On Maximizing Welfare when Utility Functions are Subadditive

Uriel Feige *

October 8, 2007

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 2006], 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.

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,

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

*Weizmann Institute, Rehovot, Israel. uriel.feige@weizmann.ac.il

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 [3] 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 [3] for more details.

In [3] 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. 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 3.2.1).

Another result of this paper is that if the utility functions are further restricted to be *fractionally subadditive* (see definition in 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 [4] for a class of utility functions known as XOS, and as we show (see Proposition 1.1), 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 [4], 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 1.2), 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.

The other known rounding technique achieving an approximation ratio of $1 - 1/e$ for fractionally subadditive utility functions [4] is not oblivious. Implementing the rounding technique of [4] 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. In a preliminary version of this paper [5], the oblivious nature of our rounding techniques was used (in combination with other ideas) in order to design a truthful mechanism that recovers at least an $\Omega(\log\log m/\log m)$ fraction of the total optimum welfare when utility functions are subadditive. However, as this result involves multiple assumptions about how players would react to various incentives, it is omitted from the current version of this manuscript.

1.1 Classes of utility functions

In this section we discuss some classes of utility functions that we shall refer to throughout this work. For more details, see [8].

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.

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

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 $T \subset S$, $w(j \cup S) - w(S) \leq w(j \cup T) - w(T)$.

Fractionally subadditive. $w(S) \leq \sum \alpha_i w(T_i)$ with $0 \leq \alpha_i \leq 1$ for all i, whenever the following condition holds: for every item $j \in S$, $\sum_{i | j \in T_i} \alpha_i \geq 1$. (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 [8]. This fact will be proved in Proposition 1.1.

Subadditive (a.k.a. complement free). $w(S \cup T) \leq w(S) + w(T)$, 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 $w(S) \geq 1$. For subadditive utility functions, we have in addition $w(S) \leq 2$, because $w(S) \leq w(\{1, 2\}) + w(3)$. For fractionally subadditive utility functions we have $w(S) \leq 3/2$, because $w(S) \leq (w(\{1, 2\}) + w(\{2, 3\}) + w(\{1, 3\}))/2$. For submodular utility functions we have $w(S) = 1$, because $w(S) + w(1) \leq w(\{1, 2\}) + w(\{1, 3\})$. The function cannot be a linear utility function at all, because $w(\{1, 2\}) \neq w(1) + w(2)$.

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 $T_1, \ldots, T_k$ 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.

Another class of utility functions that is considered in [8, 3, 4] 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 1.1 shows that the class XOS is the same as the class of fractionally subadditive utility functions.

Proposition 1.1 *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). $\square$

The reader may note that Proposition 1.1 is a straight-forward variation of the Bondareva-Shapley Theorem [2, 9]. 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.

1.2 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 [8, 3, 4].)

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 $\sqrt{m}$. Observe that for these parameters, $a = \frac{|T|}{m}$. Assume that w is given in form of a table, with T unknown (chosen at random). For $k = \sqrt{m}$, the set with maximum utility is T, and its utility is $\sqrt{m}$. A random set of $\sqrt{m}$ elements would have utility only $m^{1/3} + 1$, because the expected size of $|S \cap T|$ is $\frac{|T|}{m}|S| = 1$, and the term $a|S| + b$ dominates. 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, because regardless of the size of S, the term $m^{1/3}$ is almost surely larger than $|S \cap T| - \frac{|T|}{m}|S|$. It follows that polynomially many queries do not suffice in order to learn anything about the best choice T.

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 the utility function happens to have a compact (polynomial space) representation. One may be tempted to assume 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).

As explained in the introduction, in our work we do not need to deal with the subtleties involved with the representation of utility functions.

1.3 Integrality gaps

It is shown in [3] 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 1.2 *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, n players and utility functions as in Proposition 1.2, 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.

As shown in [3] (by a reduction from the max k-coverage problem), for every $\epsilon > 0$, it is NP-hard to approximate the maximum welfare for XOS utility functions (which as we noted, are the same as fractionally subadditive utility functions) within a ratio better than $1 - 1/e + \epsilon$. Likewise, the integrality gap of the welfare maximizing LP in this case is $1 - 1/e + \epsilon$.

