected component C in G_j that is a tree, orient the edges in C such that every edge of C points in at least one direction, and every vertex of C has at most one edge pointing to it. For concreteness,

here is an explicit way of doing so. If the diameter of C is even, say $2d$, then C (being a tree) has a unique central vertex v of distance at most d from every other vertex. Orient all edges away from v. If the diameter of C is odd, say $2d + 1$, then C has a unique central edge (u, v) of distance at most d from all vertices of C. Orient every edge other than (u, v) away from (u, v), and keep the edge (u, v) bi-directional.

8. **Random center.** Pick a vertex $u \in G$ uniformly at random. We shall use u_i to denote the neighbor of u connected by edge $(u, u_i) \in M_i$. The set labelling the edge (u, u_i) will be called the *tentative set* of player i.

9. **Item allocation.** For every item j, if some edge (u, u_i) in G_j points at u, then this edge must be unique, and item j is allocated to the player i. (If no edge in G_j points at u, allocate item j to an arbitrary player. This will not be used in the analysis.) Observe that a player i may receive item j only if item j belongs to her tentative set, as otherwise edge (u, u_i) is not in subgraph G_j. This completes the description of the final sets $S_1, \ldots, S_n$ of items allocated to each player.

Theorem 5 *The welfare of the final solution found by the guiding graph rounding technique described above is in expectation at least $w(LP)/2$.*

The proof of Theorem 5 is a consequence of the following three propositions.

Proposition 6 *Fix an arbitrary fractional solution to the LP, and consider an arbitrary vertex u in the guiding graph G. Then for every $j \in \{1, \ldots, m\}$, the probability (over the choice of random edge labelling) that the connected component of vertex u in graph G_j contains a cycle (is not a tree) tends to 0 as the girth g of G tends to ∞.*

Proof: For player i and item j, let $p_i = \sum_{S|j \in S} x_{i,S}$. By the item constraints, $\sum p_i \leq 1$. Let us use ϵ to denote $\min_{i|p_i \neq 0}$. Assume first that the girth g is infinite, and hence that the connected component of u in G_j is a tree. Let us upper bound the expected size of this tree, where probability is taken over choice of random edge labels. We develop the connected component $G_j(u)$ in breadth first search fashion, starting at u. The expected degree of u is $\sum p_i \leq 1$. Thereafter, for every vertex already in the connected components, the expected number of children it has is at most $\sum p_i - \epsilon \leq 1 - \epsilon < 1$. The distribution of connected components containing v behaves like a branching process with at most $(1 - \epsilon)$ expected children at each node. By linearity of expectation, an upper bound N on the expected number of nodes generated by such a process can be derived by the recurrence relation $N \leq 1 + (1 - \epsilon)N$, implying $N \leq 1/\epsilon$. Hence the expected size of $G_j(u)$ is at most $1 + 1/\epsilon$. By Markov's inequality, the probability that its size exceeds $2k/\epsilon$ is at most $1/k$.

Observe that for the above analysis, all that is needed is that the girth of graph G is larger than $2k/\epsilon$, rather that the girth is infinite. Hence the probability the $G_j(u)$ is a tree is at least $1 - 2/g\epsilon$, which tends to 1 as g tends to ∞. $\square$

Proposition 7 *Consider an arbitrary vertex $u \in G$. Then the expectation (over choice of random edge labels) of the sum of utilities of the respective tentative sets satisfies:*

$$E[\sum_{i \in \{1,\ldots,n\}} w_i(S_{(u,u_i)})] = w(LP)$$

Proof: Set S labels edge (u, u_i) with probability $x_{i,S}$. Hence the expected sum of utilities satisfies:

$$E[\sum_{i\in\{1,\ldots,n\}} w_i(S_{(u,u_i)})] = \sum_i \sum_S x_{i,S} w_i(S) = w(LP)$$

□

Proposition 8 *Consider an arbitrary player i and an arbitrary set S such that $x_{i,S} > 0$, and an arbitrary labelling of the guiding graph. Then conditioned on the center u chosen such that:*

1. *S is the tentative set of player i,*
2. *for every $j \in S$, the connected components of u in all subgraphs G_j are trees,*

the expected utility of the final set of items allocated to player i satisfies $E[w_i(S_i)] \geq w_i(S)/2$. Here probability is taken over choice of center vertex u.

Proof: Let $M_{i,S}$ be the set of edges controlled by player i that are labelled by set S and for which the connected components of u in subgraphs G_j are trees, for all $j \in S$. Then one may choose a center vertex with the probability distribution specified by the proposition by first picking at random an edge $(u, v) \in M_{i,S}$, and then picking the center vertex to be one of its endpoints. Let S_u be the final set that player i receives when u is the center, and let S_v be the final set that player i receives when v is the center. Observe that every item $j \in S$ must be in either S_u or S_v, depending on the orientation of the edge (u, v) in the subgraph G_j. Hence $S \subset S_u \cup S_v$. Now we use subadditivity of the utility function to conclude that $w_i(S_u) + w_i(S_v) \geq w_i(S)$. Summing over all edges of $M_{i,S}$ and averaging, the proposition follows. □

The three propositions above imply Theorem 5.

