Improved Integrality Gap in Max-Min Allocation: or Topology at the North Pole*

Penny Haxell[†] Tibor Szabó[‡]

January 16, 2023

Abstract

In the max-min allocation problem a set P of players are to be allocated disjoint subsets of a set R of indivisible resources, such that the minimum utility among all players is maximized. We study the restricted variant, also known as the Santa Claus problem, where each resource has an intrinsic positive value, and each player covets a subset of the resources. Bezáková and Dani [15] showed that this problem is NP-hard to approximate within a factor less than 2, consequently a great deal of work has focused on approximate solutions. The principal approach for obtaining approximation algorithms has been via the Configuration LP (CLP) of Bansal and Sviridenko [12]. Accordingly, there has been much interest in bounding the integrality gap of this CLP. The existing algorithms and integrality gap estimations are all based one way or another on the combinatorial augmenting tree argument of Haxell [26] for finding perfect matchings in certain hypergraphs.

Our main innovation in this paper is to introduce the use of topological methods, to replace the combinatorial argument of [26] for the restricted max-min allocation problem. This approach yields substantial improvements in the integrality gap of the CLP. In particular we improve the previously best known bound of 3.808 to 3.534. We also study the $(1, \varepsilon)$-restricted version, in which resources can take only two values, and improve the integrality gap in most cases. Our approach applies a criterion of Aharoni and Haxell, and Meshulam, for the existence of independent transversals in graphs, which involves the connectedness of the independence complex. This is complemented by a graph process of Meshulam that decreases the connectedness of the independence complex in a controlled fashion and hence, tailored appropriately to the problem, can verify the criterion. In our applications we aim to establish the flexibility of the approach and hence argue for it to be a potential asset in other optimization problems involving hypergraph matchings.

*A preliminary version of this paper appeared in the proceedings of the *Symposium on Discrete Algorithms (SODA) 2023)*.

[†]Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON, Canada. Partially supported by NSERC and by a Mercator Fellowship of the Research Training Group *Facets of Complexity*.

[‡]Freie Universität Berlin, Berlin, Germany. Research partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689) and the Research Training Group *Facets of Complexity*.

1 Introduction

In this paper we consider the *restricted max-min allocation* problem. An instance $\mathcal{I} = (P, R, v, \{L_p : p \in P\})$ of the problem consists of a set P of players, a set R of indivisible resources, where each resource $r \in R$ has an intrinsic positive value $v_r > 0$, and each $p \in P$ covets a set $L_p \subseteq R$ of resources. An *allocation* of the resources is a function $a : P \to 2^R$, with $a(p) \subseteq L_p$ for each $p \in P$, such that every resource is allocated to (at most) one player, that is $a(p) \cap a(q) = \emptyset$ for every $p \neq q$. The *min-value of allocation* a is $\min_{p \in P} v(a(p))$, where for a set $S \subseteq R$ of resources $v(S) = \sum_{r \in S} v_r$ represents the total value of S. The objective is to maximize the min-value over all allocations of resources. This value will be denoted by $OPT = OPT(\mathcal{I})$.

The choice of a max-min objective function is arguably a good one for achieving overall individual "fairness" in the distribution of a set of indivisible resources that are considered desirable by the players.[1] Since the seminal paper of Bansal and Sviridenko [12], the restricted max-min allocation problem often goes under the name *Santa Claus Problem*, where the players represent children, and the resources are presents to be distributed by Santa Claus. One imagines each present r having a "catalogue" value v_r, but some presents may not be interesting to some children.[2] To be fair[3], Santa might wish to distribute the presents so that the smallest total value received by any child is as large as possible.

The problem of how to find an optimal solution efficiently was studied first in the special case when $L_p = R$ for every player $p \in P$. In this case Woeginger [40] and Epstein and Sgall [22] gave polynomial time approximation schemes (PTAS), and Woeginger [41] gave an FPTAS when the number of players is constant. For the general case however, Bezáková and Dani [15] showed that the problem is hard to approximate up to any factor < 2. On the positive side, there has been a great deal of progress towards finding good approximations. In [15] an approximation ratio of $|R| - |P| + 1$ is achieved, as well as an additive approximation algorithm using the standard assignment LP relaxation of the problem. This finds a solution of value at least $T_{ALP} - \max_{r \in R} v_r$, where T_{ALP} is the optimal value of the assignment LP. This algorithm however does not offer any approximation factor guarantee when $\max_{r \in R} v_r$ is large.

To address the fact that the assignment LP can have arbitrarily large integrality gap in general, Bansal and Sviridenko [12] introduced the important innovation of using a stronger LP, called the configuration LP for the problem, which we now describe. Given a problem instance $\mathcal{I}$ and $T \geq 0$, for each player $p \in P$ we define the family $\mathcal{C}_p(T) = \{C \subseteq L_p : v(C) \geq T\}$ of *configurations for* p. The *configuration LP* for $\mathcal{I}$ with *target* T has a variable $x_{p,S} \geq 0$ for every player $p \in P$ and configuration $S \in \mathcal{C}_p(T)$, and a constraint

$$\sum_{S \in \mathcal{C}_p(T)} x_{p,S} \geq 1$$

for every player $p \in P$ and a constraint

$$\sum_{p \in P} \sum_{S \in \mathcal{C}_p(T), S \ni r} x_{p,S} \leq 1$$

for every resource $r \in R$.

[1] This is in contrast with the situation where resources are considered rather "chores", when one would usually aim to minimize the maximum values of the subsets of resources allocated to each player. That would be the setup for example in the classical makespan minimization problem, where various jobs have to be allocated to a set of machines.

[2] ... since perhaps they already secured the latest edition of their favorite smartphone for their birthday.

[3] ... and to avoid criticism from jealous parents

We will refer to this LP as CLP(T) for $\mathcal{I}$. Formally we minimize the objective function 0, but the main point is whether CLP(T) is feasible. For a given instance $\mathcal{I}$, let $T^* = T^*(\mathcal{I})$ be the maximum T for which CLP(T) is feasible. It is a striking fact from [12] that even though CLP(T) has exponentially many variables, T^* can be approximated up to any desired accuracy in polynomial time. Note that any allocation for $\mathcal{I}$ of min-value T' gives an (integer) feasible solution to CLP(T'). Hence $OPT \leq T^*$. We will refer to T^*/OPT as the *integrality gap*.

Using their configuration LP, Bansal and Sviridenko [12] obtained an $O(\log\log|P|/\log\log\log|P|)$-approximation algorithm for the Santa Claus problem. They also formulated a combinatorial conjecture and connected it to the problem of finding an allocation with large min-value given a feasible solution of CLP(T). Feige [23] proved this conjecture via repeated applications of the Lovász Local Lemma and hence established a constant integrality gap for the CLP. This was later made algorithmic by Haeupler, Saha, and Srinivasan [25] using Local Lemma algorithmization, which provided the first (huge, but) constant factor approximation algorithm for the Santa Claus problem.

Asadpour, Feige, and Saberi [10] formulated the problem in terms of hypergraph matching and proved an upper bound of 4 on the integrality gap of the CLP. Via the machinery of [12] this result implies an efficient algorithm to estimate the value of OPT up to a factor $(4+\delta)$. The approach of [10] is based on a local search technique introduced by Haxell [26], where the corresponding procedure is not known to be efficient. Polacek and Svensson [37] modified the local search of [10] and were able to prove a quasi-polynomial running time for a $(4+\delta)$-approximation algorithm. Finally, Annamalai, Kalaitzis, and Svensson [9] managed to adapt the local search procedure to terminate in polynomial time, introducing several influential novel ideas, which resulted in a polynomial time 12.33-approximation algorithm. Subsequently Cheng and Mao [18] altered the algorithm to establish a $(6+\delta)$-approximation guarantee, improving further in [20] to obtain a $(4+\delta)$-approximation algorithm. Davies, Rothvoss, and Zhang [21] also gave an $(4+\delta)$-approximation algorithm, working in a more general setting, where a matroid structure is imposed on the players. The integrality gap of the configuration LP was further improved by Cheng and Mao [19] and Jansen and Rohwedder [32] to 3.833 and then to 3.808 by Cheng and Mao [20] by better and better analysis of the procedure of [10].

A special case of the problem, that already captures much of its difficulty, comes from limiting the number of distinct values taken by resources to two. In the $(1,\varepsilon)$-restricted allocation problem resources can take only two values 1 or ε, where $0 < \varepsilon \leq 1$. The relevance of this case is also underlined by the fact that a key reduction step in the foundational result of [12] required an approximation algorithm for the $(1,\varepsilon)$-restricted allocation problem for arbitrarily small $\varepsilon > 0$.

Chan, Tang, and Wu [17], extending work of Golovin [24] and Bezáková and Dani [15], show that approximating OPT up to a factor less than 2 is already NP-hard for the $(1,\varepsilon)$-restricted problem, for any fixed $\varepsilon \leq 1/2$. Note that when $\varepsilon = 1$, so each resource has the same value, the problem can be solved exactly and easily via applications of a bipartite matching algorithm. This algorithm can also be used to give a $1/\varepsilon$-approximation, which is better than 2-approximation for $\varepsilon > 1/2$. In [17] it was proved that the integrality gap of the CLP for the $(1,\varepsilon)$-restricted allocation problem is at most 3, for every ε. The paper also gives a quasipolynomial-time algorithm that finds a $(3+4\varepsilon)$-approximation.

1.1 Our contributions

The existing algorithms and integrality gap estimation for the Santa Claus problem are, one way or another, based on the combinatorial augmenting tree argument of [26] for finding perfect matchings in certain hypergraphs. Many of them are sophisticated variants of the local search technique of [10] and its efficient algorithmic realization in [9].

Our main innovation in this work is to introduce the use of topological methods for the Santa Claus problem, and replace the combinatorial argument of [26]. This approach yields substantial improvements in the integrality gap of the CLP.

Our first main result improves the integrality gap from 3.808 to 3.534.

Theorem 1.1. *The integrality gap of the CLP is at most* $\frac{53}{15}$.

For our approach we make use of a criterion of Aharoni and Haxell [7] and Meshulam [36] for the existence of independent transversals in graphs, using the (topological) connectedness of the independence complex. In our application we apply this to an appropriately modified line graph of the multihypergraph of all those subsets that are valuable enough to be potentially allocated to the players. In order to show that the connectedness of the independence complex is large enough, we run a graph theoretic process, which is based on a theorem of Meshulam [36]. In the process we dismantle our line graph, but control the topological connectedness of the independence complex throughout, to make sure that the process runs for long enough. This necessitates that we choose our dismantling process with care and apply intricate analysis of the underlying structures, carefully tailored to the specifics of the problem. We employ the dual of the CLP to certify the length of the process.