1.4 Notation and conventions

Throughout we use the following notation. We assume that we are given an arbitrary feasible solution (though not necessarily optimal) to the welfare maximizing LP. For every player i and set S, we use $x_{i,S}$ to denote the (fractional) value assigned to variable $x_{i,S}$ in this particular solution. (This is a slight abuse of notation, as previously $x_{i,S}$ denoted a name of a variable rather than a value given to it.) The value of the objective function under this particular solution, $\sum_{i,S} w_i(S)x_{i,S}$, will be denoted by $w(LP)$. For every i, the contribution of player i to the objective function, namely $\sum_S w_i(S)x_{i,S}$, will be denoted by $w_i(LP)$.

It will be convenient to also establish special notation for $\sum_{S|j\in S} x_{i,S}$, which can be interpreted as the fraction of item j assigned to player i by the solution of the LP. This quantity will be denoted by $f_{i,j}$.

To simplify the presentation, we shall assume that the solution given to the LP is such that all constraints are satisfied with equality. This convention can be made without loss of generality. For example, if the item constraint associated with item j is not satisfied with equality, we may add a special player P_j with a utility function that is identically 0, and set the value of variable $x_{P_j,j}$ to $1 - \sum_{1\leq i\leq n} f_{i,j}$. Likewise, a player constraint for player i can be satisfied with equality by setting the value of the variable $x_{i,\phi}$ (where ϕ is the empty set) to a value of $1 - \sum_{S\neq\phi} x_{i,S}$.

2 Basic oblivious rounding techniques

This section contains known results [3], 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. The item constraints of the welfare maximizing LP (recall the convention from Section 1.4) imply that for every item j, $\sum_i f_{i,j} = 1$, and hence the $f_{i,j}$ define a probability distribution over players. Allocate item j to one player, by selecting player i with probability $f_{i,j}$. This gives a feasible allocation. When utility functions are additive, then in expectation a player's utility will be the same as $w_i(LP)$, and the expected welfare generated by this rounding technique is equal to $w(LP)$. However, when utility functions are not additive, this is far from true.

Consider the following example in which $n = \sqrt{m}$. All items are partitioned into n equal size sets $T_1, \ldots, T_n$. All players have the same utility function $w(S) = \max_{j=1}^{n} |S \cap T_j|$, which is fractionally subadditive. A feasible fractional solution to the LP has $x_{i,j} = 1/n$ for every player i and set T_j, and $x_{i,j} = 0$ otherwise. For this fractional solution, $w_i(LP) = \frac{1}{n}\sum_{j=1}^{n} |T_j| = m/n = n$, and $w(LP) = n^2 = m$. However, the randomized rounding procedure described above is unlikely to ever allocate more than $O(\log n)$ items from the same set T_j to a 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 $w(LP)$.

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 [3] (which was not oblivious).

1. **Tentative allocation.** For every player i, recall that the player constraints (and our convention from Section 1.4) imply that $\sum_S x_{i,S} = 1$, and hence the $x_{i,S}$ values may be thought of as defining a probability distribution over sets. Each player chooses a *tentative* set of items, where player i chooses set S with probability $x_{i,S}$. The expected utility to player i of her tentative set is exactly $w_i(LP)$. However, her chosen set might intersect with sets chosen by other players. Hence, the solution might not be feasible.

2. **Uniform contention resolution.** 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.

To analyse the quality of the final solution, we use the following known proposition.

Proposition 2.1 *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 $Pr[\sum_i x_i = k] = \binom{t}{k}(1 - 1/t)^{t-k}/t^k \leq 1/k!$. □

Proposition 2.1 implies 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 [3] only a weaker bound of $O(\log m)$ 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 2.2, which together with monotonicity implies the desired bound. (Monotonicity is appealed to because the probability per item in Proposition 2.2 is exactly $1/k$, whereas we are interested in the case where the probability is at least $1/k$.)

Proposition 2.2 *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 $w_i(LP)$. 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 2.2. By linearity of expectation, it follows that the expected welfare of the allocation delivered by the two step rounding technique is at least $\frac{1-1/m}{k} w(LP) = \Omega(\frac{\log\log m}{\log m}) w(LP)$.

2.3 Fractionally subadditive utility functions

For fractionally subadditive utility functions we can use a strengthening of Proposition 2.2. The difference between the two propositions is the removal of the requirement for statistical independence among items.

Proposition 2.3 *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 2.3 follows from the equivalence between fractionally subadditive and XOS utility function (Proposition 1.1), 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.

Recall that the expected value of the tentative set S_i to player i is $w_i(LP)$. 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 2.3 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. The second step resolves 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. But potentially, we could do better if 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. This principle was indeed used in [7]. There the setting is that utility functions are additive, and it is straightforward to determine which player derives the highest marginal utility from an item. The same principle was used in [4] for XOS utility functions (maximum of additive utility functions), under the assumption that one can determine for every player which additive utility function maximizes the utility of its tentative set.