Proof: In the limit, as g tends to ∞, Proposition 6 implies that condition 2 of Proposition 8 holds with probability 1. Then Proposition 8 implies that $E[w_i(S_i)] \geq w_i(S_{(u,u_i)})/2$, which together with Proposition 7 implies:

$$E[\sum_{i\in\{1,\ldots,n\}} w_i(S_i)] \geq w(LP)/2$$

□

To complete our discussion of the rounding procedure, let us explain how it can be implemented efficiently. The guiding graph G has degree n and girth g, which implies that it must contain at least $n^{g/2}$ nodes. Hence graph G might be too large so as to be represented efficiently. Luckily, this is not needed. The rounding procedure uses only parts of the guiding graph, namely the connected components $G_1(u), \ldots, G_m(u)$. As shown in Proposition 6, their expected size is $O(1/\epsilon)$, where $\epsilon = \min_{i,j} \sum_{S|j\in S} x_{i,S}$ (conditioned on $\sum_{S|j\in S} x_{i,S} > 0$). The relevant portion of the guiding graph (the portion that contains the union of $G_1(u), \ldots, G_m(u)$) can be generated on demand using (for example) a breadth first search procedure starting at u, and assigning labels only to those edges that are not cut off from u by labels of previously assigned edges. This leads to an expected polynomial time rounding procedure when $1/\epsilon \leq \text{poly}(n, m)$. In fact, more careful analysis shows that this last condition is not needed, because for every j, the average value (over choice of $k \in \{1, \ldots, n\}$) of $\sum_{i \neq k; S | j \in S} x_{i,S} \leq 1 - 1/n$. This careful analysis can be avoided if one accepts a modest loss in the approximation ratio. Pick some new small $\epsilon > 0$, and scale the values of all variables in the solution to the LP by a factor of $1 - \epsilon$. The solution remains feasible, and its value decreased by a factor of only $1 - \epsilon$. Now $1/\epsilon$ is simply a constant. Hence the total expected number of nodes in $G_1(u), \ldots, G_m(u)$ is $O(m/\epsilon)$, and each vertex has degree n (in G). Altogether, the expected number of edges of the guiding graph which need to be visited is $O(nm/\epsilon)$, and the rounding procedure runs in expected polynomial time.

4 Some final remarks

The most innovative result presented in this paper is the factor 2 approximation for the case when utility functions are subadditive. In an effort to present this main result in an easily readable way, some other results are only mentioned in the main body of the paper, but the details are moved to the appendix. This section briefly comments on some of the results appearing in the appendix.

1. Proposition 9 shows the equivalence between fractionally subadditive and XOS utility functions. The proof is not difficult, but apparently this equivalence was not observed before. The class XOS was studied in [9, 4, 5], and the class of fractionally subadditive utility functions plays a central role in economics theory [2, 10].

2. Section A.4 presents an example showing that many rounding techniques would not give a constant approximation ratio for the case of subadditive utility functions. This example was not easy to design. In fact, for the two step rounding technique and subadditive utility functions, there still remains a gap between the positive result of $\Omega(\log \log m / \log m)$ approximation ratio (Section 2.2) and the negative example of Section A.4 showing an $O(\sqrt{\log \log m / \log m})$ upper bound on its performance.

3. In Sections A.5 and A.6 we present oblivious rounding techniques for fractionally subadditive utility functions when there are more than two players. These are interesting rounding techniques and not merely extensions of the simple rounding technique that was presented in Section 3.1 for the case of two players. Section A.5 uses flow techniques, whereas Section A.6 uses probabilistic insights.

4. In Section A.7 we present an incentive compatible mechanism based on oblivious rounding techniques. This extends a previous approach of [8]. This result has some shortcomings, outlined towards the end of Section A.7.

Acknowledgements

The author has benefited from discussions with Konstantin and Yuri Makarychev, Yishai Mansour, Vahab Mirrokni and Noam Nisan . Special thanks to Kamal Jain who introduced the author to the class of fractionally subadditive utility functions and their properties. Sections A.5 and A.7 were motivated by comments given by Shahar Dobzinski and Michael Schapira on an earlier version of this manuscript.

References

[1] S. Arora, S. Safra. "Probabilistic Checking of Proofs: A New Characterization of NP." *Journal of ACM* , 45(1):70–122, 1998.

[2] O. N. Bondareva. "Some applications of linear programming to cooperative games." *Problemy Kibernetiki*, 10:119–139, 1963.

[3] E. Clarke. "Multipart pricing of public goods." *Public Choice* 8:17–33, 1971.

[4] S. Dobzinski, N. Nisan, M. Schapira. "Approximation algorithms for combinatorial auctions with complement free bidders." *Proceedings of 37th STOC*, 610–618, 2005.

[5] S. Dobzinski, M. Schapira. "An Improved approximation algorithm for combinatorial auctions with submodular bidders." To appear in the proceedings of SODA 2006.

[6] L. Fleischer, M. Goemans, V. Mirrokni, M. Sviridenko. "Tight approximation algorithms for maximum general assignment problems." To appear in the proceedings of SODA 2006.

[7] T. Groves. "Incentives in teams." *Econometrica* 41:617–631, 1973.

[8] R. Lavi, C. Swamy. "Truthful and near-optimal mechanism design via linear programming." To appear in the proceedings of FOCS 2005.

[9] B. Lehmann, D. Lehmann, N. Nisan. "Combinatorial auctions with decreasing marginal utilities." *Proceedings of EC'01*, 18–28, 2001.

[10] L. S. Shapley. "On balanced sets and cores." *Naval research logistics quarterly*, 14:453–460, 1967.

[11] W. Vickrey. "Counterspeculation, auctions and competitive sealed tenders." *Journal of Finance*, 1961.

A Appendix

A.1 More on utility functions, and the Bondareva-Shapley Theorem

We remark that submodular utility functions are as those that are subadditive in a stronger sense. Even if a player already has some items, her marginal utility with respect to the rest of the items remains subadditive.