Our approach is conceptually different from that of all previous work on the Santa Claus problem. The topological theorems in the background provide an incredibly rich family of independent sets in the modified line graph, that is geometrically highly structured via a triangulation of a high-dimensional simplex. This is in sharp contrast to the much simpler sparse spanning tree-like structure at the heart of the combinatorial approach.

This general strategy to show the existence of a solution of large minimum utility seems quite flexible and we expect it to be a useful asset for other algorithmic problems of interest involving hypergraph matchings.

Our machinery developed for the proof of Theorem 1.1 can also be used to improve significantly the known results on the integrality gap of the CLP for the $(1,\varepsilon)$-restricted allocation problem. In the next theorem we highlight some of the main consequences of this aspect of our work.

Theorem 1.2. *Let* $\varepsilon < \frac{1}{2}$ *and let* $\mathcal{I}$ *be an instance of the* $(1,\varepsilon)$*-restricted Santa Claus problem with maximum CLP-target* $T^* := T^*(\mathcal{I})$. *Then the integrality gap of* $\mathcal{I}$ *is at most* $f(\frac{\varepsilon}{T^*})$, *where* $f : (0,1] \to \mathbb{R}^+$ *is a function satisfying*

- $f(x) < 3$ *unless* $x = \frac{1}{6}$ *or* $x = \frac{1}{3}$,
- $f(x) \leq 2.75$ *for all* $x \in (0, \frac{1}{6}) \cup [\frac{2}{11}, \frac{1}{3}) \cup [\frac{4}{11}, 1]$, *and*
- $\lim_{x\to 0} f(x) < 2.479$.

One important message of this theorem is the identification of a couple of specific instances that seem especially hard to crack. For example, we would be delighted to see a $(1, 1/3)$ instance with an optimal CLP target of 1 and no allocation of min-value $2/3$. Furthermore, we see that as long as $\frac{\varepsilon}{T^*}$ is not too close to either of the two problematic values, the integrality gap is substantially below 3.

As observed in [17] (and also explained in the proof of Theorem 1.2), the assumption $1 \leq T^* < 2$ captures the challenging case of the problem. Under this assumption, the last part tells us that the integrality gap is less than 2.479 when $\varepsilon \to 0$. This estimate compares favorably with an instance of the problem given in [17], that has integrality gap 2 for arbitrarily small ε.

We remark that the restriction on ε in the theorem is not crucial since, as mentioned earlier, there is a simple $\frac{1}{\varepsilon}$-approximation algorithm based on bipartite matchings, which gives an approximation ratio ≤ 2 if $\varepsilon \geq \frac{1}{2}$. Moreover the restriction $x \leq 1$ is also natural as $T^* \geq \varepsilon$ whenever T^* is positive.

Finally, we note that our proofs in this paper can be turned into an algorithmic procedure that constructs an allocation with the promised min-value, but at the moment we have no control over the running time. Thus our results are in the same spirit as those of [23, 10, 17, 32, 19, 20] in which the strongest estimate on the integrality gap did not come with a corresponding efficient algorithm to find an allocation. Nevertheless, together with the machinery of [12], our work can be used to efficiently estimate the min-value of an optimal allocation. As an application of such a theorem we can imagine a scenario where Santa Claus might be prone to favoritism. Having supernatural powers, he can certainly calculate an optimal allocation, yet may choose a suboptimal one benefitting his favorites. Our Theorem 1.1 combined with [12] leads to a polynomial time algorithm that parents can use to uncover any bias Santa might have that is more blatant than $(\frac{15}{53} - \delta)$-times the optimum.

1.2 Related work

The max-min allocation problem is also widely studied in the more general case, where different players p might have different utility value v_{pr} for resource $r \in R$. The Santa Claus problem corresponds to the case when $v_{pr} \in \{0, v_r\}$. This scenario was first considered by Lipton, Markakis, Mossel, and Saberi [34]. The NP-hardness result of Bezáková and Dani [15] about approximating with a factor less than 2 is still the best known for the general case. Bansal and Sviridenko [12] showed that their CLP has an integrality gap of order $\Omega(\sqrt{|P|})$ for the general problem. Asadpour and Saberi [11] could match this with an $O(\sqrt{|P|}\log^3|P|)$-approximation algorithm using the CLP. Chakrabarty, Chuzhoy, and Khanna [16] give an $|R|^{\varepsilon}$-approximation algorithm for any constant ε, that works in polynomial time, as well as a $O(\log^{10}|R|)$-approximation algorithm that works in quasipolynomial time.

The special case where each resource is coveted by only two players is interesting algorithmically. In this case Bateni, Charikar, and Guruswami [13] showed that the Santa Claus problem is NP-hard to approximate to within a factor smaller than 2. Complementing this, Chakrabarty, Chuzhoy, and Khanna [16] give a 2-approximation algorithm, even if the values are unrestricted. The case when resources can be coveted only by three players is shown to be equivalent to the general case [13].

For the classical dual scenario of min-max allocation Lenstra, Shmoys, and Tardos [33] gave a 2-approximation algorithm and showed that it is NP-hard to approximate within a factor of 3/2. Using a configuration LP and a local search algorithm inspired by those developed for the Santa Claus problem, Svensson [39] managed to break the factor 2-barrier for the integrality gap of the restricted version of the min-max allocation problem. Once more, this result comes with an efficient algorithm to estimate the optimum value up to a factor arbitrarily close to $\frac{33}{17}$, but not with an efficient algorithm to find such an allocation. The approximation factor was subsequently improved to $\frac{11}{6}$ by Jansen and Rohwedder [30], who later [31] also provided an algorithm that finds such an allocation in quasipolynomial time.

Organization of the paper In Section 2 we present our topological tools and describe our proof strategy. In Section 3 we demonstrate how our method works by giving a clean proof of the fact that the integrality gap is at most 4. In Section 4 we introduce the two innovations that make our improvement on the integrality gap possible, and we use them in Section 5 to prove Theorem 1.1. In the subsequent Section 6 we give the proof of the two main statements from Section 4. Finally, in Section 7 we prove Theorem 1.2 on the two-values problem. Background and intuition for the topological notions we use are provided for the interested reader in the Appendix.

2 Topological tools and the proof strategy

2.1 The setup

Let $\mathcal{I} = (P, R, v, \{L_p : p \in P\})$ be an instance of the Santa Claus problem. A subset $e \subseteq L_p$ of coveted resources of some player $p \in P$ with the property that $v(e) \geq \alpha T$ and $v(e') < \alpha T$ for every proper subset $e' \subset e$ is called an *α-hyperedge*. We say that p is the *owner* of e or e is an α-hyperedge of p. To indicate this we might write e^p if necessary. Note that the hypergraph consisting of all α-hyperedges is a multihypergraph, since the same subset e may be an α-hyperedge of several players p. For example if an α-hyperedge $e \subseteq L_p \cap L_q$ with $p \neq q$, we will have both e^p and e^q in the multihypergraph. An allocation with min-value at least αT constitutes choosing for every player $p \in P$ an α-hyperedge of p, such that they are pairwise disjoint.

Let $T \in \mathbb{R}$ be a target such that CLP(T) is feasible. For $\alpha \in \mathbb{R}$, the *α-approximation allocation graph* $H(\mathcal{I}, T, \alpha) = H(\alpha)$ is the auxiliary $|P|$-partite graph with vertex set

$$V(H(\alpha)) = \cup_{p \in P} V_p, \text{ where } V_p = \{e^p : \ e \subseteq R \text{ is an } \alpha\text{-hyperedge of } p\},$$

and edge set

$$E(H(\alpha)) = \{e^p f^q : p \neq q, e \cap f \neq \emptyset\}.$$

An *independent transversal* in a vertex-partitioned graph such as $H(\alpha)$ is an independent set (i.e. one that induces no edges) that is a *transversal*, i.e. it consists of exactly one vertex in each partition class. Thus a problem instance $\mathcal{I}$ with feasible CLP(T) has an allocation with min-value at least αT for some $\alpha > 0$ if and only if the α-approximation allocation graph $H(\mathcal{I}, T, \alpha)$ has an independent transversal. Hence our Theorem 1.1 can be reformulated as follows.

Theorem 2.1. *Let $(P, R, \{L_p : p \in P\}, v)$ be an instance of the Santa Claus problem and let $T \in \mathbb{R}$ be such that the CLP(T) is feasible. Then the corresponding α-approximation allocation graph $H(\alpha)$ has an independent transversal with $\alpha = \frac{15}{53}$.*

2.2 Topological tools

In this section we introduce the main topological tools needed and describe how we use them in our arguments.

For a given graph G, let $\mathcal{J}(G) = \{I \subseteq V(G) : I \text{ is independent}\}$ be its *independence complex*. Following Aharoni and Berger [2] we define $\eta(G)$ to be the (topological) connectedness of $\mathcal{J}(G)$ plus 2. An advantage of this shifting by 2 is that the formulas for the following simple properties of η simplify (see e.g. [1, 2, 6]). (In fact Part (2) is true in much greater generality, see e.g. [2], but this simple statement is all we require.)

Fact 1. *Let G be a graph.*

(1) $\eta(G) \geq 0$ with equality if and only if G is the empty graph (i.e. the graph with no vertices).

(2) If graph G is the disjoint union of G_1 and a non-empty graph G_2 then $\eta(G) \geq \eta(G_1) + 1$. Moreover, if G_2 is a single (isolated) vertex then $\eta(G) = \infty$.

Intuitively, $\eta(G)$ represents the smallest dimension of a "hole" in the geometric realization of the abstract simplicial complex $\mathcal{J}(G)$. For the purposes of this paper it suffices to regard η strictly as a graph parameter satisfying Fact 1 and the upcoming Theorems 2.2 and 2.3. However, for the interested reader we provide the formal definition, background and some intuition in the Appendix.

Our proof of Theorem 2.1 is based on two key theorems involving the parameter η. The first one provides a sufficient Hall-type condition for the existence of independent transversals. This result was implicit already in [7] and [35], and was first stated explicitly in this form in [36] (see also [2]). Let I be an index set and J be an $|I|$-partite graph with vertex partition $V_1, \ldots, V_{|I|}$. For a subset $U \subseteq I$ we denote by $J|_U$ the induced subgraph $J[\cup_{i\in U} V_i]$ of J defined on the vertex set $\cup_{i\in U} V_i$.

