On Allocations that Maximize Fairness

Uriel Feige*

Abstract

We consider a problem known as the restricted assignment version of the max-min allocation problem with indivisible goods. There are n items of various nonnegative values and m players. Every player is interested only in some of the items and has zero value for the other items. One has to distribute the items among the players in a way that maximizes a certain notion of fairness, namely, maximizes the minimum of the sum of values of items given to any player. Bansal and Sviridenko [STOC 2006] describe a linear programming relaxation for this problem, and present a rounding technique that recovers an allocation of value at least $\Omega(\log\log\log m/\log\log m)$ of the optimum. We show that the value of this LP relaxation in fact approximates the optimum value to within a constant factor. Our proof is not constructive and does not by itself provide an efficient algorithm for finding an allocation that is within constant factors of optimal.

1 Introduction

We study a problem related to fair allocation of indivisible goods. There are m players and n items. Let p_{ij} denote the value of item j to player i. We assume throughout that $p_{ij} \geq 0$ for all $1 \leq i \leq m$ and $1 \leq j \leq n$. For a bundle S of items, its value to player i is $\sum_{j \in S} p_{ij}$, namely, a linear function of the values of the items for player i. The goal is to allocate all items to the players (each item to exactly one player) in a way that optimizes some objective function. Various objective functions have been studied in this context. Some of them attempt to maximize value (e.g., when items are goods that the players really appreciate having), and some attempt to minimize value (e.g., when the items are jobs, the players are machines, and the p_{ij} are running times). The objective in the current work is a maximization one, and is motivated by notions of fairness, ensuring that every player gets at least some target level of "quality of service". Namely, the objective will be to partition the n items into m disjoint bundles, $S_1, \ldots, S_m$, such that $\min_i \sum_{j \in S_i} p_{ij}$ is maximized. This problem is known as the max-min allocation problem with indivisible goods, and also as the "Santa Claus" problem (each player is a kid, each item is a gift, some kids do not want some of the gifts, and Santa Clause wants all kids to be happy with the bundle of gifts that they get). The decision version of this problem is as follows. Given

*Department of Computer Science and Applied Mathematics, the Weizmann Institute, Rehovot 76100, Israel. uriel.feige@weizmann.ac.il

some threshold t, determine whether there is an allocation of the items to the players such that each player receives a value of at least t.

The max-min allocation problem is NP-hard. (This is true for indivisible goods.) Our goal will be to design good approximation algorithms for it. For the decision version of the problem, an *estimation algorithm* needs to estimate the maximum value t (up to a multiplicative factor that is the estimation ratio) for which the underlying decision problem has a feasible solution. For the optimization version (the search version), an *approximation algorithm* needs to also find an allocation for which the value of the objective function is close to optimum (up to a multiplicative factor that is the approximation ratio). All previous approximation algorithms that we shall refer to here deal with the more difficult goal of approximating the search version of the paper. All previous hardness results apply also to the easier goal of estimating the decision version.

For the general case of the max-min fair allocation problem, Bezakova and Dani [5] show an additive approximation of $\max_{ij} p_{ij}$. That is, their algorithm is guaranteed to output a solution of value at least $(\mathrm{opt} - \max_{ij} p_{ij})$, where opt denotes the value of the optimal solution. Though this result appears to be very strong, it turns out that the challenging cases of the problem are when $\mathrm{opt} \leq \max_{ij} p_{ij}$, in which case the results of [5] do not offer any guarantees at all. Bansal and Sviridenko [4] considered a certain linear programming relaxation for the problem (the *configuration* LP that will be reviewed shortly), and showed how it can be used in order to find a solution of value $\Omega(\mathrm{opt}/m)$ (recall that m is the number of players). They also showed an integrality gap of $O(1/\sqrt{m})$ for this LP, implying in particular that its value is no better than an $O(1/\sqrt{m})$ approximation to opt. Asadpour and Saberi [3] improved over the results of [4], and showed how to round the configuration LP so as to get a solution of value at least $\Omega(\mathrm{opt}/\sqrt{m}(\log m)^3)$.

A special case of the max-min allocation problem, known as the *restricted assignment* version, is one in which every item j has an intrinsic value p_j, but there are players to which this item cannot be assigned. Equivalently for our purpose, for every player i, either $p_{ij} = p_j$ or $p_{ij} = 0$. For the restricted assignment case, Bansal and Sviridenko [4] showed that the configuration LP can be rounded to give a feasible solution of value $\Omega(\mathrm{opt} \log\log\log m/ \log\log m)$. Bezakova and Dani [5] showed that already this case is NP-hard to approximate within ratios better than 1/2.