Another class of utility functions that is considered in [9, 4, 5] is called XOS. This is the class of utility functions that can be expressed as a maximum of linear utility functions. If one allows the number of linear utility functions in the XOS representation to be arbitrarily large (exponential in the number of items), then Proposition 9 shows that the class XOS is the same as the class of fractionally subadditive utility functions.

Proposition 9 *A utility function is in the class XOS if and only if it is fractionally subadditive.*

Proof: Let w be an XOS utility function. Then by definition there are additive utility functions $w_1, w_2, \ldots$ such that for every set of items S, $w(S) = \max_j\{w_j(S)\}$. Now we show that w is fractionally subadditive. Consider an arbitrary fractional cover of a set S by sets T_i, namely S is covered by $\sum \alpha_i T_i$. For set S, let j^* be such that $\max_j\{w_j(S)\} = w_{j^*}(S)$. Since w_{j^*} is an additive function and the T_is form a cover of S, it follows that $\sum \alpha_i w_{j^*}(T_i) \geq w_{j^*}(S)$. But for every T_i we have $w(T_i) \geq w_{j^*}(T_i)$. Putting these inequalities together we have:

$$w(S) = \max_j\{w_j(S)\} = w_{j^*}(S) \leq \sum \alpha_i w_{j^*}(T_i) \leq \sum \alpha_i \max_j\{w_j(T_i)\} = \sum \alpha_i w(T_i).$$

To show that every fractionally subadditive utility function is in the class XOS consider the following linear program associated with a set S and a utility function w.

minimize $\sum_T x_T w(T)$
subject to:

- $\sum_{T|j\in T} x_T \geq 1$ for every item $j \in S$
- $x_T \geq 0$ for every set T.

The fact that w is fractionally subadditive implies that the optimum of the above LP is at least $w(S)$. Hence, it is exactly $w(S)$ (by setting the variable $x_S = 1$ and all other $x_T = 0$).

The dual of the above LP is:
maximize $\sum_{j\in S} y_j$
subject to:

- $\sum_{j\in T} y_j \leq w(T)$ for every set T.
- $y_j \geq 0$ for every item j.

By linear programming duality, the value of the dual is also $w(S)$. The optimal values y_j^* of the dual variables define a linear function w_S in a natural way, where the value of a set T is $w_S(T) = \sum_{j\in T} y_j^*$.

Now w can be represented as an XOS utility function using $w = \max_S\{w_S\}$. Indeed, for every set T we have that for every S, $w_S(T) \leq w(T)$ (a consequence of the dual linear program for S), and $\max_S\{w_S(T)\} \geq w_T(T) = w(T)$ (a consequence of the equality between primal and dual). □

The reader may note the Proposition 9 is a straight-forward variation of the Bondareva-Shapley Theorem [2, 10]. This theorem is sketched below for completeness.

Suppose that there is a set S of n players that jointly receive some service. For each set of players $T \subseteq S$ there is a cost $c(T) \geq 0$ for providing the service to that set. A *cost sharing scheme* f allocates nonnegative shares of the cost $c(S)$ to each of the players so that the service is payed for. The cost sharing scheme is said to be in the *core* if no subset T of players has incentive to defect from S, receive the service on their own, and pay for it the cost of serving T alone. That is, for every $T \subset S$, $\sum_{i\in T} f(i) \leq c(T)$. The Bondareva-Shapley theorem says that the core is nonempty if and only the cost function c is fractionally subadditive with respect to S.

A.2 Single player problems

We are interested in this work in efficient (polynomial time) algorithms. Intuitively, one may imagine that the complexity of the allocation problem is the result of having multiple players with conflicting wishes. But in fact, even single player problems might involve computationally difficult tasks. We elaborate on this below. (More details can be found in [9, 4, 5].)

A utility function specifies a nonnegative value to every set of items. Representing a utility function as a table requires space exponential in the number of items m. This representation is incompatible with standard notions of efficient algorithms. As a way of coping with this exponential complexity, one may consider the *value query* model. The allocation algorithm is assumed to be able to access each entry of the utility table at unit cost. That is, for every set S, the algorithm may obtain $w(S)$ as an answer to a query, and this is considered to cost one computation step.

When there are n players and m items, each player gets on average m/n items. Hence one of the most basic pieces of information that we would like to deduce about a player is which k (e.g., $k = m/n$) items would bring her maximum utility. We call this a k-query. Unfortunately, even for the case of fractionally subadditive utility functions, a polynomial number of value queries do not suffice in order to answer a k-query, even in an approximate sense. Consider for example the fractionally subadditive utility function

$$w(S) = \max[a|S| + b, |S \cap T|]$$

where $a = 1/\sqrt{m}$, $b = m^{1/3}$ and T is some fixed set of size $\sqrt{m}$. Assume that w is given in form of a table, with T unknown (chosen at random). For $k = \sqrt{m}$, the set T has maximum utility $\sqrt{m}$. A random set of $\sqrt{m}$ elements would have utility only $m^{1/3} + 1$. Querying the value of set S, the reply is affected by the set T only if $|S \cap T| > a|S| + b = \frac{|T|}{m}|S| + m^{1/3}$. But the probability of choosing such a set S is smaller than the inverse of any polynomial in m.

Hence representing the utility function as a table and charging for value queries does not capture properly our intention that single player problems should be easy.