Theorem 2.2. *Let I be an index set and J be an $|I|$-partite graph with vertex partition $V_1, \ldots, V_{|I|}$. If for every subset $U \subset I$ we have $\eta\left(J|_U\right) \geq |U|$, then there is an independent transversal in J.*

The formal resemblance of Theorem 2.2 to Hall's Theorem for matchings in bipartite graphs is no coincidence: the latter is a consequence of the former. Indeed, for a bipartite graph $B = (X \cup Y, E)$ satisfying Hall's Condition we can define an $|X|$-partite (simple) graph $J(B)$, where for every $x \in X$ there is a part $V_x = \{y^x : y \in N_B(x)\}$ and $y_1^{x_1} y_2^{x_2}$ is an edge if and only if $y_1 = y_2$. Then a matching of B saturating X corresponds to an independent transversal in $J(B)$. For a subset $U \subseteq X$, the subgraph $J(B)|_U$ is the union of $|N(U)|$ disjoint cliques, so $\eta(J(B)|_U) \geq |N(U)| \geq |U|$ by Properties (1) and (2) in Fact 1 and Hall's condition.

Our second tool is a theorem of Meshulam [36], reformulated in a way that is particularly well-suited for our arguments. Let G be a graph, and let e be an edge of G. We denote by $G - e$ the graph obtained from G by deleting the edge e (but not its end vertices). We denote by $G \divideontimes e$ the graph obtained from G by removing both endpoints of e and all of their neighbors. The graph $G \divideontimes e$ is called G with e *exploded*.

Theorem 2.3. *Let G be a graph and let $e \in E(G)$, such that $\eta(G - e) > \eta(G)$. Then we have that $\eta(G) \geq \eta(G \divideontimes e) + 1$.*

Inspired by Meshulam's Theorem we call an edge e of G *deletable* if $\eta(G - e) \leq \eta(G)$ and *explodable* if $\eta(G \divideontimes e) \leq \eta(G) - 1$. By the theorem, if an edge is not deletable then it is explodable. A *deletion/explosion sequence, or DE-sequence, starting with graph G_{start}* is a sequence of operations, which, starting with G_{start}, in each step either deletes a deletable edge or explodes an explodable edge in the current graph. The *length* $\ell(\sigma)$ of the sequence σ is the number of explosions in σ. The following are simple yet crucial properties of DE-sequences.

Observation 2.4. *Let G be the outcome of a DE-sequence σ of length ℓ, starting with G_{start}. Then the following are true.*

(i) If G has an edge then σ can be extended by a further deletion or explosion.

(ii) $\eta(G_{start}) \geq \eta(G) + \ell$.

(iii) If G has an isolated vertex then $\eta(G_{start}) = \infty$.

We call a DE-sequence σ satisfying the property of part (iii) a *KO-sequence*. Note the obvious: appending a KO-sequence after any DE-sequence results in a KO-sequence. In Appendix 9.2 we give a small concrete example demonstrating how to use DE-sequences to obtain a lower bound on η.

Proof. By Meshulam's Theorem in any graph G every edge is either deletable or explodable. Hence (i) follows. Part (ii) follows since during performing the DE-sequence σ the deletion of a deletable edge does not increase the value of η and the explosion of an explodable edge decreases the value of η by at least 1. For (iii), by (ii) and by Fact 1(2) we have $\eta(G_{start}) \geq \eta(G) = \infty$. □

2.3 The proof strategy

Let T be such that $CLP(T)$ of instance $\mathcal{I}$ with the target T has a feasible solution. Our proof strategy is to take, for our chosen α, the graph $H(\alpha)$ defined in the introduction and use Theorem 2.2 to derive the existence of an independent transversal in it.

Those α-hyperedges that contain a single resource will have a special status. A resource $r \in R$ is called *fat* if $v_r \geq \alpha T$, otherwise it is called *thin*. The set $F = F(\alpha) := \{r \in R : v_r \geq \alpha T\}$ is the set of fat resources. Any set $S \subseteq R$ of resources with $S \cap F = \emptyset$ is called thin. We will in particular be speaking of *thin α-hyperedges* and *thin configurations*. Note that an α-hyperedge is thin if and only if it contains at least two elements. The corresponding vertices of $H(\alpha)$ are also called thin. For a fat resource $r \in R$, the singleton $\{r\}$ is called a *fat α-hyperedge*, and if $r \in L_p$ then r^p is called a *fat vertex* of $H(\alpha)$. Each fat resource $r \in F$ corresponds to a clique $C_r := \{r^p : r \in L_p\}$ in $H(\alpha)$ which forms a component, since no other α-hyperedge contains r (due to their minimality).

As we show next, we can shift our main focus to the subgraph $J(\alpha) := H(\alpha) - \cup_{r\in F} C_r$ of $H(\alpha)$ induced by the set of thin vertices. To verify the condition of Theorem 2.2 we need to consider an arbitrary subset $U \subseteq P$ of the players and the corresponding induced subgraph $H(\alpha)|_U$ of $H(\alpha)$. By Fact 1(2) the disjoint clique components corresponding to fat vertices $r \in F_U := F \cap (\cup_{p\in U} L_p)$ each contribute at least one to the value of $\eta(H(\alpha)|_U)$. We thus need to prove that for the remaining graph we have $\eta(J(\alpha)|_U) \geq |U| - |F_U|$.

To that end, starting with $G_{start} = J(\alpha)|_U$ we will specify a DE-sequence σ and prove that either σ is a KO-sequence or $\ell(\sigma) \geq |U| - |F_U|$. In the former case Observation 2.4(iii) implies $\eta(J(\alpha)|_U) = \infty$. In the latter case, denoting by G_{end} the final graph of σ, Observation 2.4(ii) and Fact 1(1) imply $\eta(J(\alpha)|_U) \geq \eta(G_{end}) + |U| - |F_U| \geq |U| - |F_U|$. In both cases we have that

$$\eta(H(\alpha)|_U) \geq \eta(J(\alpha)|_U) + |F_U| \geq |U|,$$

so the condition of Theorem 2.2 is verified. Hence there exists an independent transversal in $H(\alpha)|_U$ and we are done. We have just proved the following.

Theorem 2.5. *Let $\mathcal{I} = (P, R, v, \{L_p : p \in P\})$ be a problem instance and $T \in \mathbb{R}$ such that CLP(T) has a feasible solution. Suppose for every $U \subseteq P$ there exists a DE-sequence σ starting with $G_{start} = J(\alpha)|_U$ such that either σ is a KO-sequence, or $\ell(\sigma) \geq |U| - |F_U|$. Then $H(\alpha)$ has an independent transversal.*

We remark that this approach to proving the existence of an independent transversal using η was described in terms of a game in [6], and used in many settings, see e.g. [4, 5, 8, 27, 28, 29].

With Theorem 2.5 we have reduced our task to constructing, for every $U \subseteq P$, a DE-sequence σ starting with $G_{start} = J(\alpha)|_U$ such that either σ is a KO-sequence, or $\ell(\sigma) \geq |U| - |F_U|$. To prove lower bounds on the length of a DE-sequence σ that starts with $J(\alpha)|_U$, we will maintain a cover $W \subseteq R$ of all α-hyperedges that correspond to vertices of $J(\alpha)|_U$, that disappeared during explosions of σ, and control the size of W. If we are able to do this, then the complement of W is large, allowing us to find an α-hyperedge in it and hence extend the DE-sequence further. Note that deletions do not remove any vertices of $J(\alpha)|_U$.

More generally, we say W is a *cover of the DE-sequence* σ starting with a subgraph $G_{start} \subseteq J(\alpha)|_U$ and ending with G_{end} if

($\star$) every vertex e^p of G_{start} with $e \cap W = \emptyset$ is present in G_{end}.

The natural choice to cover the α-hyperedges corresponding to vertices that disappeared from $G_{start} \subseteq J(\alpha)|_U$ during the explosions in a DE-sequence σ is $\bigcup(e \cup f)$, where the union is over all edges $e^p f^q$ of G_{start} exploded in σ. This will be called the *basic cover* of σ. Note that for the basic cover W_σ, every vertex h^s of G_{start} with $h \cap W_\sigma = \emptyset$ is unaffected by each explosion that happened during σ and hence is still present in the graph G_{end}.

In the next subsection we will demonstrate how the simple accounting by adding up the values of the basic covers of the explosions of an arbitrary DE-sequence starting with $J(\alpha)|_U$ and ending with a graph with no edges is already sufficient to derive the existence of an allocation of min-value at least $\frac{1}{4}T$. To achieve our improved bounds in Theorem 2.1, in Sections 4 and 5 we will choose our DE-sequences and account for their accompanying covers more carefully.

3 The demonstration of the method

In this section, as a warm-up, we set $\alpha = \frac{1}{4}$ and construct for every $U \subseteq P$, a DE-sequence σ starting with $G_{start} = J(\alpha)|_U$ such that either σ is a KO-sequence, or $\ell(\sigma) \geq |U| - |F_U|$. Then by Theorem 2.5 we have proved Theorem 2.1 with $\frac{1}{4}$ in place of $\frac{15}{53}$.

Fix $U \subseteq P$. We choose an *arbitrary* DE-sequence σ starting with $G_{start} = J(\alpha)|_U$ and ending with a graph G_{end} with no edges. This is possible because of Observation 2.4(i). If G_{end} contains a vertex then σ is a KO-sequence and we are done, so we may assume that G_{end} has no vertices. We are left to show that $\ell(\sigma) \geq |U| - |F_U|$.

To estimate the value of covers the following definition will be useful. Let $m = m(\alpha) := \max\{v(s) : s \subseteq L_p \text{ for some } p \in P \text{ and } v(s) < \alpha T\}$ be the maximum value of a "non-α-hyperedge". Note that in particular $m < \alpha T$ and the value of every configuration S is greater than m/α. A subset $s \subseteq R$ is called a *block* if $v(s) \leq m$. Note that any proper subset of an α-hyperedge is a block.

We estimate the value of the basic cover W_σ by simply adding up estimates for the basic covers of its individual explosions.

Observation 3.1. *For the explosion of edge $e^p f^q$ the value of its basic cover $e \cup f$ is at most $3m$.*

Proof. The cover $e \cup f$ is a subset of the union of three blocks: $(e \setminus \{x\}) \cup \{x\} \cup (f \setminus e)$, where $x \in e$ is arbitrary. Indeed, $f \setminus e$ is a block since it is a proper subset of f, and both $e - x$ and $\{x\}$ are blocks since they are proper subsets of e. For this recall that all α-hyperedges under consideration are thin. Consequently $v(e \cup f) \leq 3m$. $\square$

Hence $v(W_\sigma) \leq 3m\ell(\sigma)$.