An even more special case is the *uniform* case, in which each item j has an intrinsic value p_j, and $p_{ij} = p_j$ for all players i. Woeginger designs a PTAS for the uniform case [7].

In the current paper, we revisit the restricted assignment case. We show the following.

Theorem 1.1 *In the restricted assignment case, there is a polynomial time algorithm that estimates the optimal value of max-min allocation problem within a constant factor.*

Our proof shows that in the restricted assignment case, there is a feasible solution of value at most some constant factor smaller than the optimal value of the configuration LP. Combining this with the facts that the value of the configuration LP is an upper bound on opt, and that the configuration LP can be solved in polynomial time [4], this shows that the value of the configuration LP provides the approximation desired by Theorem 1.1.

Unlike previous approximation algorithms for the max-min allocation problem, our approximation applies only to the decision version of the problem, but not to the search version. We only prove the existence of a feasible solution of value close to that of the configuration LP, but do not provide an efficient (polynomial time) algorithm for finding such a solution.

Subsequently to the work reported here, a different non-constructive proof for Theorem 1.1 was found [2]. It shows that the integrality gap for the configuration LP in the restricted assignment case is no worse than 5.

1.1 The configuration LP

Though the main technical contributions of this paper can be understood without referring to the configuration LP, we present the configuration LP for completeness.

Let t be a threshold value for the decision problem. The configuration LP is feasible whenever $t \leq \text{opt}$ (and sometimes also for some values $t > \text{opt}$). To use the LP most effectively, one first needs to select a feasible value for t which is (nearly) as large as possible. In the restricted assignment version of the problem, this can be done by sorting all items in order of decreasing p_j, and finding the smallest index k such that there is a matching involving all players and (some of) the first k items, where player i can be matched to item j if $p_{ij} = p_j$. It then follows that $p_k \leq \text{opt} \leq p_k n$. Thereafter, one may start with $t = p_k$, and by repeated doubling find a factor two approximation for the maximum value of t for which the LP is feasible.

The configuration LP has exponentially many variables and exponentially many constraints. For every player i and bundle S of items, there is a variable x_{iS}. The intention is that these will be indicator 0/1 variables, having value 1 iff bundle S is allocated to player i. However, to make the LP solvable in polynomial time, the 0/1 constraints are relaxed to nonnegativity constraints. In the configuration LP, p_{iS} is used as shorthand notation for $\sum_{j \in S} p_{ij}$. The linear constraints of the configuration LP are the following:

1. Every player gets a bundle of value at least t: $X_{iS} = 0$ whenever $p_{iS} < t$.
2. Every player gets one bundle: $\sum_S x_{iS} = 1$ for every player i.
3. Every item is allocated to at most one player: $\sum_{i,S|j \in S} x_{iS} \leq 1$ for every item j.
4. Relaxation of 0/1 constraints: $x_{iS} \geq 0$ for every player i and set S.

Every feasible allocation of the items is a feasible solution to the LP (by setting the corresponding variables x_{iS} to 1). Hence t can be raised to a value of at least opt. For every value of t, the feasibility of the configuration LP can be checked in polynomial time. Details appear in [4].

1.2 The key technical theorem

Bansal and Sviridenko [4] analyzed the configuration LP for the restricted assignment case. They showed a rounding technique that produces a feasible solution of value at least $\Omega(t \log\log\log m / \log\log m)$. Moreover, they identified a key combinatorial problem that characterizes (up to constant factors) the true integrality gap of the LP. This combinatorial problem can be viewed as a special case of the restricted assignment version of the max-min allocation problem. We now proceed to describe this special case.

The m players are partitioned into m/ℓ groups, where each group has ℓ players. For each group, there are $\ell - 1$ *expensive* items of value k that only members of the group value, but no one else. In addition, there are many *cheap* items of value 1. Every player values exactly k of the cheap items (and has value 0 for the remaining cheap items). Every cheap item is valued by exactly ℓ of the players. (The number n' of cheap items is chosen so as to satisfy $n'\ell = mk$.)