A more general class of queries that has been considered is that of *demand queries*. In this model one may set prices p_j for items, and obtain in one query the set S that maximizes $w(S) - \sum_{j \in S} p_j$. One advantage for this model is that prices come up naturally as dual variables to linear programs for the allocation problem, and demand queries offer a level of generality that allows one to solve linear program relaxations to the allocation problem.

In some cases (as in the examples given in Section 1.1) the utility function happens to have a compact (polynomial space) representation. One may be tempted to say that in these cases we are better off than in the cases in which one needs to resort to a query model. However, this is not always true. Consider for example a utility function that is defined as follows. There is a d-regular graph on m vertices. Every vertex corresponds to an item. The value of a set of items is the number of edges covered (incident with) by the corresponding vertices. This is a compact representation of a submodular utility function. However, it is NP-hard to answer demand queries on this representation. For example, there is a set of k items with value dk if and only if the graph has an independent set of size k.

The set cover utility function example given in Section 1.1 is a compact representation for which even value queries are NP-hard to answer (as they require solving a minimum set cover instance).

Luckily, in our work we do not need to deal with the subtleties involved with the representation of utility functions. The reason, which was explained in the introduction, is that we consider a situation in which a solution for a linear programming relaxation to the allocation problem is given to us. We do not care how this solution was found, or whether it is optimal or not. Thereafter we design oblivious rounding techniques that produce feasible allocations. These rounding techniques do not use any information about the utility functions of the players. Hence our results hold regardless of how utility functions are represented, and what type of queries (if any) are allowed.

A.3 Analysis of the two-step rounding technique

We present here a (known) proposition that can be used in order to analyse the two step rounding technique of Section 2.2.

Proposition 10 *Let x_i for $1 \leq i \leq t$ be independent indicator random variables, with $Pr[x_i = 1] = p_i$ and $\sum_i p_i = 1$. Then for every nonnegative integer k, $Pr[\sum_i x_i = k] \leq 1/k!$.*

Proof: Clearly the proposition is true when $t = 1$. Hence we may assume that $t \geq 2$. We now sketch a shifting argument that shows that we may assume that all p_i are equal. Assume that $p_1 \neq p_2$, and let $p = p_1 + p_2$. Replace x_1 and x_2 by two new indicator random variables, where $Pr[x_1 = 1] = y$ and $Pr[x_2 = 1] = p - y$, where $0 \leq y \leq p$. We need to choose y so as to maximize $Pr[\sum_i x_i = k]$. It is not hard to see that once $p_3, \ldots, p_t$ are fixed, this probability is some quadratic function $f(y)$, with $f(0) = f(p)$. Under these circumstances, the maximum is attained either when $y = p/2$ (and then $p_1 = p_2$) or when y is either 0 or p (and then one of the variables drops out). It follows that if not all p_i are equal, we did not maximize $Pr[\sum_i x_i = k]$.

Given that $p_i = 1/t$ for all i, we have that $Pr[\sum_i x_i = k] = \binom{t}{k}(1 - 1/t)^{t-k}/t^k \leq 1/k!$. $\square$

A.4 A negative example

We present an explicit example showing that the two step randomized rounding procedure described in Section 2.2 does not produce a constant approximation ratio for subadditive utility functions. Many other rounding techniques also fail to give a constant approximation ratio on this example.

There are m items. There are $n = \log^2 m$ players. With every player i we associate a random canonical subset S_i, where every item j belongs to set S_i independently at random with probability $1/\log m$. Items not in set S_i have no utility for player i. The utility of a subset $S' \subset S_i$ to player i is based on the set cover paradigm of Section 1.1: it is the smallest number of ground sets of type i that can cover S'. The definition of ground sets is a bit tricky. It involves two types of ground sets (easy and hard) and parameters $\ell = \sqrt{\log m / \log\log m}$ and $t = \ell/3$.

Definition 1 *A set U is an* easy *ground set of type i if there is some collection of $c \log m$ canonical sets S_j of other players such that U contains those items of S_i that appear in at most ℓ sets S_j in the collection. Here c is some explicit constant that will be defined later,* *satisfying* $0 < c \leq 1$. *A set* V *is a* hard *ground set of type* i *if it satisfies the following conditions:*

1. $V \subset S_i \setminus U$ *for some* easy *set* U *of type* i.
2. *For every collection of* t *additional easy sets* $U_{j_1} \dots U_{j_t}$ *of type* i,

$$|V \cap (S_i \setminus (U \cup U_{j_1} \cup \ldots \cup U_{j_t}))| \leq 2|S_i \setminus (U \cup U_{j_1} \cup \ldots \cup U_{j_t})|/\ell$$

This completes the description of the example.

The following technical lemma explains some of the properties of our example.

Lemma 11 *With high probability over the random choice of canonical subsets* S_i*, for every set* S_i *and every* $t+1$ *easy sets* $U_0, \dots, U_t$ *with respect to* i, $|S_i \setminus \cup_{k=0}^{t} U_k| \geq m^{1/4}$.

Proof: There are $\binom{n}{c \log m} < 2^{\log^2 m}$ ways of choosing the indices of the $c \log m$ sets in a collection that defines an easy set. Hence there are at most $2^{\log^3 m}$ ways of choosing t easy sets. For an item $j \in S_i$, we now compute the probability that it is in none of t easy sets. We use the fact that the sets S_j are chosen at random. The probability that j is not in a particular easy set is at least the probability that the first $\ell + 1$ sets from the respective collection all contain j, which is at least $(1/\log m)^{\ell+1}$. Applying the same principle to all t easy sets, the probability of j not being in any of the easy sets is at least $(1/\log m)^{t(\ell+1)}$ (and even higher, if the respective collections share canonical sets). This is at least $1/\sqrt{m}$, for our choice of parameters. Hence the expected number of items from S_i not covered by any of the easy sets is at least $\sqrt{m}/\log m$. Large deviations bounds now imply that with probability $1 - 2^{-m^\delta}$ (for some $\delta > 0$) there will be at least $m^{1/4}$ uncovered elements. As there are only n ways of choosing i and at most $2^{\log^3 m}$ ways of choosing the collections defining the easy sets, we can apply the union bound to prove the lemma. □

We now compute the utility of the random subset S_i to player i.