In contrast to [4] (and to [7]), 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 1.2) and the integrality gap of the welfare maximizing LP in this case.

We alert the reader to an implicit distinction between worst case instances and "typical" instances in the above discussion. If we are only interested in the worst case approximation ratios, then our oblivious rounding techniques are indeed optimal (up to low order terms). However, it is not difficult to design specific instances on which rounding techniques that do take into account the actual utility functions of players do better than our oblivious rounding techniques.

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 [4].

We first describe an instance in which the two step randomized rounding procedure of Section 2.2 does not produce an approximation ratio significantly better than $1/2$. There are two players and $2m$ items partitioned into two equal size sets S and T. Player P_1 has utility 1 if she gets at least one item from S (items in T have negligible utility for P_1). Player P_2 has a utility 1 if she gets at least one item from T (the items from S have negligible utility for P_2).

For $1 \leq j \leq m$, let S_j denote the set containing the jth item from S and all but the jth item from T. Let T_j denote the set containing the jth item from T and all but the jth item from S. An optimal solution to the LP sets $x_{1,S_j} = 1/m$ for all sets S_j, and $x_{2,T_j} = 1/m$ for all sets T_j. All other variables are 0. The value of the LP is $2m/m = 2$. However, the two step rounding procedure will produce a solution of expected value $1 + 1/m$.

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. (The reader is advised to review notation from Section 1.4. Here we shall not use the convention that item constraints are satisfied with equality, as this involves an increase in the number of players, and our rounding technique is specific for two players.)

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 2.

   (c) If $j \in S_1 \cap S_2$, then allocate j to player 1 with probability $\frac{f_{2,j}}{f_{1,j}+f_{2,j}}$ and to player 2 with probability $\frac{f_{1,j}}{f_{1,j}+f_{2,j}}$.

   (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 3.1 *For every player $i \in \{1, 2\}$, 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 $\frac{3}{4} w_i(LP)$.*

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 $w_1(LP)$. 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

$$1 - f_{2,j}\frac{f_{1,j}}{f_{1,j} + f_{2,j}} \geq 3/4$$

where the inequality follows from the fact that $f_{1,j} + f_{2,j} \leq 1$ (the item constraints). Now the proof follows from Proposition 2.3. $\square$

We now consider the case when the number of players is $n > 2$.

Theorem 3.2 *There is an oblivious rounding technique that for every player i that has a fractionally subadditive utility function guarantees an expected utility that is at least a $(1 - (1 - 1/n)^n)w_i(LP)$.*

An earlier version of this manuscript contained a proof of Theorem 3.2 that involved setting up a flow problem of size exponential in n, and hence resulted in an efficient algorithm only when the number of players is small. (When n is large, $(1 - (1 - 1/n)^n)$ approaches $1 - 1/e$ from above, and this case will be handled below.) The proof of Theorem 3.2 is omitted from the current version of the paper, because subsequently a better proof (not requiring time exponential in n) was discovered by Jan Vondrak and is given in [6].

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 having probability $1/N$ of choosing a tentative set that contains item j. Then the probability that no player gets item j is $(1 - 1/N)^N \leq 1/e$. By symmetry among players, it follows that player i gets item j with probability at least $(1 - 1/e)/N$. Observe that player i gets item j only if the tentative set that player i chooses happens to contain item j, and this happens with probability $1/N$. It follows that conditioned on player i choosing a tentative set that contains item j, player i in fact gets item j with probability at least $(1 - 1/e)$.

Let us now consider a more general case. Consider an item j and recall our notation of $f_{i,j} = \sum_{S|j \in S} x_{i,S}$ for the fraction of item j assigned by the LP solution to player i. Fixing j in this discussion, we omit the subscript j and use f_i to denote $f_{i,j}$. Recall also our convention from Section 1.4 that $\sum_i f_i = 1$. The previous paragraph corresponds to the case when there are N players and $f_i = 1/N$ for all i. Now we shall not make these assumptions. Instead, let N be an integer such that $1/N$ divides all f_i. Every player i gets to control $f_i N$ *coins* (associated with item j). Each of the coins in biased, and comes up 1 independently with probability $1/N$. Every player i tosses her $f_i N$ coins independently, and those coins that come up 1 are placed into a bag B_j. If bag B_j is empty at the end of the process, no player gets item j. If the bag contains at least one coin, than one coin is chosen independently at random from this bag, and the player who owned this winner coin gets item j.