In the above setting, the configuration LP has the following feasible solution of value $t = k$. The variables x_{iS} have value $1/\ell$ in two cases: when S is the set of k cheap items valued by player i, and when S is a set containing exactly on expensive item valued by player i. Hence for every player i, there are exactly $1 + (\ell - 1) = \ell$ sets S for which $x_{iS} = 1/\ell$, for every expensive item j, there are exactly ℓ players i for which $x_{i\{j\}} = 1/\ell$, and for every cheap item j there are exactly i players that have a set S of cheap items that contains j for which $x_{iS} = 1/\ell$.

If the configuration LP is to have an integrality gap no worse than a constant (in the restricted assignment case), then there must be a feasible solution in which each player gets a value of $\Omega(k)$. In each group, $\ell - 1$ of the players can be satisfied by an expensive item. Hence we need to find a way of choosing one player from each group, and giving this player $\Omega(k)$ cheap items (without giving the same cheap item to two players). Bansal and Sviridenko [4] show that this problem is the only obstacle to proving a constant integrality gap for the configuration LP. Namely, if in the special case there always is a feasible solution of value $\Omega(k)$, then the integrality gap of the configuration LP is no worse than constant for every instance of the restricted assignment case.

In the above description of the special case, we slightly simplified the combinatorial problem that was characterized in [4]. The following definition gives a more complete description. It introduces an additional parameter β, and allows the instance to be somewhat less regular than in the version presented above. Setting $\beta = 1$ and $\gamma < 1$ to be a universal constant, it should not be difficult to see that sets in the definition below correspond to players in the special case explained above.

Definition 1.2 *A (k, ℓ, β) system is defined as follows. There are m sets parti-*

tioned into m/ℓ groups of sets, where each group contains exactly ℓ sets. Every set contains exactly k items. Every item is contained in at most $\beta\ell$ sets. A (k, ℓ, β) system is γ-good if there is a choice of one set from each group and $\lfloor \gamma k \rfloor$ items from every chosen set such that all chosen items are distinct.

All that we require from the parameter m in the above definition is that it is divisible by ℓ. We do not carry it around as part of the notation, because the notation is established for use in Section 2, and there the value of parameter m has no effect whatsoever on the proofs.

In [4] it is shown that for $\beta = O(1)$, every (k, ℓ, β) system is γ-good for $\gamma = \Omega(\log \log \log m / \log \log m)$. Our main result is as follows.

Theorem 1.3 *Every (k, ℓ, β) system is γ-good for $\gamma = \frac{\alpha}{\max[1,\beta]}$, where $\alpha > 0$ is some fixed universal constant. In particular, when $\beta = O(1)$, γ is some fixed positive constant independent of k, ℓ and m.*

In our proof of Theorem 1.3 we will not present an explicit value of α, especially as we prefer simplicity over trying to optimize α. Observe however that for the theorem to be correct, α needs to be at most $\frac{1}{2}$. For example, when $\beta = 1$ and $k = \ell$, we may have a case where for each set S of the first group, there is some companion group (that depends on S) in which each set contains $k - 1$ of the items of S. Then regardless of which set S is chosen from the first group, the number of distinct items in the union of S and the set chosen from its companion group is $k + 1$, giving $\gamma \leq \lfloor \frac{k+1}{2} \rfloor \frac{1}{k}$. When k is small and $\beta > 1$, one can design examples were α is strictly smaller than $\frac{1}{2}$. For $\beta = \frac{4}{3}$ and $k = \ell = 3$, having four identical groups forces γ to be $\frac{1}{3}$, giving $\alpha \leq \frac{4}{9}$.

Our proof of Theorem 1.3 combines ideas from [4] with iterative applications of the Lovasz local lemma, an approach inspired by [6]. Throughout we use known probabilistic techniques without proof. The reader is referred to standard references (such as [1]) for more details on these techniques.

2 Proof of the main result

For the proof of Theorem 1.3 we may assume that $\beta \geq 1$ (because of the term $\max[1, \beta]$). Moreover, we may assume that $k \geq \beta$, because for $k < \beta$ the value of γ claimed by the theorem is such that $\lfloor \gamma k \rfloor = 0$ and there is nothing to prove. Hence throughout we assume that $k \geq \beta \geq 1$, without explicitly repeating this assumption in every place where it is used.