Proposition 12 *With overwhelming probability over the choice of the random canonical sets* S_i, $w_i(S_i) = \Omega(\ell)$.

Proof: We show that $t = \ell/3$ ground sets of type i do not suffice in order to cover S_i. Consider an arbitrary collection of t ground sets of type i. Some of them are *easy* sets U_j, and some of them are *hard* sets V_k. For each *hard* set V_k, add to the collection also the *easy* set U_k that corresponds to V_k by condition 1 of the definition of *hard* sets. (This will be needed when we later apply condition 2 of the definition of *hard* sets.) If there is more than one such easy set U_k that corresponds to V_k, pick one arbitrarily. The union of all *easy* sets in the resulting collection does not cover S_i, by Lemma 11. Every *hard* set in the collection can cover a fraction of at most $2/\ell$ of the remaining items, by condition 2 of the definition of *hard* sets. As the number of *hard* sets is at most $t < \ell/2$, some item of S_i must remain uncovered. □

We now present a feasible fractional solution to the LP. For every i we set $x_{i,S_i} = \epsilon/\log m$ for some sufficiently small $\epsilon > 0$, and all other variables to 0. The player constraints trivially hold. The item constraints also hold, because with overwhelming probability, no item belongs to more $(\log m)/\epsilon$ of the random canonical subsets S_i. The value of this fractional solution is at least $n\frac{\epsilon}{\log m}t = \Omega((\log m)^{3/2}/\sqrt{\log\log m})$, for our choice of parameters. We now contrast this value with the expected welfare of the feasible solution that is obtained after the two step randomized rounding procedure.

Proposition 13 *After applying the two step randomized rounding procedure, with overwhelming probability the feasible solution that is obtained has welfare at most $O(\log m)$.*

Proof: After the first step of the randomized rounding, with high probability between $\epsilon\log m/2$ and $2\epsilon\log m$ players remain. Take c in Definition 1 to be $c = \epsilon/2$. Consider an arbitrary remaining player i and its tentative set S_i, and apply the second step of randomized rounding to obtain a final set S' for player i. The final set S' (and likewise, S_i) will contain two types of items. The *easy* items are those that were contained in at most ℓ other tentative sets. The *hard* items are those that were contained in more than ℓ tentative sets. By Definition 1, one *easy* set U covers all easy items of S_i (and perhaps also some of the hard items, because the number of tentative sets might be larger than $c\log m$). With overwhelming probability, one *hard* set V covers those hard items that end up in S' (and were not in U). This follows from condition 2 in Definition 1 as follows. The set of hard items in $S' \setminus U$ is composed of items not in U, each chosen with probability at most $1/\ell$. Hence from each set of the form $S_i \setminus (U \cup U_{j_1} \cup \ldots \cup U_{j_t})$ they are expected to contain at most a $1/\ell$ fraction of the items. This number is not much smaller than $m^{1/4}$, by Lemma 11. Hence bounds on large deviations make it highly unlikely that the fraction would exceed $2/\ell$, even if one takes the union bound over all possible choices of $S_i \setminus (U \cup U_{j_1} \cup \ldots \cup U_{j_t})$.

Hence the integral solution will most likely have value at most $4\epsilon\log m$. $\square$

We have shown a gap of $\Omega(\sqrt{\log m/\log\log m})$ between the fractional solution to the LP and the solution obtained by the two step randomized rounding procedure. The alert reader may have noticed that for our particular example, the fractional solution that we presented for LP is far from optimal. Simply giving each player a single item has welfare $n = (\log m)^2$, which is better than the value of our solution to the LP. Moreover, the one step randomized rounding procedure of Section 2.1 would in fact recover such a solution from the LP.

To overcome this issue, we slightly modify our example. We create $(\log m)/\epsilon$ identical copies of the above example with the same set of players but disjoint sets of items. For a set S, let S^k be its items that are in copy k. We define the subadditive utility functions w' as $w'_i(S) = \max_k[w_i(S^k)]$, where w_i is defined as in the previous example. Now a fractional solution to the LP assigns value $\epsilon/\log m$ to each of the $(\log m)/\epsilon$ variables $x_{i,(S_i)^k}$. This fractional solution has value at least $t(\log m)^2$ and is optimal. For this modified example, both the one step and the two step randomized rounding procedures produce a feasible solution of value $O((\log m)^2)$, giving a gap of $\Omega(\sqrt{\log m/\log\log m})$.

A.5 An approximation ratio that depends on n

Theorem 14 *There is an oblivious rounding technique that for every player that has a fractionally subadditive utility function guarantees an expected utility that is at least a $(1 - (1-1/n)^n)$ fraction of the utility offered to the player by the LP solution.*

Proof: In the first step of the rounding technique each player i chooses a tentative set S with probability $x_{i,S}$. In the second step, for each item j independently, we employ the following procedure to resolve contention among players whose tentative sets contain j.

For item j, represent the solution to the LP by following flow problem. There is a directed graph with four levels.