Observe that the probability that bag B_j is empty is precisely $(1 - 1/N)^N < 1/e$. Hence with probability at least $1 - 1/e$, one player gets item j. By symmetry among the coins, every coin is equally likely to be the winning coin, and hence player i has probability at least $(1-1/e)(f_i N)/N = (1 - 1/e)f_i$ of getting item j.

The above discussion was lacking in the sense that we ignored the question of whether j was in the tentative set of player i. A preliminary fix to that is to do the following: if player i happened to place a coin in bag B_j (this happens with probability $1 - (1 - 1/N)^{f_i N} \leq f_i$), then she chooses a tentative set that contains item j. To have the right marginal probability of choosing a tentative set that contains item j, then conditioned on placing no coin in bag B_j, player i is required to choose such a set with probability $(f_i - 1 + (1 - 1/N)^{f_i N})/(1 - 1/N)^{f_i N}$. As before, player i gets item j with probability at least $(1-1/e)f_i$, but now, player i gets item j only if item j is in her tentative set (which happens with probability f_i). Hence conditioned on having item j in her tentative set, player i gets item j with probability at least $1 - 1/e$, as desired. Note that this last probability is independent of which is the actual tentative set containing item j chosen by player i, so it remains $1 - 1/e$ for each such set.

We are not quite done. The problem is that there are m different items to consider. We cannot have the tentative set chosen by player i depend on what happened in the m coin tossing experiments that player i performs, because this might place too many constraints on the choice of her tentative set. To get around this last problem, we switch the order of two probabilistic events, while exactly preserving their joint distribution. Rather than first randomly choosing how many coins player i places in bag B_j and then randomly choosing a tentative set for player i, we first randomly choose the tentative set, and then based on the outcome randomly choose the number of coins, with the appropriate marginal distribution conditioned on whether the tentative set contains item j.

Let us now explain how this marginal probability is computed. If the tentative set does not contain item j, then the probability of player i placing a coin in bag B_j becomes 0. If the tentative set does contain item j, then the distribution over the number t of coins that player i places in bag B_j is computed as follows. The probability that t of the f_iN coins come up 1 is $\binom{f_iN}{t}(1-\frac{1}{N})^{f_iN-t}\frac{1}{N^t}$, which in the limit (when N tends to ∞) can be taken to be $f_i^te^{-f_i}/t!$. This probability is now scaled by $1/f_i$, to cancel out the fact that player i has probability f_i (rather than probability 1) of having item j in her tentative set. Hence for $t \geq 1$, the probability for t coins is $\frac{1}{f_i}f_i^te^{-f_i}/t!$. The probability that player i places no coins in bag B_j (even though her tentative set does contain item j) then becomes $1-\frac{1}{f_i}\sum_{t\geq 1}f_i^te^{-f_i}/t! = 1-\frac{(e^{f_i}-1)e^{-f_i}}{f_i} = 1-\frac{1-e^{-f_i}}{f_i}$. (This probability is indeed nonnegative because $e^{-x} \geq 1-x$ for every x.)

Observe that having chosen the tentative set first, player i can enforce for every item separately the correct marginal distribution for its bag conditioned on the choice of tentative set. In summary, we have the following three step rounding technique:

1. **Tentative allocation.** Each player chooses a tentative set of items, where player i chooses set S with probability $x_{i,S}$.

2. **Assigning weights to competing players.** A player i is said to *compete* for item j if item j is in her tentative set. Consider an arbitrary item j, and to simplify notation use f_i as shorthand notation for $f_{i,j}$. A competing player i is assigned at random an integer nonnegative weight $c_{i,j}$ with respect to item j as follows. For $t \geq 1$, $c_{i,j} = t$ with probability $\frac{1}{f_i}f_i^te^{-f_i}/t!$. The probability that $c_{i,j} = 0$ is $1-\frac{1-e^{-f_i}}{f_i}$.

3. **Weighted contention resolution.** 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.)

The proof of the following theorem is implicit in the intuitive introduction that we gave to the three step rounding technique, but we repeat it for completeness (sometimes using terminology from the intuitive introduction).

Theorem 3.3 *For fractionally subadditive utility functions, the three step randomized rounding procedure obtains a random feasible solution with expected welfare at least a* $(1 - 1/e)w(LP)$.