Before embarking on the proof of Theorem 1.3, let us discuss shortly a simple approach that does not work. Assume that one selects from each group one set at random. A selected set T contains k items. In expectation, each one of these items is expected to be contained in $\beta\ell/\ell = \beta$ other selected sets. Resolving conflicts at random among the competing sets, T is expected to keep $k/(\beta + 1)$ of its items. (In fact, the expectation is higher than $k/(\beta + 1)$, due to variance in the number of players competing for an item.) This suggests, that one should be able to choose $\alpha \simeq 1/2$ in Theorem 1.3.

Technically, the above simple approach is lacking in the following sense. It does not suffice to analyze the expectation. One needs to also analyze the probability of large deviations from expectation. This probability might be large regardless of whether k itself is large or small. Large overlap between different sets can cause competition for different items in T to be highly correlated, defeating attempts to use bounds on large deviations that are true for sums of independent random variables.

In Section 2.1 we explain how to prove Theorem 1.3 in the case that either k or ℓ are constant. In Section 2.2 we present without proofs three lemmas and explain how they imply the proof of Theorem 1.3 in the general case (when k and ℓ depend on n). In Section 2.3 we show how each of the three lemmas of Section 2.2 can be proved by using the Lovasz local lemma.

2.1 Constant k or ℓ

The following lemma proves Theorem 1.3 in cases where $k = O(1)$ and more generally, when $k = O(\beta)$.

Lemma 2.1 *Every (k, ℓ, β) system is γ-good for $\gamma = \frac{1}{k}$.*

Proof: Consider the bipartite graph with items on the left-hand side and groups of sets on the right-hand side. We connect item i to group g by p parallel edges, where p (which is possibly 0) is the number of sets in group g that contain item i. Observe that the degree of every lefthand side vertex is at most $\beta\ell$, whereas the degree of every righthand side vertex is $k\ell \geq \beta\ell$ (by our assumption that $k \geq \beta$). Hall's theorem then implies that there is a matching involving all right-hand side vertices. This matching selects one item from each group, with no conflicts. From each group, pick the set that contains the chosen item. $\square$

The following lemma is useful when ℓ is a constant.

Lemma 2.2 *Every (k, ℓ, β) system is γ-good for γ satisfying*

$$\gamma k = \lfloor \frac{k}{\lceil \beta\ell \rceil} \rfloor$$

Proof: Pick the first set from each group. Partition every such set into subsets of size $\lceil \beta\ell \rceil$, removing any items leftover if $\lceil \beta\ell \rceil$ does not divide the size of the set. Consider the bipartite graph with items on the left-hand side and subsets on the right-hand side, with edges connecting items to subsets that contain them. The degree of every lefthand side vertex is at most $\beta\ell$, whereas the degree of every righthand side vertex is $\lceil \beta\ell \rceil$. Hall's theorem then implies that there is a matching involving all right-hand side vertices. This matching selects one item from each subset, with no conflicts. $\square$

2.2 The proof plan

Our plan for proving Theorem 1.3 is to reduce ℓ below a value of c (where c is some universal constant, related to the inverse of α from Theorem 1.3) without significantly increasing β and γ, and then to appeal to Lemma 2.2. This plan is implemented by iteratively switching between two transformations, one transformation (stated in Lemma 2.4) that reduces k (while slightly increasing γ) and one transformation (stated in Lemma 2.3) that reduces ℓ (while slightly increasing β). The reason for using two types of transformations is that the one for reducing k requires ℓ to be smaller than k (or at least not much larger than k), and the one for reducing ℓ requires k to be smaller than ℓ (or at least not much larger than ℓ).

It will turn out that applying the plan above in the case that β is not a constant (but some function of n) results in γ having a slightly super-linear dependence on $1/\beta$, rather than a linear dependence as claimed in Theorem 1.3. This has to do with the fact that Lemma 2.4 does not effect γ in a significant way only if $k > \frac{1}{\gamma}\log^2\frac{1}{\gamma}$, implying that it can be applied only when is k is super-linear in β. Hence the approach described above will fail to make $k = O(1/\gamma)$, and will also fail to make $\ell < c$. To overcome this problem, we will use a transformation that reduces the value of β (stated in Lemma 2.5).

We now present our three key lemmas. Their proofs will appear in Section 2.3. Throughout, we use c to denote a sufficiently large constant.

When k is not too large compared to ℓ, then the following lemma shows that ℓ can be reduced significantly.