To give a lower bound on this value we invoke the dual DCLP(T) of the configuration LP. In DCLP(T) there is a variable $y_p \geq 0$ for each player $p \in P$, a variable $z_r \geq 0$ for each resource $r \in R$, and for each configuration $S \in \mathcal{C}_p(T)$ there is a constraint $y_p \leq \sum_{r \in S} z_r$. The objective function $\sum_{p \in P} y_p - \sum_{r \in R} z_r$ has to be maximized. The following proposition, that, with foresight, is stated here under general conditions, provides feasible solutions to the dual LP and a lower bound on the value of W_σ.

Proposition 3.2. *Let $U \subseteq P$ be arbitrary and let $Y \subseteq R \setminus F$ such that $v(Y \cap S) \geq c$ for every thin configuration $S \in \mathcal{C}_p(T)$ for $p \in U$. Then*

$$y_p = \begin{cases} 0 & p \notin U \\ c & p \in U \end{cases} \qquad z_r = \begin{cases} c & r \in F_U \\ v_r & r \in Y \\ 0 & \textit{otherwise} \end{cases}$$

is a feasible solution of DCLP(T). In particular

$$v(Y) \geq c(|U| - |F_U|).$$

Proof. To check the feasibility of the solution, let $S \in \mathcal{C}_p(T)$ be an arbitrary configuration. If $p \notin U$, then $y_p = 0$ and the corresponding constraint holds by the non-negativity of the z_r.

If $p \in U$, then $y_p = c$. If there is a fat resource $s \in S$, then $s \in F_U$, so $\sum_{r \in S} z_r \geq z_s = c = y_p$. Otherwise $S \cap F = \emptyset$ and hence $\sum_{r \in S} v_r \geq \sum_{r \in S \cap Y} v_r \geq c = y_p$. So the solution is feasible.

That means that the value $\sum_{p \in P} y_p - \sum_{r \in R} z_r = |U|c - |F_U|c - \sum_{r \in Y} v_r$ is a lower bound on the value 0 of the primal CLP(T) and the second claim follows. □

To obtain a lower bound on $v(W_\sigma)$ we apply Proposition 3.2 with U, $Y = W_\sigma$ and $c = 3m$. To that end we need to check $v(S \cap W_\sigma) \geq 3m$ for every thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$. Note that since $\alpha = 1/4$, the value of every configuration S is $v(S) > 4m$. Hence it is enough to verify that $v(S \setminus W_\sigma) \leq m$. Since G_{end} has no vertices, Property $(\star)$ of W_σ implies that $R \setminus F \setminus W_\sigma$ should contain no α-hyperedge of any $p \in U$. Consequently, for any thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$, the value of $S \setminus W_\sigma$ should not be large enough for an α-hyperedge. Hence $v(S \setminus W_\sigma) \leq m$ as needed.

Proposition 3.2 then implies $v(W_\sigma) \geq 3m(|U| - |F_U|)$. Combining this with $v(W_\sigma) \leq 3m\ell(\sigma)$, we obtain $\ell(\sigma) \geq |U| - |F_U|$ and we are done by Theorem 2.5.

4 Economical DE-sequences

In this section we start our proof of Theorem 2.1 by establishing a couple of important tools. Our improvement on Section 3 relies on finding DE-sequences whose accounting (through their covers) is done more economically when some of the explosions are packed together. We use two different approaches, one based on total value (treated in Section 4.1) and the other based on total cardinality (Section 4.2).

For our setup in this section we let $\mathcal{I} = (P, R, v, \{L_p : p \in P\})$ be a problem instance and $T \in \mathbb{R}$ be such that CLP(T) has a feasible solution. We fix $\alpha, 0 < \alpha < 1$, and subset $U \subseteq P$. We let $G^* \subseteq J(\alpha)|_U$ be a subgraph of the graph of interest and $W \subseteq R \setminus F$ be a subset of resources such that $(\star)$ holds with $G_{start} = J(\alpha)|_U$ and $G_{end} = G^*$.

4.1 Cheap DE-sequences

We say that a DE-sequence σ is *cheap* if there exists a cover of σ of value at most $2m\ell(\sigma)$. Note that any sequence of deletions is a cheap DE-sequence, hence the following holds by Meshulam's Theorem.

Observation 4.1. *Suppose there is no cheap DE-sequence starting with graph G^*. Then every edge of G^* is explodable.*

In practice we often demonstrate that a DE-sequence σ is cheap by exhibiting a cover that is a subset of the union of at most $2t$ blocks. Recall that the value of any block is at most m and hence any proper subset of an α-hyperedge is a block.

We will extensively use the following simple fact about KO-sequences.

Observation 4.2. *Suppose there is no KO-sequence starting with graph G^*. Let G be the current graph in some DE-sequence σ starting with G^*. For any α-hyperedge $e^p \in V(G)$, there exists a (possibly empty) sequence of deletions, after which e^p can be exploded (with a neighbor in G).*

Proof. Note that appending a sequence of deletions to the end of a DE-sequence starting with G^* would form a KO-sequence starting with G^* if these deletions would make e^p an isolated vertex in G. Recall also that by Meshulam's Theorem for every edge of G it is possible to perform either a deletion or an explosion. □

Before stating our results we introduce a few conventions. For an element x and set e we write $e - x$ and $e + x$ as shorthand for $e \setminus \{x\}$ and $e \cup \{x\}$, respectively. We use the term *α-edge* for an α-hyperedge with exactly two elements.

Our main theorem of this section states that if there is no cheap DE-sequence and no KO-sequence starting with G^*, then configurations with substantial value outside W can only contain α-edges (i.e., α-hyperedges of size 2) outside W.

Theorem 4.3. *Suppose there is no KO-sequence and no cheap DE-sequence starting with $G^* = (V, E)$. Let $C \in \mathcal{C}_p(T)$ be a configuration of player $p \in U$, such that $C \setminus W$ contains an α-hyperedge e with $|e| \geq 3$. Then $C \setminus (W + s)$ contains no α-hyperedge, where s is the second-most valuable element of e (ties broken arbitrarily). In particular $v(C \setminus W) \leq 3m/2$.*

4.2 Cardinality based economical DE-sequences

In this subsection we describe short DE-sequences that have an economical cover in terms of cardinality.

A DE-sequence σ is called a *γ-DE-sequence* if it has a cover of cardinality at most $\gamma\ell(\sigma)$. In our proofs γ will be either 7/3 or 5/2.

Lemma 4.4. *Let $j = 2$ or 3. Suppose there is no KO-sequence and no cheap DE-sequence starting with G^*. If there is a thin configuration $C \in \mathcal{C}_p(T)$ with $p \in U$ and $v(C \cap W) < T - jm$, then there exists a $(2j + 1)/j$-DE-sequence starting with G^*.*

The proofs of both Theorem 4.3 and Lemma 4.4 are quite intricate and postponed to Section 6.

5 Proof of Theorem 2.1

In this section we prove Theorem 2.1.

Proof. Recall that as we saw in Section 2 it is sufficient to verify the conditions of Theorem 2.5 with $\alpha = 15/53$. That is, for any fixed subset $U \subseteq P$ of the players we will prove that there exists a DE-sequence σ starting with $G_{start} = J(\alpha)|_U$ such that either σ is a KO-sequence, or $\ell(\sigma) \geq |U| - |F_U|$.

We define σ in four phases. Here G denotes the current graph of the sequence.

- Phase 1. WHILE a KO-sequence or cheap DE-sequence τ exists in G, DO perform τ
- Phase 2. WHILE there exists a 7/3-DE-sequence τ in G, DO perform τ, and then iteratively perform any KO-sequence or cheap DE-sequence until no further one exists.
- Phase 3. WHILE there exists a 5/2-DE-sequence τ in G, DO perform τ, and then iteratively perform any KO-sequence or cheap DE-sequence until no further cheap DE-sequence exists.
- Phase 4. WHILE G has an edge DO perform a deletion or an explosion in G.

When the procedure terminates, the final graph G_{end} has no edge. If G_{end} contains a vertex, then σ is a KO-sequence starting with G_{start}, as desired. Therefore we may assume that G_{end} has no vertices. Note that in this case we did not perform a KO-sequence at any point during our procedure.

Let n_1 denote the total number of explosions performed in the cheap DE-sequences throughout Phases 1, 2, and 3, and W_1 be the union of all the covers associated to these cheap DE-sequences. By definition of cheap DE-sequence we know that

$$v(W_1) \leq 2mn_1. \tag{5.1}$$

For $j = 2, 3$, let n_j denote the number of explosions performed in 7/3-DE-sequences during Phase 2 and 5/2-DE-sequences during Phase 3, respectively, and let W_j be the union of their corresponding covers. For a $(2j+1)/j$-DE-sequence the number of resources in the cover is $2j+1$ and the number of explosions is j. Hence

$$|W_3| \leq \frac{7}{3}n_2 \text{ and } |W_2| \leq \frac{5}{2}n_3. \tag{5.2}$$

For these sets it will also be useful to estimate their values. For this, recall that each thin resource is of value at most m. Therefore

$$v(W_3) \leq \frac{7}{3}mn_2 \text{ and } v(W_2) \leq \frac{5}{2}mn_3. \tag{5.3}$$

Let n_4 be the number of explosions performed in Phase 4 and W_4 the union of the basic covers corresponding to them. Then by Observation 3.1 we have

$$v(W_4) \leq 3mn_4. \tag{5.4}$$

We will use the dual DCLP(T) to take snapshots at various points during σ in order to derive lower bounds involving linear combinations of the quantities n_j, $j = 1, 2, 3, 4$. A couple of times we will find our estimates using a more refined version of Proposition 3.2.

Proposition 5.1. *Let $U \subseteq P$ be arbitrary. Let $0 \leq c \leq 2d$, $Y \subseteq R \setminus F$, such that if $S \in \mathcal{C}_p(T)$ is a thin configuration owned by $p \in U$ with $|Y_{>d} \cap S| \leq 1$ then*

$$v(Y_{\leq d} \cap S) \geq \begin{cases} c & \text{if } Y_{>d} \cap S = \emptyset \\ c - d & \text{if } |Y_{>d} \cap S| = 1. \end{cases}$$

Then

$$y_p = \begin{cases} 0 & p \notin U \\ c & p \in U \end{cases} \qquad z_r = \begin{cases} c & r \in F_U \\ d & r \in Y_{>d} \\ v_r & r \in Y_{\leq d} \\ 0 & \textit{otherwise} \end{cases}$$

is a feasible solution of DCLP(T) and

$$c|U| - c|F_U| \leq d|Y_{>d}| + v(Y_{\leq d}).$$

Moreover for any partition of $Y = Y_1 \cup Y_2$ we have

$$c|U| - c|F_U| \leq d|Y_1| + v(Y_2).$$

Note that we can recover the final conclusion of Proposition 3.2 by setting $d = c$ and using that $d|Y_{>d}| \le v(Y_{>d})$.