The vertices. The first level contains a source node a. The second level contains n vertices, labelled 1 to n, one vertex for each player. The third level contains 2^n vertices, one vertex for every set of players. These vertices are labelled by sets T. The fourth level contains a sink node b.

The arcs (directed edges). From every vertex, there is an arc to every vertex in the next level. Hence there are all possible arcs of the form (a,i), (i,T) and (T,b).

Capacities. We now describe how to assign capacities to edges. For this we introduce a parameter $0 < \rho \leq 1$ which will eventually correspond to our approximation ratio. We let $p_i = \sum_{S|j\in S} X_{i,S}$ denote the probability that the tentative set of player i contains item j. Arcs of the form (a,i) have capacity ρp_i. Arcs of the form (i,T) have capacity 1 if $i \in T$ and capacity 0 if $i \notin T$. Hence arcs connecting a second level vertex i to a third level vertex T can simply be removed from the graph if player i is not in the set T. Arcs of the form (T,b) have capacity $c_T = (\prod_{i\in T} p_i)(\prod_{i\notin T}(1-p_i))$. Note that this capacity c_T is equal to the probability that T is the set of players that have j in their tentative set.

Proposition 15 *If for a given ρ there is a flow from a to b that saturates all arcs exiting a, then there is a randomized rounding scheme that ensures for every player i that conditioned on item j being in its tentative set, player i receives item j with probability at least ρ.*

Proof: Fix some flow f from a to b that saturates all arcs leading out of a. We use $f(i,T)$ to denote the flow on the arc (i,T). For every set T and player i, given that T is the set of players that have item j in their tentative set (which happens with probability c_T), we allocate item j to player i with probability $p(i,T) = f(i,T)/c_T$. This indeed is a probability distribution because the flow condition ensures that $\sum_i p(i,T) = \sum_i f(i,T)/c_T \leq c_T/c_T = 1$. Hence player i receives item j with probability $\sum_T c_T \frac{f(i,T)}{c_T} = \sum_T f(i,T) = \rho p_i$, where the last equality follows from the fact that the flow saturates the arc (a,i). Hence conditioned on j being in the tentative set of player i (which happens with probability p_i), player i receives item j with probability ρ. This is true regardless of which tentative set it is that player i selected (as long as it contains item j), because the analysis does not distinguish between different such sets. □

Proposition 16 *When $\rho = 1-(1-1/n)^n$, there is a saturated flow from a to b.*

Proof: We use the max-flow min-cut relation. We show that the graph has no a,b-cut of capacity less than $\rho\sum p_i$ for ρ as in the proposition.

Observe that a minimum cut cannot contain any arc of the form (i,T), because those arcs have capacity 1, whereas the cut that takes all arcs of the form (a,i) has capacity $\sum \rho p_i = \rho \sum p_i \leq \rho < 1$. Let W be the set of vertices of level two that are on the side of a in the minimum (a,b)-cut. Then for every third level vertex T we must cut the arc (T,b) if $W \cap T$ is nonempty. Hence the sum of capacities of all (T,b) edges that are not cut is exactly the probability that no player from W has j in its tentative set, namely $\prod_{i\in W}(1-p_i)$. As the sum of capacities entering b is exactly 1, the capacity entering b that is cut is at least $1 - \prod_{i \in W}(1 - p_i)$. Letting $p_W = \sum_{i \in W} p_i$ we have:

$$1 - \prod_{i \in W}(1 - p_i) \geq 1 - (1 - p_W/|W|)^{|W|} \geq p_W(1 - (1 - 1/|W|)^{|W|}) \geq p_W(1 - (1 - 1/n)^n) = \rho p_W$$

The second inequality can be verified by checking that the derivative of $(1 - (1 - p_W/|W|)^{|W|})/p_W$ with respect to p_W is positive for $0 < p_W < 1$. Hence a cut which is at least as small can instead cut all arcs (a, i) for $i \in W$, because the capacity of these arcs is ρp_W. This establishes that the minimum a, b-cut is simply to cut all arcs exiting a, which by the max-flow min-cut theorem implies that there is an a, b-flow that saturates all these arcs. $\square$

The two propositions above imply that there is a randomized rounding technique that for every player i and every item j guarantees that i receives j with probability at least $1 - (1 - 1/n)^n$ conditioned on item j being in player's i tentative set. Together with Proposition 3, this completes the proof of Theorem 14. $\square$

The complexity of the rounding technique presented in the proof of Theorem 14 is exponential in n. However, this is not a major drawback, because Theorem 14 is interesting mainly in the case that n is a small constant. The proof contains several inequalities. A necessary condition for the theorem to be tight is that the fractional solution to the LP is such that for some item j all p_i are $1/n$.

A.6 An approximation independent of n

We present an oblivious rounding technique that achieves an approximation ratio of at least $1 - 1/e$ when utility functions are fractionally subadditive, whose running time is polynomial regardless of the number of players. The rounding technique was designed in a way that makes its analysis simple. As intuition, consider the performance of the two step randomized rounding technique of Section 2.2 when there are N players, each choosing a tentative set that contains item j with probability $1/N$. Then the probability that no player gets item j is $(1 - 1/N)^N \leq 1/e$. By symmetry, it follows that player 1 gets item j with probability at least $(1 - 1/e)/N$. Observe that player 1 gets item j only if the tentative set that player 1 chooses happens to contain item j, and this happens with probability $1/N$. It follows that conditioned on player 1 choosing a tentative set that contains item j, player 1 in fact gets item j with probability at least $(1 - 1/e)$. The following three step rounding technique is designed in a way that simulates the above situation for all items (regardless of the question of whether the variables $x_{i,S}$ have value $1/N$ or not).