Proof: Consider an arbitrary player, which for simplicity we will rename to be player 1. The expected value of her tentative set is $w_1(LP)$. By Proposition 2.3, 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 (regardless of whether $j \in S_1$), and recall that $\sum_i f_i = 1$ (where f_i is shorthand notation for $f_{i,j}$, and we use here the convention from Section 1.4). Observe that $\sum_i c_{i,j}$ (which we denote here by c_j) is distributed *exactly* as a random variable that is (the limit as N tends to infinity of) the sum of $N \sum f_i = N$ indicator random variables, each with probability $1/N$ of being 1. Item j is allocated if and only if $c_j > 0$, which has probability $1 - (1 - 1/N)^N > 1 - 1/e$. Moreover, conditioned on item j being allocated, every coin has exactly the same probability of winning item j (this is a consequence of the marginal distribution that is enforced on $c_{i,j}$ conditioned on the tentative set containing item j), and hence this probability is at least $(1 - 1/e)/N$. As player 1 controls $f_1 N$ coins, her probability of winning item j is at least $(1 - 1/e)f_1 N/N = (1 - 1/e)f_1$. However, player 1 may win j only if she competes for j, and this happens with probability exactly f_1. Hence conditioned on having item j in her tentative set, player 1 gets item j with probability at least $1 - 1/e$, 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. $\square$

We remark that the rounding technique as described assumes computation with infinite precision. A slight loss in the approximation ratio might result from rounding errors when finite precision arithmetic is used. However, this loss will not bring the approximation ratio below $1 - 1/e$, because it can be compensated for by slackness in the analysis. (For example, as the f_i are solutions of an LP, they are rational, and hence N is finite, and $(1 - 1/N)^N$ is strictly smaller than $1/e$.) We omit the tedious details from this manuscript.

3.2 Subadditive utility functions

3.2.1 A negative example

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 many other rounding techniques).

There are m items. There are $n = \log^2 m$ players. With every player i we associate one *canonical* subset S_i. We shall later explain how this set is chosen, but here only remark that its size is roughly $m/\log m$. 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 3.4 *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 least one but 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,*

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

To complete the description of the example, it remains to explain how the canonical sets S_i are chosen. We require the choice of canonical sets to satisfy **three properties**:

1. For a sufficiently small $\epsilon > 0$ (independent of m), no item belongs to more than $(\log m)/\epsilon$ canonical sets.
2. For every canonical set S_i and every $t + 1$ easy sets $U_0, \ldots, U_t$ with respect to i, $|S_i \setminus \cup_{k=0}^{t} U_k| \geq m^{1/4}$.
3. For every canonical set S_i, $w_i(S_i) = \Omega(\ell)$.

A simple consequence of property 1 is that for every type i, every item is in some easy set of type i. In particular, this implies that every canonical set can be completely covered by easy sets, and hence every utility function is well defined.

We do not show an explicit choice of canonical sets that satisfies the three properties. Instead, we choose these sets independently at random. Specifically, for every item j and player i, item j belongs to set S_i independently at random with probability $1/\log m$. We shall show that positive probability (in fact, probability very close to 1) this creates an instance on which the three properties hold.

The following proposition establishes property 1.

Proposition 3.5 *Let $\epsilon > 0$ be a sufficiently small constant (independent of m). Than with high probability (say, at least 0.99) over the choice of canonical sets, no item belongs to more than $(\log m)/\epsilon$ canonical sets.*

Proof: In expectation, an item belongs to $n/\log m = \log m$ canonical sets. Standard bounds on large deviations for sums of independent random variables show that for sufficiently small ϵ, an item has probability at most $1/100m$ of belonging to more than $(\log m)/\epsilon$ canonical sets. Taking the union bound over all items, with probability at least 0.99 no item is in more than $(\log m)/\epsilon$ canonical sets. $\square$

The following lemma establishes property 2.

Lemma 3.6 *With high probability (say, 0.99) over the random choice of canonical sets S_i, for every canonical set S_i and every $t+1$ easy sets $U_0, \ldots, U_t$ 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 canonical sets 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 $\sqrt{m}/\log m$. Moreover, the events of not being covered by the easy sets are independent among items, because for each item they depend only on whether the item is contained in the canonical sets that define the easy sets, and each item is placed in canonical sets independently of other items. Large deviation 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. □

The following proposition establishes property 3.