Proof. First we check feasibility of the given solution. Let $S \in \mathcal{C}_p(T)$ be an arbitrary configuration. If $p \notin U$, then $y_p = 0$ and the corresponding constraint holds by the non-negativity of the z_r.

Otherwise $p \in U$ and we must check $\sum_{r \in R} v_r \ge y_p = c$.

If there is a fat resource $s \in F \cap S \subset F_U$, then $\sum_{r \in S} z_r \ge z_s = c = y_p$. Otherwise $S \cap F = \emptyset$ and we make a case distinction based on $|Y_{>d} \cap S|$. If $Y_{>d} \cap S = \emptyset$, then $\sum_{r \in S} z_r \ge \sum_{r \in S \cap Y_{\le d}} z_r = v(S \cap Y_{\le d}) \ge c$.

If $Y_{>d} \cap S = \{s\}$, then $\sum_{r \in S} z_r \ge \sum_{r \in S \cap Y_{>d}} z_r + \sum_{r \in S \cap Y_{\le d}} z_r = z_s + v(S \cap Y_{\le d}) \ge d + c - d = c$.

Finally, if $|Y_{>d} \cap S| \ge 2$, then $\sum_{r \in S} z_r \ge \sum_{r \in S \cap Y_{>d}} z_r \ge 2d \ge c$.

So in all cases the solution is feasible. That means that the value $\sum_{p \in P} y_p - \sum_{r \in R} z_r = c|U| - c|F_U| - d||Y_{>d}| - v(Y_{\le d})$ is a lower bound on the value 0 of the primal LP and our second conclusion follows. To derive the final conclusion note that

$$\begin{aligned} d|Y_{>d}| + v(Y_{\le d}) &= d|(Y_1)_{>d}| + v((Y_1)_{\le d}) + d|(Y_2)_{>d}| + v((Y_2)_{\le d}) \\ &\le v((Y_1)_{>d}) + v((Y_1)_{\le d}) + d|(Y_2)_{>d}| + d|(Y_2)_{\le d}| \\ &= v(Y_1) + d|Y_2|. \end{aligned}$$

□

We derive our first inequality after Phase 2 is complete.

Lemma 5.2.

$$|U| - |F_U| \le \frac{2m}{T - 3m} n_1 + \frac{7}{3} n_2. \tag{5.5}$$

Proof. Let G be the current graph after the end of Phase 2 and set $W = W_1' \cup W_2$, where W_1' is the union of the covers associated with cheap DE-sequences up to the end of Phase 2. We use Proposition 5.1 with U, $c = d = T - 3m$, and $Y = W$. For this we only need to check for every thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$ and $W_{>T-3m} \cap S = \emptyset$ that $v(W_{\le T-3m} \cap S) = v(W \cap S) \ge T - 3m$. This follows from Lemma 4.4 applied with W, $j = 3$, $G^* = G$, and $C = S$, since after Phase 2 is complete there is no KO-sequence, cheap DE-sequence, or 7/3-DE-sequence starting with G. By Proposition 5.1 we conclude that

$$(T - 3m)(|U| - |F_U|) \le v(W_1') + (T - 3m)|W_2|.$$

By (5.1) we have $v(W_1') \le v(W_1) \le 2mn_1$ and by (5.2) we have $|W_2| \le \frac{7}{3} n_2$. This completes the proof. □

In our next lemma we take a snapshot after Phase 3 and derive two inequalities.

Lemma 5.3.

$$|U| - |F_U| \le \frac{m}{T - 3m} n_1 + \frac{7}{6} n_2 + \frac{5}{4} n_3 \tag{5.6}$$

$$|U| - |F_U| \le \frac{2m}{T - 2m} n_1 + \frac{7m}{3(T - 2m)} n_2 + \frac{5m}{2(T - 2m)} n_3. \tag{5.7}$$

Proof. Let G be the current graph at the end of Phase 3 and set $W = W_1 \cup W_2 \cup W_3$ to be the corresponding cover. As there is no KO-sequence, cheap DE-sequence, or 5/2-DE-sequence at the end of Phase 3, Lemma 4.4, applied with W, $j = 2$, $G^* = G$, and $C = S$, implies $v(W \cap S) \geq T - 2m$ for any thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$.

For the first inequality we use Proposition 5.1 with U, $c = 2(T - 3m), d = T - 3m$, and $Y = W$. We need to check for every thin $S \in \mathcal{C}_p(T)$ with $p \in U$ and $|W_{>T-3m} \cap S| \leq 1$ that its value is large enough. If $W_{>T-3m} \cap S = \emptyset$ then $v(W_{\leq T-3m} \cap S) = v(W \cap S) \geq T - 2m \geq 2(T - 3m)$, since $4m > T$ (otherwise we are done by Section 3). If $W_{>T-3m} \cap S = \{s\}$ then $v(W_{\leq T-3m} \cap S) = v(W \cap S) - v_s \geq T - 2m - m$, since s is thin.

By Proposition 5.1 we conclude that

$$2(T - 3m)(|U| - |F_U|) \leq v(W_1) + (T - 3m)|W_2 \cup W_3|.$$

Using (5.1) and (5.2), we are done.

For the second inequality we use Proposition 3.2 with U, $c = T-2m$, and $Y = W$. Recall that by the application of Lemma 4.4, we already know that $v(W \cap S) \geq T-2m$ for every thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$. Hence by Proposition 3.2 we conclude that $(T - 2m)(|U| - |F_U|) \leq v(W) \leq v(W_1) + v(W_2) + v(W_3)$. Using (5.1) and (5.3), the inequality follows. □

Finally, after Phase 4, we also measure the covers.

Lemma 5.4.

$$|U| - |F_U| \leq \frac{2m}{T - m} n_1 + \frac{7m}{3(T - m)} n_2 + \frac{5m}{2(T - m)} n_3 + \frac{3m}{T - m} n_4. \tag{5.8}$$

Proof. Set $W = W_1 \cup W_2 \cup W_3 \cup W_4$ and apply Proposition 3.2 with U, $c = T - m$, and $Y = W$. This is possible as after Phase 4 is complete, there are no vertices left in the final subgraph G_{end} of $J(\alpha)|_U$. Consequently the value of resources in $S \setminus W$ is at most m for any thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$. That means $v(W \cap S) \geq T - m$ holds.

By Proposition 3.2 we then find that $(T - m)(|U| - |F_U|) \leq v(W) \leq v(W_1) + v(W_2) + v(W_3) + v(W_4)$. Using (5.1), (5.3) and (5.4), the inequality follows. □

Now it is easy to check that if $T \geq \frac{53}{15}m$, then the convex combination of the inequalities (5.5), (5.6), (5.7), and (5.8) with coefficients $\frac{1}{35}$, $\frac{26}{245}$, $\frac{46}{2205}$, and $\frac{38}{45}$, respectively, imply that

$$|U| - |F_U| \leq n_1 + n_2 + n_3 + n_4 = \ell(\sigma).$$

This completes the proof of Theorem 2.1. □

6 Proofs of the existence of economical DE-sequences

In what follows, when our attention is focused on the multihypergraph of α-hyperedges, we often refer to a vertex e^p of $J(\alpha)|_U$ as α-hyperedge e of p or owned by p. When the identity of the owner is irrelevant or already established, we often omit the reference to the owner. In particular we also sometimes say the pair of α-hyperedges e and f are explodable, without specifying their owners. We refer to α-hyperedges e of p and f of q as *decoupled* in the current graph G of our DE-sequence if $e^p f^q$ is not an edge of G. We also say that an α-hyperedge g of p *survives* an explosion if the vertex g^p is still present in the current graph G after the explosion. We say that α-hyperedges e and f are *explodable at resource* r if $e \cap f = \{r\}$ and the pair e and f are explodable.

Lemma 6.1. *Suppose there is no KO-sequence and no cheap DE-sequence starting with $G^* = (V, E)$. Let $f^q \in V$ be an α-hyperedge of player q. Then the following hold.*

(i) If for player $p \in U$ there exists an α-hyperedge $g^p \in V$ such that $g^p f^q \in E$ then for every α-hyperedge $e^p \in V$ of p we have $|f \cap e| \leq 1$.
In particular, for any $e^p \in V$ with $|e \cap f| \geq 2$, we have $e^p f^q \notin E$.
Even more in particular, if $e^p, e^q \in V$, then $e^p e^q \notin E$.

(ii) For every resource $r \in f$ there exists an α-hyperedge g in V that is explodable with f at r.

Proof. Recall that since there is no cheap DE-sequence starting with G^*, every edge of G^* is explodable.

To prove (i), first we show that if $e^p \in V$ with $|e \cap f| \geq 2$ then $e^p f^q \notin E$. Indeed, otherwise $e^p f^q$ is explodable and $(e - x) \sqcup ((f \setminus e) + x))$, where $x \in e \cap f$, is a partition of the basic cover $e \cup f$ into two blocks. This demonstrates the explosion of the edge $e^p f^q$ is cheap, a contradiction. Here we use that $e - x$ is a proper subset of the α-hyperedge e since $x \in e$, and $(f \setminus e) + x$ is a proper subset of the α-hyperedge f, since $|e \cap f| \geq 2$.

Let now $g^p \in V$ be an α-hyperedge of p that is explodable with f^q. Suppose on the contrary that there exists an α-hyperedge e^p of p such that $|e \cap f| \geq 2$. By the above $e^p f^q \notin E$ and $e \neq g$. We define a DE-sequence of length two (starting with G^*) and show that it is cheap, giving a contradiction. First explode f^q with g^p. Note that e^p survives, since f^q is decoupled from e^p and g^p and e^p are owned by the same player. Then, since there is no KO-sequence isolating e^p, after possibly some deletions, we can explode e^p with some neighbor h. (See Observation 4.2.) Then the basic cover $f \cup g \cup e \cup h$ of this DE-sequence of length two has a partition $(e - x) \sqcup ((f \setminus e) + x) \sqcup (g \setminus f) \sqcup (h \setminus e)$, where $x \in e \cap f$, into four blocks. Hence this DE-sequence is cheap, as claimed. This verifies the main statement of (i), which in turn implies the second statement.