Lemma 2.3 *For $\ell > c$, every (k, ℓ, β) system with $k \leq \ell$ can be transformed into a (k, ℓ', β') system with $\ell' \leq \log^5 \ell$ and $\beta' \leq \beta(1 + \frac{1}{\log \ell})$.*

Observe that even repeated applications of Lemma 2.3 will not cause β' to increase beyond twice the original value of β (if one stops when $\ell' < c$).

Lemma 2.3 can be applied only when k is not too large compared to ℓ. We couple Lemma 2.3 with another lemma that allows us to reduce k.

Lemma 2.4 *Every (k, ℓ, β) system with $k \geq \ell \geq c$ can be transformed into a (k', ℓ, β) system with $k' \leq k/2$ and with the following additional property: if the original system is not γ-good, then the new system is not γ'-good for $\gamma' = \gamma(1 + \frac{3\log k}{\sqrt{\gamma k}})$. Conversely, if the new system is γ'-good, then the original system was γ-good.*

Observe that even with repeated applications of Lemma 2.4, the value of γ' reaches at most twice the original value of γ, as long as $k' > \max[c, \frac{1}{\gamma^2}]$.

The following lemma allows us to reduce β. As noted above, this will be used in order to get a linear dependency of γ in $1/\beta$.

Lemma 2.5 *For $\ell > c$, every (k, ℓ, β) system can be transformed into a (k', ℓ, β') system with $k' = \lfloor \frac{k}{2} \rfloor$ and $\beta' \leq \frac{\beta}{2}(1 + \frac{\log(\beta\ell)}{\sqrt{\beta\ell}})$.*

We can now complete the proof of Theorem 1.3.

Proof: Let c be a sufficiently large constant, as required in lemmas 2.4, 2.3 and 2.5. If $\ell \leq c$, then the theorem follows from Lemma 2.2. Hence we may assume that $\ell > c$. Likewise, if $k \leq \max[c, 4\beta]$ then the theorem follows from Lemma 2.1. Hence we may assume that $k > c$ and $k > 4\beta$.

Starting from the original system, repeatedly apply Lemma 2.5 until the value of β is reduced below 1. In this process, k is reduced essentially at the same rate as β, and hence reaches a value $k' \geq \frac{k}{3\beta}$. Any value of γ that we prove for the new system translates into at least $\frac{\gamma}{3\beta}$ for the original system. Hence it suffices to prove a value of $\gamma = \Omega(1)$ for the new system.

Starting from the new system, as long as $\ell > c$, apply the transformation of Lemma 2.3 when $\ell > k$ and the transformation of Lemma 2.4 when $k \geq \ell$. (Here, unlike the statements of Lemma 2.3 and 2.4, we assume that the new parameters k, ℓ, β retain the names of older parameters after each iteration, even though they do not retain their value.) For the final system, $\ell < c$ and $k \geq \ell^3$. This last property follows because the last transformation must have been that of Lemma 2.3 which reduces ℓ to polylogarithmic factors, whereas the transformations of Lemma 2.4 reduce k by a factor of roughly two, and hence ℓ was at most roughly $2k$ before the last transformation. (Technical remark relating to the above argument: If the first transformation to be applied reduces ℓ, one might need to temporarily increase ℓ' in Lemma 2.3 so that it does not happen that ℓ is reduced below c before a transformation reducing k gets to be applied. Details left to the reader.) Observe that during the transformations that reduce ℓ the value of β increased by a total factor of at most 2, and hence its value is at most 2. Lemma 2.2 implies that (up to rounding errors) for the final system $\gamma \geq \frac{1}{2\ell} \geq \frac{1}{2c}$. Moreover, $k \geq \ell^3 > \gamma^{-2}$ so that the total increase in γ from the new system to the final system (due to transformations that decrease k) is at most a factor of 2. Hence for the new system $\gamma \geq \frac{1}{4c}$, and for the original system $\gamma \geq \frac{1}{12c\beta}$. □

2.3 Proofs of Lemmas

In this section we provide the proofs of the lemmas that were stated in Section 2.2.

Proof of Lemma 2.3.

Proof: The proof follows easily from the uniform version of the Lovasz local lemma. Pick $\lfloor \log^5 \ell \rfloor$ sets at random from every group. The value of k remains unchanged by this choice, whereas ℓ changes to ℓ' as desired. It remains to analyze the new value of β. Per item, in expectation, the value of β remains unchange. Standard application of the Chernoff bound implies that for any item, the probability that more than $\beta\ell' + t\sqrt{\beta\ell'}$ copies of it remain is exponentially small in t^2. In particular, the probability of an item leading to $\beta' \geq \beta(1 + \frac{1}{\log \ell})$ is exponentially small in $\log^3 \ell$.