Proposition 3.7 *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 is associated with 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 is associated with V_k, pick one arbitrarily. The union of all *easy* sets in the resulting collection does not cover S_i, by property 2 (which was already established in Lemma 3.6). 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 assume that we have an instance in which the above three properties hold. The following is a feasible fractional solution to the LP. For every i we set $x_{i,S_i} = \epsilon/\log m$ for some sufficiently small $\epsilon > 0$ (as in property 1), and all other variables to 0. The player constraints trivially hold. The item constraints also hold, by property 1. 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 3.8 *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 3.4

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 3.4, 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 3.4 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 property 2. 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 $4\epsilon \log m$. □

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 $\Omega(\sqrt{\log m / \log \log m})$.

3.2.2 A constant approximation ratio

We now present a rounding procedure for the LP that has an approximation ratio of $1/2$ when utility functions are subadditive. As our formal description of the rounding technique is based on a lot of hindsight, let us first give some intuition regarding what the rounding technique is aiming to achieve.

The first step of the rounding technique is one that we have already seen – each player chooses a tentative set. Consider now an arbitrary player, say player 1, and her tentative set, say S_1. Player 1 would need to give up some of the items from S_1, since S_1 is likely to intersect other sets. Our goal would be to devise a method by which after giving up some items, player 1 still retains (in expectation) at least half the utility of $w_1(S_1)$. But as we have seen in Section 3.2.1, even if every item of S_1 has probability at least $1/2$ of staying with player 1, still correlations among the items may cause the expected remaining utility to be much lower than $w_1(S_1)/2$.

Our way of deciding which items player 1 should give up (and which items she should retain) is based on the idea of pairing different rounding scenarios. By scenario we mean here the collection of tentative sets chosen by the players. Two scenarios can be paired with respect to player 1 if in both scenarios player 1 chooses the same tentative set. Other players may choose different tentative sets in these two scenarios. Assume that we paired two scenarios in which the tentative set for player 1 is S_1, and moreover, assume for simplicity that the following property holds: every item of S_1 is in tentative sets of other players only in one of the two scenarios. In this case, in each of the two scenarios it suffices to give player 1 those items of S_1 that are under no contention. Hence every item of S_1 is given to player 1 in one of the two scenarios, and by subadditivity of her utility function, the sum of utilities that player 1 obtains in the two scenarios is at least $w_1(S_1)$. This implies that on average over these two scenarios, player 1 gets utility at least $w_1(S_1)/2$. (Technically, this last argument requires both scenarios to have the same probability of being generated by the randomized rounding procedure. This point can be overcome by associating a probability with each scenario and a weight with each pairing, and reducing the probability of each of the paired scenarios by the weight of the pairing.) If for every player, all rounding scenarios could be paired in such a way (eventually reducing all probabilities to 0), this would imply that there is a feasible solution of value at least $w(LP)/2$.

In the above approach the pairing function has to be player dependent. Consider for example a scenario A in which item j belongs to the tentative sets of three different players, say 1,2 and 3. Let B_i be the scenario that player i pairs with scenario A. Then it cannot be that $B_1 = B_2 = B_3$ because then one of the three players will not get item j neither in scenario A nor in the paired scenario $B = B_1 = B_2 = B_3$.

Let us examine more closely the issues involved in finding a good pairing function. Observe that in the example given in Section 3.2.1, if one chooses two random scenarios in which player 1 chooses set S_1, it is likely that a constant fraction of the items of S_1 will appear in other tentative sets in both scenarios. Moreover, some items of S_1 will appear in many (a growing function of n) tentative sets in both scenarios. Nevertheless, in our rounding technique we will pair random scenarios (sharing S_1). We will make sure that every item of S_1 will be given to player 1 in at least one of the two scenarios. For some items (say item j), and some players (say player 2) this will prevent player 2 from getting item j in one of the scenarios (call it the *bad* scenario) even though it is in her tentative set. Hence in the (random) scenario that player 2 pairs with the bad scenario (lets call it the *worse* scenario), we are committed to give item j to player 2. This might exclude some other player from getting item j in the worse scenario, and this player needs to be compensated in her own pairing (with respect to the worse scenario). However, this commitment to give item j to certain players will not propagate forever: the item constraints imply that in random scenarios only one player is expected to have item j in her tentative set, and this will cause the chain (or rather tree) of commitments to eventually die out.

Our rounding technique is based on the idea above, and is described in a way that makes its analysis simple. It involves 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.

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.

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.*

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.

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

$x_{i,S}$.

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.

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

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. This can be done by choosing an arbitrary vertex of C as a root and orienting every edges away from it. We note that the root vertex can be chosen in a way that depends only on the topology of the tree, independently on the names of vertices. 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.