For the last statement of (i) note that in our setting V contains only thin hyperedges, so $|e| \geq 2$.

For (ii) note that since there is no KO-sequence starting with G^*, f has some neighbor in G^*. Once again, every edge of G^* is explodable. For a contradiction assume that every α-hyperedge g with $f \cap g = \{r\}$ is not explodable with f. Explode f with an arbitrary neighbor h. By (i) we know that $f \cap h = \{s\}$ for some s. By our assumption $s \neq r$. The key observation here is that the set $W^* = (f \cup h) - r$, i.e. something less than the basic cover, is also a cover of the explosion of the edge fh. For this, let g be an α-hyperedge of G^* that did not survive the explosion of the edge fh. If g is a neighbor of h in G^* then $W^* \cap g \supseteq h \cap g \neq \emptyset$, since $r \notin h$. Otherwise g is a neighbor of f and is covered by W^* unless $g \cap f = \{r\}$. However our starting assumption was that f had no such neighbor in G^*. Then the cover $W^* = (f - r) \cup (h \setminus f)$ of the single explosion of the edge fh is the union of two blocks, which makes it cheap, a contradiction. □

We are now ready to prove Theorem 4.3.

Proof of Theorem 4.3. To begin we observe that from our setup it follows that if a subset $S \subseteq C \setminus W$ has value $v(S) > m$ then S contains some α-hyperedge h of p and by property $(\star)$ of W we have $h^p \in V$. Throughout this proof, when talking about α-hyperedges contained in $C \setminus W$, we mean those owned by p, unless otherwise specified.

To derive the last statement from the main conclusion note first that the second-most valuable element s of e satisfies $v_s \leq m/2$, otherwise the value of the two most valuable elements of e would exceed m, contradicting that $|e| \geq 3$. Then $v(C \setminus W) = v(C \setminus (W + s)) + v_s \leq m + m/2$.

Suppose on the contrary that $C \setminus (W + s)$ contains an α-hyperedge. We claim that some such α-hyperedge g intersects $e - s$ nontrivially. Let $f \subseteq C \setminus (W + s)$ be an α-hyperedge and suppose it is disjoint from $e - s$. We know f contains at least two elements, so removing the least valuable element r from f results in $v(f-r) \geq v(f)/2 > m/2$. Similarly, since s is not the unique most valuable element of e we find that $v(e-s) \geq v(e)/2 > m/2$. Therefore $(f-r)\cup(e-s) \subseteq C\setminus(W+s)$ is valuable enough to contain an α-hyperedge g of p, and since $f-r$ is not valuable enough to contain it, we have $g\cap(e-s) \neq \emptyset$, as claimed.

Recall that since $g \subseteq C\setminus W$, we know $g^p \in V$. Let a be the most valuable element in $g\cap(e-s)$.

Case 1. $|g\cap(e-s)| \geq 2$. Let $b \neq a$ such that $b \in g\cap(e-s)$. Then $v_a \geq v_b$, so $v_s \geq v_b$ and therefore $v((g\cup e)-b) \geq v(g-b+s) \geq v(g) > m$. Hence $(g\cup e)-b$ contains an α-hyperedge h of p and $h^p \in V$.

We achieve a contradiction by defining a DE-sequence of length two starting with G^*, which turns out to be cheap. First explode g^p at b with some α-hyperedge d_1 in G^*, which exists by Lemma 6.1(ii). By part (i) $|d_1\cap e| \leq 1$ and hence $d_1\cap e = \{b\}$. Consequently h^p survives the explosion of the edge gd_1, since $h\cap d_1 \subseteq ((g\cup e)-b)\cap d_1 = \emptyset$ and h has the same owner as g. Since there is no KO-sequence isolating h^p starting with G^*, after possible deletions now we can explode h^p with some α-hyperedge d_2. We claim that the basic cover of this DE-sequence of length two (starting with G^*) is a subset of the union of four blocks:

$$d_1\cup g\cup h\cup d_2 \subseteq (d_1\setminus g)\cup(e-b)\cup(g-a)\cup(d_2\setminus h).$$

Here we used that $(h\setminus g)+a$ is contained in $e-b$. This contradiction completes Case 1.

Case 2. $g\cap(e-s) = \{a\}$.

Case 2.a. $v((g\cup e)-a) > m$.

Let us choose an α-hyperedge $f \subseteq (g\cup e)-a$ of p, as follows. If there is a resource $b \in e-a$ such that $v(g-a+b) > m$, choose $f \subseteq g-a+b \subseteq (g\cup e)-a$, and otherwise choose one arbitrarily.

By Lemma 6.1(ii) g^p has an explodable neighbour d at a in G^*. Since g and e are both α-hyperedges of p, by Lemma 6.1(i) we have that $|d\cap e| \leq 1$ and consequently $d\cap e = \{a\}$.

We achieve a contradiction by defining a DE-sequence of length two starting with G^*, which turns out to be cheap. First we explode g^p with d. The α-hyperedge f^p survives this explosion since $f\cap d \subseteq ((g\cup e)-a)\cap d = \emptyset$ and f and g are both α-hyperedges of p. Secondly, (after some possible deletions) we explode f^p with some neighbour h. This is possible since there is no KO-sequence isolating f^p.

We claim that the basic cover $d\cup g\cup h\cup f$ of this DE-sequence of length two is the subset of the union of four blocks, providing the contradiction we seek. There is a slight difference in the accounting depending how f was chosen.

If b is such that $v(g-a+b) > m$ and $f \subseteq g-a+b$, we take the blocks $(d\setminus g)\cup(h\setminus f))\cup(g-a)\cup\{a,b\}$. Note that $\{a,b\}$ is a block since it is a proper subset of the α-hyperedge e with at least three elements.

Otherwise $f \subseteq (g\cup e)-a$ and $g-a+b$ is a block for every $b \in e-a$. Then we take the blocks $(d\setminus g)\cup(h\setminus f))\cup(g-a+b)\cup(e-b)$, where $b \in e-a$ is arbitrary. Here note that the union of the third and the fourth term is $g\cup e$, which in turn contains f. This completes Case 2.a.

Case 2.b $v((g\cup e)-a) \leq m$.

Let x denote the most valuable element in $g-a$ (here $\{x\} = g-a$ is possible). By Lemma 6.1(ii) there is an α-hyperedge h, that is explodable with g^p at x in G^*.

We claim that $h\cap(e-a) \neq \emptyset$. Suppose not. We define a cheap DE-sequence starting with G^*, giving a contradiction. We start by exploding g^p with h. The α-hyperedge e^p survives this explosion, since $h\cap e = \emptyset$ (note that $a \notin h$ since $a \in g-x = g\setminus h$) and g and e are both α-hyperedges of p. Again, since there is no KO-sequence isolating e^p, after possibly some deletions, we explode

e^p with some neighbour d. The basic cover $g \cup h \cup e \cup d$ of this DE-sequence is the subset of the union of four blocks: $(h \setminus g) \cup \{a\} \cup ((g \cup e) - a) \cup (d \setminus e)$.

This contradiction establishes $h \cap (e - a) \neq \emptyset$.

Since e is an α-hyperedge of the owner of g^p, by Lemma 6.1(i) we have $|h \cap e| \leq 1$, and hence $h \cap e = \{b\}$ for some $b \neq a$. To complete the proof, we will argue that the explosion of the edge hg is cheap by establishing that the basic cover $h \cup g$ can be partitioned into two blocks: $(h - b) \cup (g - x + b)$. The rest of the proof is concerned with demonstrating that $g - x + b$ is indeed a block (the first term is clearly a block).

If $g - a = \{x\}$ is a singleton, then $g - x + b = \{a, b\}$ which is a proper subset of the α-hyperedge e of size at least three and hence is a block.

Otherwise fix a resource $y \in g - a - x$, and suppose on the contrary that $v(g - x + b) > m$. We will find an α-hyperedge d of p with $h \cap d \supseteq \{b, x\}$, which would contradict Lemma 6.1(i) since h is explodable with an α-hyperedge of p, namely g.

Since $y \in g$, we have $v_x \geq v_y$, and hence for $X := g - y + b$ we have $v(X) \geq v(g - x + b) > m$. Note however, that both $X - a$ and $X - b$ are blocks, since $v(X - a) \leq v(g \cup e) - a) \leq m$, and $X - b = g - y$ is a proper subset of the α-hyperedge g. Hence $a, b \in f$ for any α-hyperedge $f \subseteq X$ of p.

If $x \in f$ as well, then we are done. Otherwise let us fix an α-hyperedge $f = \{a, b, u_1, \ldots, u_t\} \subseteq X$ and modify it slightly to obtain the appropriate d.

Note that $t \geq 1$ since $\{a, b\}$ is a proper subset of e. Then $f' = f - u_1 + x = \{a, b, x, u_2, \ldots, u_t\}$ contains an α-hyperedge d of p because $v(x) \geq v(u_1)$ since $u_1 \in f - a - b \subseteq g - y - a$. Note that since $d \subseteq f - u_1 + x \subseteq X + x$ and $x \in g - a \subseteq X$ we find $d \subseteq X$, and so $\{a, b\} \subset d$ (since $X - a$ and $X - b$ are blocks). Furthermore d must also contain x, since otherwise $d \subseteq f' - x = \{a, b, u_2, \ldots, u_t\}$ is a block (as a proper subset of f). This completes the proof of Case 2.b, and that of the theorem. □

To end this section we give the proof of Lemma 4.4.

Proof of Lemma 4.4. Since $v(C \setminus W) = v(C) - v(C \cap W) > T - (T - jm) = jm$ there are strictly more than j resources in $C \setminus W$; let $s, t_1, t_2, \ldots, t_j$ be the $j + 1$ most valuable ones, in increasing order of value. Then $v(C \setminus (W \cup \{t_2, \ldots, t_j\})) > m$ and consequently $C \setminus (W \cup \{t_2, \ldots, t_j\})$ must contain an α-hyperedge.

Since $v(C \setminus W) > jm > \frac{3}{2}m$, by Theorem 4.3 every α-hyperedge in $C \setminus W$ has cardinality two. So in particular the two most valuable elements s and t_1 of $C \setminus (W \cup \{t_2, \ldots, t_j\})$ form an α-edge of p and $(st_1)^p \in V$. Then, since $v_{t_1} \leq v_{t_i}$ for $i \in \{2, \ldots, j\}$, we also know that each st_i is an α-edge of p and $(st_i)^p \in V$.