To use the local lemma, let a bad event be the event that more than

$\beta\ell'(1+\frac{1}{\log \ell})$ copies of some item survive. Consider the dependency graph among items. Its degree is bounded by $k\ell\beta\ell \leq \ell^4$, because an item participates in at most $\beta\ell$ groups, and each group contains at most $k\ell$ other items. For large enough c, the local lemma applies. $\square$

Proof of Lemma 2.4.

Lemma 2.4 will be proved by using a general version of the Lovasz local lemma, which we state here without proof.

Lemma 2.6 *Let $B_1, \ldots, B_t$ be "bad" events, and let $G(V,E)$ be a dependency graph for them, in which for every i, event B_i is mutually independent of all events B_j for which $(i,j) \notin E$. Let x_i for $1 \leq i \leq t$ be such that $0 < x_i < 1$ and $Pr[B_i] \leq x_i \prod_{(i,j)\in E}(1-x_j)$. Then with positive probability no event B_i holds.*

We shall also use a characterization related to goodness of (k, ℓ, β) systems. In the following, we call a collection of sets *connected* if their intersection graph is connected. (The intersection graph has the sets as its vertices, and edges connect sets that share an item.)

Lemma 2.7 *Consider a collection of n sets and a positive integer q.*

1. *If for some $1 \leq i \leq n$, there is a connected subcollection of i sets whose union contains less than iq items, then there is no choice of q items per set such that all items are distinct.*

2. *If for every i, $1 \leq i \leq n$, the union of every connected subcollection of i sets contains at least iq (distinct) items, then there is a choice of q items per set such that all items are distinct.*

Proof: Case 1 follows from the fact that the number of distinct items in the subcollection is too small so as to give q items to every set in the subcollection. In order to prove case 2, consider the following bipartite graph. The lefthand side vertices are items. The righthand side vertices are the sets, where each set has q copies. Edges connect items to sets that contain them. Every set of righthand side vertices has at least as many neighbors on the lefthand side, and hence by Hall's theorem there is a matching involving all righthand side vertices. This gives a choice of q distinct items per set. $\square$

We are now ready to prove Lemma 2.4.

Proof: Consider an arbitrary (k, ℓ, β) system. Perform a random experiment in which every item survives independently with probability $1/2$. Consider the following two types of bad events.

1. B_1: some set has less than $k' = (1 - \frac{\log k}{\sqrt{k}})\frac{k}{2}$ items surviving.

2. B_i for $i \geq 2$: there is a connected collection of i sets from distinct groups whose union originally contained at most $i\gamma k$ items, of which more than $i\delta'\frac{k}{2}$ items survive, where $\delta' = (1 + \frac{\log k}{\sqrt{\gamma k}})\gamma$.

If none of the bad events happened, we can remove from every set all but its first k' items. Then every set will contain exactly k' items, and moreover, bad events of the second type still do not occur. In combination with Lemma 2.7, this implies that every collection of sets that shows that the new system is γ'-good (for $\gamma' = \delta' \frac{k}{2} \frac{1}{k'} \le (1 + \frac{3 \log k}{\sqrt{\gamma k}})\gamma$) also certifies that the the original system was γ-good.

The proof of Lemma 2.4 will follow by showing (using Lemma 2.6) that with positive probability none of the bad events specified above happens.

Observe that by the Chernoff bound, any event of type B_1 happens with probability at most $2^{-\Omega(\log^2 k)} \le 2^{-20 \log k}$, where the inequality follows because k is sufficiently large, as dictated by assuming that c is sufficiently large. Any event of type B_i for $i \ge 2$ happens with probability at most $2^{-\Omega(i \log^2 k)} \le 2^{-20 i \log k}$ (again the inequality follows from our choice of large enough c). We now associate values x_i with the bad events. For events of type B_1 we set $x_1 = 2^{-10 \log k}$. For events of type B_i for $i \ge 2$ we set $x_i = 2^{-10 i \log k}$.

It remains to analyze the dependency graph. Consider an arbitrary event E_i of type B_i, for $i \ge 1$. It involves i sets and at most ik items (and in fact, at most $i\gamma k$ items if $i > 1$). The number of events of type B_1 that share an item with it is at most $ik\beta\ell \le ik^3$. All other events of type B_1 are mutually independent of E_i.