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.

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.

The following theorem implies (among other things) that with positive probability, the guiding graph rounding technique produces a solution of welfare at least $w(LP)/2$. See also the remarks that follow the proof of the theorem.

Theorem 3.9 *For every $\varepsilon > 0$ (that may depend on the fractional solution of the LP), there is some sufficiently large girth $g(\varepsilon)$ such that if the guiding graph described above is chosen to have girth at least $g(\varepsilon)$, then the welfare of the final solution found by the guiding graph rounding technique is in expectation at least $(1 - \varepsilon)w(LP)/2$.*

The proof of Theorem 3.9 is a consequence of the following three propositions.

Proposition 3.10 *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 3.11 *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. □

Proposition 3.12 *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, recall our notation of $f_{i,j} = \sum_{S|j\in S} x_{i,S}$. As we shall fix j throughout the discussion, we use f_i to denote $f_{i,j}$. By the item constraints, $\sum f_i = 1$. Let us use ϵ to denote $\min_{i|f_i\neq 0}[f_i]$. 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 $\sum f_i = 1$. Thereafter, for every vertex already in the connected components, the expected number of children it has is at most $\sum f_i - \epsilon \leq 1 - \epsilon < 1$. 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 ∞. □

The three propositions above imply Theorem 3.9.

Proof: When g is sufficiently large, Proposition 3.12 implies that condition 2 of Proposition 3.11 holds with probability at least $1-\varepsilon$. Then Proposition 3.11 implies that $E[w_i(S_i)] \geq (1-\varepsilon)w_i(S_{(u,u_i)})/2$, which together with Proposition 3.10 implies:

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

□

A few remarks concerning Theorem 3.9 are in order here.

Integrality gap. Theorem 3.9 implies that the integrality gap of the welfare maximizing LP cannot be worse than 1/2 (when utility functions are subadditive). This follows from the fact that there are only finitely many possible ways of allocating items to players, and hence only finitely many possible values for the welfare. If all these values were strictly below $w(LP)/2$, we could set ε to be sufficiently small so that all these values would also be strictly below $(1-\varepsilon)w(LP)/2$, and this would contradict Theorem 3.9.

An expected polynomial time version. The proof of Theorem 3.9 does not by itself imply that an integer solution of value $w(LP)/2$ can be found 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 3.12, 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 (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 $\sum_{i\neq k; S|j\in S} x_{i,S} \leq 1 - 1/n$. Details omitted.

A faster version. If one is satisfied with an approximation ratio of $(1/2-\epsilon)$, one may speed up the implementation of the guiding graph rounding technique. One may take $g = O(m/\epsilon^2)$ as the girth of the guiding graph G (though the graph need not be constructed explicitly). 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 the total expected number of nodes in $G_1(u), \ldots, G_m(u)$ is $O(m/\epsilon)$, and by Markov's inequality there is only probability ϵ of exceeding the girth of G (and if this happens we may abort). Each vertex has degree n (in G). Altogether, the total number of edges of the guiding graph that are visited is at most $O(nm/\epsilon^2)$, and the approximation ratio is $1/2-\epsilon$.

A slight improvement over $w(LP)/2$. Here we sketch how one can show that the integrality gap of the LP is slightly better than 1/2. We change the guiding graph rounding technique as follows. Rather that orienting the edges independently in each graph G_j, we introduce the following correlations in orientation among these graphs. We pick an edge e in G, and in all G_j (for those items j that are contained in the set labelling edge e) we make e bidirectional. Thereafter, for every G_j and every tree that contains e, we orient the other tree edges away from e. We may iteratively pick new edges e' (not visited in previous steps of the procedure), and repeat the procedure (making these e' bidirectional, and orienting the other edges in their trees). The advantage of this correlated orientation is that if the random center happens to be an endpoint of such a bidirectional edge e, then the player who controls the edge e gets all her tentative set, rather than only part of it. This leads to an expected welfare strictly (though only slightly) better than $w(LP)/2$. Further details are omitted.

A failed attempt. The author also considered the following approach in attempt to get an approximation ratio better than 1/2. Guess (or try all possibilities, there are polynomially many of them) one item j and one player i such that the optimal solution allocates j to i. Allocate item j to player i, and use the LP approach to solve the residual problem that remains (without item j). If one obtains an approximation ratio of 1/2 on the residual problem, then together with item j the approximation ratio is slightly better than 1/2. However, this approach by itself does not work, because in the residual problem, the marginal utility function of player i (given that she already received item j) might no longer be subadditive, and the proof of Theorem 3.9 does not apply anymore.