1. Each player chooses at most one set of items, where player i chooses her *tentative* set S with probability $x_{i,S}$.

2. For each item j, perform the following procedure. The players that happened to choose a tentative set containing item j are called the *competing* players for j. A competing player i is assigned at random an integer nonnegative weight $c_{i,j}$ with respect to item j as follows. For every competing player i, define $p_i = \sum_{S|j \in S} x_{i,S}$. Let N be some large integer such that $1/N$ divides all p_i. (We shall think of N as being ∞.) Then player i gets to control $p_i N$ virtual players, each competing for item j with probability $1/N$. The probability that $c_{i,j}$ is set to t, where $t \geq 1$, is then tentatively set to be the

same as the probability that t of the p_iN virtual players compete for item j, namely $\binom{p_iN}{t}(1-\frac{1}{N})^{p_iN-t}\frac{1}{N^t}$, which in the limit (when N tends to ∞) can be taken to be $p_i^t e^{-p_i}/t!$. This tentative probability is now scaled by $1/p_i$, to cancel out the fact that player i has probability p_i (rather than probability 1) of competing. In summary, for $t \geq 1$, $c_{i,j} = t$ with probability $\frac{1}{p_i}p_i^t e^{-p_i}/t!$. The probability that $c_{i,j} = 0$ is then taken to be $1 - \frac{1}{p_i}\sum_{t\geq 1} p_i^t e^{-p_i}/t! = 1 - \frac{(e^{p_i}-1)e^{-p_i}}{p_i} = 1 - \frac{1-e^{-p_i}}{p_i} > 0$ (as can be seen from the Taylor expansion $e^x = \sum_{t\geq 0} x^t/t!$).

3. If $\sum_i c_{i,j} > 0$, allocate item j to player i with probability $c_{i,j}/\sum_i c_{i,j}$. If $\sum_i c_{i,j} = 0$, do not allocate item j to any player. (Of course, one may allocate item j to some player also when $\sum_i c_{i,j} = 0$, but this is not used in the analysis.)

Observe that in the three step randomized rounding procedure, $\sum c_{i,j}$ is distributed *exactly* as a random variable that is the sum of $N\sum p_i$ indicator random variables, each with probability $1/N$ of being 1. When $\sum p_i = 1$, the probability that item j is not allocated at all is essentially $1/e$, which can be shown to imply that in the worst case the approximation ratio of the three step rounding technique is no better that $1-1/e$. The main theorem of this section shows that the approximation ratio is in fact always at least $1-1/e$.

Theorem 17 *For fractionally subadditive utility functions, the three step randomized rounding procedure obtains a random feasible solution with expected welfare at least a $1-1/e$ fraction of the value of the LP solution.*

Proof: Consider an arbitrary player, which for simplicity we will rename to be player 1. In expectation, in step 1 its utility does not change. By Proposition 3, it suffices to show that for every item j in the tentative set of player 1 (which we call S_1), the probability that player 1 is allocated item j in the final solution is at least $1-1/e$.

Consider an arbitrary item j, and recall the definition of $p_i = \sum_{S|j\in S} x_{i,S}$. By the item constraints, $\sum p_i \leq 1$. We may assume that conditioned on the value of p_1, $\sum_{i\neq 1} p_i$ is maximized, as this puts player 1 in maximum competition for item j. Hence we may take $\sum p_i = 1$. It follows that the probability that item j is allocated at all (without assuming yet that player 1 in fact chose set S_1) is at least $1-1/e$. Moreover, conditioned on item j being allocated, a player i has probability exactly p_i of being the player that gets item j. This is because player i controls precisely p_iN of the N potential virtual players for item j, and each one of the virtual players has equal probability of getting item j. Hence the probability that player 1 gets item j is $p_1(1-1/e)$. But the probability that player 1 gets item j is by definition the product of the probability that j is in players 1's tentative set (which is p_1), times the probability that player 1 gets item j conditioned on j being in its tentative set. Hence this last probability is $1-1/e$ as desired. Finally, note that this last probability is independent of which is the actual tentative set containing item j chosen by player 1, so it remains $1-1/e$ also when the tentative set is S_1. □

A.7 An incentive compatible mechanism

In this section we explain how the oblivious two step rounding technique of Section 2.2 can be leveraged to give incentives to players to report their true utilities. As explained earlier, the true utility functions of the players play no role when rounding the LP (the rounding technique is oblivious). However, they are necessary for obtaining an optimal fractional solution to the LP. A player who behaves selfishly might report an incorrect utility function so as to influence the fractional solution of the LP in a way that benefits her. We present a mechanism that gives players incentives to report their true utility functions. This mechanism requires players to pay for the items that they receive. The notion of incentive compatibility that we use is similar to the one used by Lavi and Swamy [8], and the approach used in this section can be viewed as an extension of the approach used in [8].

We slightly modify the two step rounding technique. Recall that for k sufficiently large ($k = O(\log m / \log \log m)$ suffices) there is only very small probability ϵ (and we can take $\epsilon < 1/m^2$) that the following "high contention" event happens: there is an item that after step 1 remains in more than k tentative sets. The modified two step rounding technique is as follows. Step 1 (the tentative allocation) is considered successful only if the high contention event did not happen. If step 1 is not successful (which is a low probability event), then the rounding procedure simply aborts and does not allocate any item. If step 1 is successful, then a modified step 2 is performed. For every item j independently, every player that has j in her tentative set gets item j with probability exactly $1/k$, regardless of the number of players that have item j in their tentative set (even if no other player wants the item). Hence some of the items are not allocated at all.