Next we derive a couple of crucial observations about the explodable neighbors of st_i at t_i.

Claim. *For every $i = 1, \ldots, j$, if an α-hyperedge f^q (of some player $q \neq p$), with $t_i \in f_i$, is explodable with $(st_i)^p$ in G^* then*

- *f is an α-edge and*
- *$f \cap (C \setminus W) = \{t_i\}$.*

Proof. If f had at least 3 resources then the explosion would be cheap, contradicting our assumption. To see this choose $a \in f - t_i$, and observe that the basic cover $\{s, t_i\} \cup f$ can be partitioned into two blocks $\{a, t_i\}$ and $f \setminus \{a, t_i\} + s$. Indeed, $\{a, t_i\}$ is a proper subset of f and hence is a block, and using $v_s \leq v_{t_i}$ we see that $v(f \setminus \{a, t_i\} + s) \leq v(f \setminus \{a, t_i\} + t_i) = v(f - a) \leq m$. Thus every such f is an α-edge.

If $f \subseteq C \setminus W$ then by our setup ($\star$) implies that $f^p \in V$. This contradicts Lemma 6.1(i) since f^q explodable with an α-hyperedge owned by p (that is st_i) and $|f \cap f| \geq 2$. So $f \setminus (C \setminus W) \neq \emptyset$ and the Claim follows. □

We now define a DE-sequence τ of length j starting with G^* that has basic cover W_τ of size $(2j+1)$.

We construct τ by finding explosions one by one. Let τ_0 be the empty DE-sequence. Suppose for some $i, 1 \leq i \leq j$ we have already found DE-sequence τ_{i-1} which performs explosions of the edges $(st_1)f_1, \ldots, (st_{i-1})f_{i-1}$, in this order, such that for every $k = 1, \ldots i-1$ we have

(a) f_k is an α-edge and

(b) $f_k \cap (C \setminus W) = \{t_k\}$.

Note that by (a) and (b) for the basic cover $W_{\tau_{i-1}} = \{s\} \cup \cup_{k=1}^{i-1} f_k$ we have $|W_{\tau_{i-1}}| = 2i-1$ and $v(W_{\tau_{i-1}}) \leq (2i-1)m$ (since every element is thin).

Our first step in constructing τ_i from τ_{i-1} is to perform, iteratively, all possible deletions, so any remaining edge in the current graph G is explodable. Note that $(st_i)^p$ survived all the explosions of τ_{i-1} since it is disjoint from $f_1, \ldots, f_{i-1}$ by (a) and has the same owner as each other $(st_j)^p$. To complete the definition of τ_i, our aim is to find an explodable neighbor f_i of $(st_i)^p$ at t_i in the current graph G, and explode it.

Note we cannot apply Lemma 6.1(ii) here directly to show the existence of such an f_i, since after executing τ_{i-1}, we do not know any more that there is no cheap DE-sequence starting with the current graph G. So we suppose that $(st_i)^p$ has no explodable neighbour at t_i in the current graph G and obtain a contradiction. Since there is no KO-sequence starting with G^* that isolates $(st_i)^p$ there exists an α-hyperedge e with $s \in e$, that is explodable with $(st_i)^p$. If we now perform this explosion then we claim that the resulting DE-sequence τ' starting with G^* would be cheap, providing a contradiction. Indeed, the basic cover of τ_{i-1} together with e forms a cover of τ'. (Recall that no neighbor of $(st_i)^p$ in the current graph gets destroyed in the last explosion just because it contained t_i.) The value $v(W_{\tau_{i-1}}) + v(e-s) \leq (2(i-1)+1)m + m = 2im$ shows that τ' is cheap. So for some player $q \neq p$ there exists an α-hyperedge $f_i^q \in V(G)$ which is explodable with $(st_i)^p$ at t_i in G.

Now we show that f_i satisfies properties (a) and (b). To see this observe that f_i^q was also explodable with $(st_i)^p$ in G^*. Otherwise, by Meshulam's Theorem, the edge $f_i^q(st_i)^p$ was deletable in G^*. But the existence of a deletable edge in G^* contradicts our assumption that G^* has no cheap DE-sequence. Hence our Claim applies to f_i and so (a) and (b) hold for $k = i$.

For the basic cover of the ultimate DE-sequence τ_j we have $|W_{\tau_j}| = 2j+1$, showing that τ_j is a $(2j+1)/j$-DE-sequence. □

7 Two values

Here we consider the $(1,\varepsilon)$-restricted version of the Santa Claus problem, in which $v_r \in \{1, \varepsilon\}$ for each $r \in R$, where ε is a constant satisfying $0 < \varepsilon < 1$. Our overall approach in this section conceptually parallels our work in the earlier sections. We will derive Theorem 1.2 from the following technical result.

Theorem 7.1. *Let $\varepsilon < \frac{1}{2}$. Let $\mathcal{I}$ be an instance of the $(1,\varepsilon)$-restricted Santa Claus problem and let $T \in \mathbb{R}$ be such that CLP(T) has a feasible solution. Suppose that $1 \leq T < 2$, and that*

$c := \lceil T/\varepsilon \rceil \geq 4$. *Fix a positive integer* r*. For each integer* $X \geq r$ *define*

$$a(X) = a_r(X) := \begin{cases} 3r - X - 1 & r \leq X \leq \frac{3r-1}{2} \\ 2r - \frac{X+1}{3} & \frac{3r}{2} \leq X \leq 2r \\ \frac{4r-1}{3} & 2r + 1 \leq X. \end{cases}$$

Suppose $r \geq 2$ *is an integer such that* $\sum_{X=r}^{c} \frac{1}{a(X)} \geq 1$*. Then for* $\alpha = \frac{r\varepsilon}{T}$*, the corresponding* α*-approximation allocation graph* $H(\alpha)$ *has an independent transversal. In particular, there is an allocation with min-value at least* $r\varepsilon$*.*

For each integer $c \geq 4$, we will apply this theorem with the largest integer $r \in \mathbb{N}$, denoted by r_c, satisfying $\sum_{X=r}^{c} \frac{1}{a(X)} \geq 1$. For convenience, we provide a table showing the triples $(c, r_c, c/r_c)$ for $1 \leq c \leq 30$ (with c/r_c truncated to two decimal places).

c	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
r_c	1	1	1	2	2	2	3	3	4	4	4	5	5	6	6
c/r_c	1	2	3	2	2.5	3	2.33	2.66	2.25	2.5	2.75	2.4	2.6	2.33	2.5
c	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
r_c	6	7	7	8	8	8	9	9	10	10	11	11	11	12	12
c/r_c	2.66	2.42	2.57	2.37	2.5	2.62	2.44	2.55	2.4	2.5	2.36	2.45	2.54	2.4	2.5

The idea of the proof of Theorem 7.1 is to use the same basic framework as that of Theorem 2.1. As before, we will define a DE-sequence by concatenating many shorter DE-sequences constructed in phases, in which each phase lasts as long as there remains a configuration C with value $v(C \setminus W)$ exceeding a certain threshold associated with that phase. Note however that in the $(1, \varepsilon)$-restricted setting we can measure value entirely in terms of cardinality. These DE-sequences are described in Section 7.1 and parallel our arguments in Section 4. We will then define a dual solution based on each phase and obtain from it a lower bound on the total length of the DE-sequences constructed during that phase. The result will be an overall lower bound on the length of the whole DE-sequence, which we then optimize to derive the conclusion of Theorem 7.1. This part of the proof corresponds to Section 5 and is done in Section 7.2.

In the following subsections we will need to refer to some simple properties of the pairs (c, r_c). These are spelled out in the following observation.

Observation 7.2. *(i)* $r_c \geq \frac{c}{4}$ *for every* $c \geq 4$*,*

(ii) $c \geq 2r_c + 1$ *for every* $c \geq 5$*,*

(iii) $c \geq 2r_c + 2$ *for every* $c \geq 10$*.*

Proof. The sum $A = \sum_{X=r}^{c} \frac{1}{a(X)}$ has $c - r + 1$ terms, that form a non-decreasing sequence. The smallest term is $\frac{1}{2r-1}$, implying that if $c - r + 1 \geq 2r - 1$ then $A \geq 1$. This is easily satisfied by $r = \lceil \frac{c}{4} \rceil$, implying (i).

The largest term in A is (at most) $\frac{3}{4r-1}$. Note that $\frac{3}{4r-1} < \frac{1}{r+1}$ when $r \geq 5$, so to reach 1 in this case there must be at least $c - r + 1 > r + 1$ terms. Hence $c \geq 2r_c + 1$ if $r \geq 5$. For values of $r \leq 4$ see the table to complete the proof of (ii).

Similarly $\frac{3}{4r-1} < \frac{1}{r+2}$ when $r \geq 8$, so to reach 1 in this case there must be at least $c-r+1 > r+2$ terms. Hence $c \geq 2r + 2$ if $r \geq 8$, and again the table completes the proof of (iii) for the smaller values. □

7.1 DE-sequences

For our setup in this subsection we fix $\varepsilon \in (0, 1/2)$, and let $\mathcal{I} = (P, R, v, \{L_p : p \in P\})$ be an instance of the $(1, \varepsilon)$-restricted problem, where $T \in [1, 2)$ is such that CLP(T) has a feasible solution. Set $c = \lceil T/\varepsilon \rceil$. By assumption $c \geq 4$. Let $r \geq 2$ be such that $\sum_{X=r}^{c} \frac{1}{a(X)} \geq 1$. For $r = 2$ this certainly satisfied (although in practice we will want to choose r as large as possible to get the strongest result). Set $\alpha = r\varepsilon/T$. Then in particular $\varepsilon \leq \frac{\alpha T}{2}$, in other words the resources of value ε are thin.

We also claim that resources of value 1 are fat, i.e. that $1 \geq \alpha T$. When $c = 4$ and $r = 2$, the assumption $\varepsilon < 1/2$ implies $\alpha T = r\varepsilon < 1$. Otherwise, by Observation 7.2(ii) we know $c \geq 2r_c + 1$, so $\frac{2}{\varepsilon} + 1 > \frac{T}{\varepsilon} + 1 > c \geq 2r + 1$, implying $r < \frac{1}{\varepsilon}$ as claimed.

Hence we conclude that c is a lower bound on the size of every thin configuration, and $r = \lceil \alpha T/\varepsilon \rceil$ is the common size of every thin α-hyperedge.