It may be instructive to compare between the rounding technique introduced in this section and the two step rounding technique of Section 2.2. In both rounding techniques, every player first chooses a tentative set, and the task that remains is that of contention resolution for items that are in more than one tentative set. In the two step rounding technique, contention resolution is performed independently for every item. Though every item has probability at most half of being under contention, the set of items that are under contention is chosen in a correlated way (as not every set can be a tentative set), a fact that is used in Section 3.2.1 to show that this independent contention resolution does not give a constant approximation ratio when utility functions are subadditive. In the approach that we use in this section, we compensate for this correlation among items that are under contention by introducing a correlation between contention resolution of different items. This last correlation is a consequence of the fact that the graphs G_j (used for contention resolution of the different items) are subgraphs of the same guiding graph.

4 Conclusions

We mention here some questions that remain open.

Submodular utility functions. As we have seen, the approximation ratio provided by the welfare maximizing LP depends on the class of utility functions that is involved. When utility functions are subadditive, this ratio is 1/2, when they are fractionally subadditive, this ratio is $1-1/e$, and when they are linear, this ratio is 1. We do not know what the ratio is when utility functions are submodular. Recent work [6] shows that there is some $\epsilon > 0$ such that this ratio is at least $1 - 1/e + \epsilon$, and the ratio is also known to be bounded away from 1, but we are still far from finding matching upper and lower bounds on the approximation ratio (or the integrality gap of the LP) in this case.

Universally good rounding schemes. Recall the rounding technique from Section 3.1 for fractionally subadditive utility functions. Let p_j denote the probability with which item j is allocated. (We took p_j to be $1 - 1/e$ or $(1 - 1/N)^N$ for some large N, but more generally, when $\sum_i f_{i,j} < 1$, the value of p_j can be different.) Then player i gets item j with probability exactly $p_j \frac{f_{i,j}}{\sum_\ell f_{\ell,j}}$. With probability $1 - p_j$ the item remains unallocated. We can modify the rounding scheme so that it does allocate every item. A natural choice would be to allocate a remaining item j to a player i with probability exactly $\frac{f_{i,j}}{\sum_\ell f_{\ell,j}}$. Hence overall, every player i receives every item j with probability $f_{i,j}$ (and even higher, if the item constraints are not satisfied with equality). If player i happened to have a linear utility function (and not just fractionally subadditive), the expected utility of all items allocated to her is at least $w_i(LP)$. That is, we have just designed one oblivious rounding technique that simultaneously guarantees expected utility $(1 - 1/e)w_i(LP)$ for those players that have fractionally subadditive utility functions and $w_i(LP)$ for those players that have linear utility functions. Can this result be extended to subadditive utility functions? Namely, is there one universal rounding technique that simultaneously achieves (in expectation) an approximation ratio of 1 for those players that have a linear utility function, an approximation ratio of $1 - 1/e$ for those players that have a fractionally subadditive utility function, and an approximation ratio of 1/2 for those players that have a subadditive utility function?

Two players. When there are only two players and subadditive utility functions, our results only promise a rounding technique that achieves a welfare of $w(LP)/2$, whereas known integrality gap examples only exclude the possibility of having a rounding technique that obtains welfare better than $\frac{3}{4}w(LP)$. Finding matching upper and lower bounds for this case may lead to new insights about subadditive utility functions.

Acknowledgements

This work was done while the author was a member of the theory group in Microsoft Research, Redmond, Washington. The author thanks Shahar Dobzinski, Kamal Jain, Konstantin Makarychev, Yuri Makarychev, Yishai Mansour, Vahab Mirrokni, Noam Nisan, Michael Schapira and Jan Vondrak for helpful discussions.

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

[4] S. Dobzinski, M. Schapira. "An Improved approximation algorithm for combinatorial auctions with submodular bidders." *Proceedings of 17th SODA*, 1064–1073, 2006.

[5] U. Feige. "On maximizing welfare when utility functions are subadditive". Preliminary version of the current paper, appeared in *Proceedings of 38th STOC*, 41–50, 2006.

[6] U. Feige, J. Vondrak. "Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$." *Proceedings of 47th FOCS*, 667–676, 2006.

[7] L. Fleischer, M. Goemans, V. Mirrokni, M. Sviridenko. "Tight approximation algorithms for maximum general assignment problems." *Proceedings of 17th SODA*, 611–620, 2006.

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

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