Consider now events of type B_j for $j \ge 2$ that share an item with E_i. There are at most ik^3 choices for the set S that shares an item with E_i. The j sets involved in the event B_j must induce a connected subgraph in the graph H that connects any two sets that share an item. Fix a tree rooted at S over this subgraph. Observe that there are less than 2^{2j} possible tree structures. For example, touring the tree in depth first search (DFS) order, we need only to specify which of the $2j-2$ moves is a backward move. Given the tree structure, for each new set encountered in a DFS traversal of the tree, there are at most $k\beta\ell \le k^3$ choices. Altogether, there are at most $ik^3 2^{2j} (k^3)^{j-1} = i(4k^3)^j$ events of type B_j that depend on E_i.

With the above choices for the x_i, the conditions of Lemma 2.6 indeed hold. This follows because for every $i \ge 1$ and every event E_i of type B_i,

$$
\begin{aligned}
Pr[E_i] &\le 2^{-20 i \log k} \\
&\le 2^{-10 i \log k} \prod_{j \ge 1} (1 - 2^{-10 j \log k})^{i(4k^3)^j} \\
&\le x_i \prod_{j \ge 1} \prod_{*} (1 - x_j)
\end{aligned}
$$

where * denotes the statement "events of type B_j that depend on E_i". □

Proof of Lemma 2.5.

Proof: In every set, arrange all items arbitrarily in pairs (discarding the last item if k is odd). From each pair, select independently at random one of the two items to remain in the set, and discard the other. (Unlike Lemma 2.4, the item is discarded only from the particular set and not from all sets.)

For every item, the expected number of copies of it that remain is at most $\frac{\beta\ell}{2}$. The probability of deviating by $t\sqrt{\beta\ell}$ is exponentially small in t^2. Each item participates in at most $\beta\ell$ pairs, and hence is independent of all but at most $\beta\ell$ other items. Setting $t = \frac{1}{2}\log(\beta\ell)$ we can apply the uniform version of the local lemma to conclude that with positive probability, for every item, at most $\frac{\beta}{2}(1 + \frac{\log(\beta\ell)}{\sqrt{\beta\ell}})$ copies remain. $\square$

2.4 Conclusions

Our proof for Theorem 1.3 is nonconstructive. It applies the local lemma for three different purposes: the uniform version to reduce β in Lemma 2.5 and to reduce ℓ in Lemma 2.3, and the general version to reduce k in Lemma 2.4. Our proof does not produce (by itself) a polynomial time algorithm for deciding which sets to choose from each group. Though there are known methods for making the local lemma constructive (see for example [1] and references therein), it is not clear whether any of them can be applied in our case. The main source of difficulty in this respect is Lemma 2.4, because there the number of bad events is exponential in the problem size, and moreover, there are bad events that involve a constant fraction of the random variables.

Recently Asadpour, Feige and Saberi [2] discovered an alternative proof for Theorem 1.1 (and hence also for Theorem 1.3). Interestingly, also this new proof is nonconstructive, even though it does not use the local lemma at all.

Acknowledgements

This work was done while the author was a member of the theory group at Microsoft Research, Redmond, Washington. I thank Arash Asadpour and Amin Saberi for useful discussions.

References

[1] N. Alon and J. Spencer, *The Probabilistic Method*, Second Edition, Wiley-Interscience, 2000.

[2] A. Asadpour, U. Feige and A. Saberi, Manuscript, 2007.

[3] A. Asadpour and A. Saberi, *An Approximation Algorithm for Max-Min Fair Allocation of Indivisible Goods*, Proc 39th STOC, 114–121, 2007.

[4] N. Bansal and M. Sviridenko, *The Santa Clause problem*, Proc 38th STOC, 31–40, 2006.

[5] I. Bezakova and V. Dani, *Allocating indivisible goods*, SIGecom Exchanges, Vol. 5.3, 2005.

[6] F. T. Leighton, B. Maggs and S. Rao, *Packet Routing and Job-Shop Scheduling in O(Congestion + Dilation) Steps*, Combinatorica 14(2): 167–186, 1994.

[7] G. Woeginger, *A Polynomial-time approximation scheme for maximizing the minimum machine completion time*, Oper. res. Lett., 20(4):149–154, 1997.