For the rest of our discussion we make the simplifying assumption that the probability of aborting is so small that it can be considered to be exactly 0. (Later we shall return to discuss this issue.) Then the modified two step rounding technique has the following property: every player receives every item of her tentative set with probability *exactly* $1/k$, independently of all other events. We now define modified utility functions for the players. For each player i and set S, we define $w_i^*(S)$ as the expected value of $w_i(S')$, where S' is chosen by independently including in it every item of S with probability $1/k$.

Now the key is to solve the welfare maximizing LP with respect to w^* rather than with respect to w. Thereafter, the modified two step rounding procedure produces an integer solution with the following guarantee for every player: the expected utility she gets (with respect to her true utility function w_i) is *exactly* the same as the utility offered to her by the fractional solution (with respect to w_i^*).

Let us make now the assumption that players are *expectation maximizers*. That is, every player wishes to maximize the expected utility that she gets (and does not care about the variance, among other things). In this case, the rounding of the LP can be viewed as exactly preserving the utility of the players. That is, a fractional solution with respect to w^* exactly characterizes the utilities that players eventually receive. Hence to maximize welfare, one should simply find the optimal fractional solution with respect to w^*, and this can be done in polynomial time (when utility functions are "digestible", see below).

Now we can use standard methodology, namely, the VCG mechanism [11, 3, 7], to construct an incentive compatible mechanism. In short, every player is required to pay her marginal influence on the set of other players, namely, the difference between the optimal total welfare had the player not existed, and the total welfare allocated to the other players when the player does exist. Let us define the *benefit* that a player receives as the utility she receive minus her payment. Under the VCG mechanism, reporting true utilities is a strategy that maximizes the benefit that a player receives. In our context, we may call a player *expected benefit maximizing*, if her objective is to maximize her expected benefit.

Observe that given a value oracle for w_i, the value of $w_i^*(S)$ can be estimated with arbitrary precision with high probability by a natural randomized algorithm (that selects random sets S', evaluates $w_i(S')$, and takes the average). This does not necessarily mean that if one can answer demand queries regarding w_i then one can also answer demand queries with respect to w_i^*. But let us call a utility function w_i *digestible* if we are given a short (polynomial size) representation of w_i that allows one to efficiently answer demand queries with respect to w_i^*.

Observe that by the results of Section 2.2 and Proposition 2 we have that for every S, $w_i^*(S) \geq w_i(S)/k$, when w_i is subadditive. Moreover, it may be interesting to note that if w_i is subadditive (or submodular, or fractionally subadditive, respectively) then so is w_i^*, though we do not use this fact.

Summarizing the above discussion, we have the following theorem.

Theorem 18 *There is a mechanism that runs in polynomial time if utility functions are digestible, and is incentive compatible with respect to players that maximize expected benefit, that recovers in expectation at least a fraction of* $\Omega(\log\log m/\log m)$ *of the optimum welfare when utility functions are subadditive.*

We remark that there are slight extensions to the above theorem. The payment mechanism can be modified to ensure that a player never pays more than the actual utility that she receives (rather than the expected utility). The requirement that utility functions are digestible can be relaxed to having a demand oracle for w^*. Every player i can serve as the oracle with respect to her own utility function, and answer demand queries with respect to w_i^*. This interactive mechanism comes at a price of a weaker notion of incentive compatibility, that of ex-post Nash equilibrium (discussed in [8]).

Let us now return to our assumption that the probability of aborting is 0 rather than at most ϵ (for some arbitrarily small ϵ). If players are not sensitive to multiplicative changes of $(1-\epsilon)$ in their benefit (which is a very reasonable assumption), then Theorem 18 continues to hold unchanged. But if players are sensitive even to arbitrarily small values of ϵ, then strictly speaking, the mechanism referred to by Theorem 18 is not necessarily truthful. The reason is that the probability of aborting is not exactly ϵ, but rather at most ϵ, and might vary (very slightly) depending on the actual fractional solution to the LP. It is in a player's interest to reduce this probability as much as possible. It is conceivable that in some situations a player can slightly reduce the probability of abort by reporting a utility function that is not truthful.

Theorem 18 was presented mainly to illustrate the relevance of oblivious rounding techniques to the design of incentive compatible mechanisms. The author is aware of multiple shortcomings of this theorem, including the fact that subadditive utility functions would often not be digestible, that incentive compatibility comes at a price of a worse approximation ratio, that the notion of incentive compatibility in expectation (especially in combination with ex-post Nash equilibrium) might be too weak, and more generally, that the question of which mechanisms induce truthful behavior has psychological aspects that are not captured by the mathematical models.

A.8 Submodular utility functions

A question left open by this paper is whether there is some rounding technique that achieves an approximation ratio better than $1 - 1/e$ when utility functions are submodular. In this context, it may be informative to note that the rounding techniques that were developed in this paper are *noncombinational*, in the sense that for every player i, the set of items that she receives is always entirely contained in some set S for which $x_{i,S} > 0$ in the given fractional solution of the LP. This property may be desirable in some contexts. For example, the solution to the LP might reflect also capacity constraints (see [6] for a setting where this is the case), and then noncombinational rounding is guaranteed to respect them. However, it is easy to construct examples in which noncombinational rounding cannot achieve an approximation ratio better than $1 - 1/e$ even for the special case of linear utility functions, whereas we do know that better approximation ratios are achievable in this special case (see Section 2.1).