As in the proof of Theorem 2.1, we fix a subset $U \subseteq P$. We let $G^* \subseteq J(\alpha)|_U$ be a subgraph of the thin part of the α-approximation allocation graph $H(\alpha)$ and $W \subseteq R \setminus F$ be a subset of resources such that $(\star)$ holds with $G_{start} = J(\alpha)|_U$ and $G_{end} = G^*$. We will also assume that there is no KO-sequence starting with G^* and there is no deletion possible, so all edges of G^* are explodable.

We say that a DE-sequence σ starting with G^* is *based in* a configuration C if one α-hyperedge out of every pair of α-hyperedges exploded during σ is in $C \setminus W$ (and is owned by the owner of C). We write $av(\sigma) = |W_\sigma|/\ell(\sigma)$ for the *average cost* of the cover W_σ of σ.

In our first lemma we describe two criteria that guarantee that a DE-sequence based in C can be continued.

Lemma 7.3. *Let $C \in \mathcal{C}_p(T)$ be a thin configuration with $p \in U$ and let σ be a DE-sequence starting with G^* based in C, with explosions of $e_1f_1, \ldots, e_\ell f_\ell$, in this order, where $e_i \subseteq C \setminus W$, $1 \leq i \leq \ell$, is owned by p.*

(a) If $e \subseteq (C \setminus W) \setminus \left(\bigcup_{i=1}^{\ell} f_i\right)$ is of size $|e| = r$, then σ can be extended to a longer DE-sequence based in C with an explosion involving e^p. In particular, if σ is maximal, then $av(\sigma) \leq r + \frac{r-1}{\ell(\sigma)}$.

(b) Let G be the current graph immediately before the ℓth explosion and suppose G has only explodable edges. If $e_\ell f_\ell \in E(G)$ was chosen such that $e_\ell \subseteq \left((C \setminus W) \setminus \bigcup_{j=1}^{\ell-1} f_j\right)$ and $|e_\ell \cap f_\ell|$ is maximized with this property, and if $f_\ell \cap e_\ell \neq f_\ell \cap \left((C \setminus W) \setminus \bigcup_{j=1}^{\ell-1} f_j\right)$, then σ can be extended to a longer DE-sequence based in C.

Proof. For part (a), note that being disjoint from all f_js, the α-hyperedge e^p survived all explosions so far. Since there is no KO-sequence starting with G^* isolating e^p, after possible deletions, there will be an explodable edge incident to e^p in the current graph. So we can extend σ with a further explosion based in C involving e^p. Therefore if σ is maximal then $\left|\left((C \setminus W) \setminus \bigcup_{j=1}^{\ell} f_j\right)\right| \leq r - 1$ and for the size of the basic cover we have

$$|W_\sigma| \leq \left|\left(\bigcup_{i=1}^{\ell} f_i\right) \cup (C \setminus W)\right| \leq r\ell(\sigma) + r - 1.$$

For part (b) let us take $y \in (f_\ell \setminus e_\ell) \cap \left((C \setminus W) \setminus \bigcup_{j=1}^{\ell-1} f_j\right) \neq \emptyset$ and $w \in e_\ell \setminus f_\ell \neq \emptyset$. Then the α-hyperedge $g = e_\ell - w + y$ is contained in $\left((C \setminus W) \setminus \bigcup_{j=1}^{\ell-1} f_j\right)$, and consequently by part (a) g^p in particular is *present* in the current graph G immediately before the explosion of $e_\ell f_\ell$. Since $|f_\ell \cap g| > |f_\ell \cap e_\ell|$, the α-hyperedge g^p was not explodable with f_ℓ immediately before the explosion of $e_\ell f_\ell$, when G had only explodable edges. Hence $g^p f_\ell$ was not an edge of G. Therefore g^p survives the ℓth explosion as well and since there is no KO-sequence isolating g^p, another explosion involving g^p is possible. $\square$

The main lemma of this section provides DE-sequences, based in C, with low average cost.

Lemma 7.4. *For every thin configuration $C \in \mathcal{C}_p(T)$ with $p \in U$ and $X := |C \setminus W| \geq r$ there exists a DE-sequence σ starting with G^* based in C with $av(\sigma) \leq a(X)$.*

Proof. First we assume that $X \leq \frac{3r-1}{2}$. If there is a DE-sequence of length two based in C, then by Lemma 7.3(a) it has average cost at most $r + \frac{r-1}{2} \leq 3r - X - 1$.

We may thus assume that there is no DE-sequence of length two based in C. Among all explodable pairs $ef \in E(G^*)$ with $e \subseteq C \setminus W$, let us choose one with $|e \cap f|$ largest. If $f \cap e \neq f \cap (C \setminus W)$, then by Lemma 7.3(b) there is a second explosion based in C, a contradiction. For this recall that G^* has only explodable edges.

Otherwise $f \cap e = f \cap (C \setminus W)$. Since no further explosion based in C is possible, $|(C \setminus W) \setminus f)| \leq r - 1$ by Lemma 7.3(a). In other words $|f \cap (C \setminus W)| \geq X - r + 1$, in which case for the size of the basic cover $f \cup e$ we have $|e| + |f| - |f \cap (C \setminus W)| \leq 2r - (X - r + 1) = 3r - X - 1$, as desired.

We consider now the range $\frac{3r}{2} \leq X$. We distinguish two cases and in each of them we either construct a DE-sequence based in C with average cost at most $2r - \frac{X+1}{3}$ or one of length at least 3, which has average cost at most $\frac{4r-1}{3}$ by Lemma 7.3(a). This shows that $av(\sigma) \leq \max\{2r - \frac{X+1}{3}, \frac{4r-1}{3}\}$. The bound on $av(\sigma)$ in the second and third ranges then follows since $2r - \frac{X+1}{3} \geq \frac{4r-1}{3}$ if and only if $2r \geq X$.

For the purposes of this proof we set $b = 2r - \frac{X+1}{3}$.

Case 1. Suppose first that some explodable neighbor f is such that $|(C \setminus W) \cap f| \leq \lceil \frac{2X-1}{3} \rceil - r = 3r - 1 - \lfloor 2b \rfloor =: t_1$ (note that this expression is non-negative for X in our range).
We perform this explosion and then deletions until no more are possible. Then at least $X - t_1 \geq r$ resources remain in $(C \setminus W) \setminus f$, so by Lemma 7.3(a) there is a further explosion possible. Among all edges hg in the current graph, we choose one with $g \subseteq (C \setminus W) \setminus f$ and $|h \cap g|$ largest. Note that hg was an edge of G^* as well, and hence we can assume $|h \cap g| \leq 2r - 1 - \lfloor b \rfloor$. Indeed, otherwise the single explosion of hg at the beginning would have constituted a DE-sequence with average cost $|h \cup g| \leq |h| + |g| - |h \cap g| \leq \lfloor b \rfloor \leq b$. Now for our second explosion we explode hg.

If $h \cap g = h \cap ((C \setminus W) \setminus f)$, then we still have $|(C \setminus W) \setminus f| - |h \cap g| \geq X - t_1 - (2r - 1 - \lfloor b \rfloor) \geq X - 5r + 2 + \lfloor b \rfloor + \lfloor 2b \rfloor \geq X - 5r + 2 + b + 2b - 1 = X - 5r + 1 + 6r - (X + 1) = r$ resources in $(C \setminus W) \setminus (f \cup h)$. Hence by Lemma 7.3(a) our sequence can be extended to a sequence of length 3.

Otherwise $h \cap g \neq h \cap (C \setminus W) \setminus f)$, in which case Lemma 7.3(b) guarantees the extension of our sequence.

Case 2. Suppose now that $|(C \setminus W) \cap f| \geq t_1 + 1$ for every explodable neighbor f of an α-hyperedge in $C \setminus W$ owned by p.
For our first explosion we choose edge hg of G^* with $g \subseteq C \setminus W$ such that $|g \cap h|$ is largest. Unless our first explosion is already a DE-sequence of the type we are looking for, we have that its basic cover $h \cup g$ satisfies $|h \cup g| \geq \lfloor b \rfloor + 1$.

If $h \cap g = h \cap (C \setminus W)$ then the set $(C \setminus W) \setminus h$ contains $X - |h \cap (C \setminus W)| = X - (|h| + g| - |h \cup g|) \geq X - (2r - (\lfloor b \rfloor + 1)) = X - \lceil \frac{X+1}{3} \rceil + 1$ resources, which is at least r for $X \geq \frac{3r}{2}$. Consequently we can apply Lemma 7.3(a) to extend our DE-sequence with a second explosion.

If $h \cap g \neq h \cap (C \setminus W)$, then by Lemma 7.3(b) we can also extend our DE-sequence with a second explosion.

Either way we have a DE-sequence σ of length two, based in C, with explosions gh and $g'h'$. We claim that the set $W_\sigma = h \cup h' \cup (C \setminus W) \setminus Y$, where $Y \subseteq (C \setminus W) \setminus (h \cup h')$ is an *arbitrary* subset of size $\min\{t_1, |(C \setminus W) \setminus (h \cup h')|\}$, is a cover. Indeed, let f be an α-hyperedge disjoint from W_σ, and let us show that f survived the explosions of σ. Since f is in particular disjoint from h and h', the only way f could have disappeared is if it had an edge of the current graph, and hence also of G^*, to g or g'. But then f is owned by some $q \neq p$ and the inequality in Case 2 applies to f, so we know that $|f \cap (C \setminus W)| \geq t_1 + 1$. Since $f \cap W_\sigma = \emptyset$ we have that $f \cap (C \setminus W) \subseteq Y$, which is a contradiction since Y is too small for that.

The DE-sequence σ is based in C and has length two, so unless its average cost is already small enough for our lemma, we have that $2b < |W_\sigma| = |h| + |h'| + |((C \setminus W) \setminus (h \cup h')) \setminus Y|$, which implies

$$|(C \setminus W) \setminus (h \cup h')| \geq \lfloor 2b \rfloor + 1 - 2r + |Y|.$$

We claim that this implies $|(C \setminus W) \setminus (h \cup h')| \geq r$ and therefore by Lemma 7.3(a) we can extend σ to a third explosion. Indeed, if $|Y| = t_1$, then this is immediate from $t_1 = 3r - 1 - \lfloor 2b \rfloor$. Otherwise we have $2r - 1 \geq \lfloor 2b \rfloor = 4r - \lceil \frac{2X+2}{3} \rceil$, or equivalently $\lceil \frac{2X+2}{3} \rceil \geq 2r + 1$, which implies $X \geq 3r$. But then $|(C \setminus W) \setminus (h \cup h')| \geq |C \setminus W| - |h| - |h'| = X - 2r \geq r$ as needed. □

7.2 Proof of Theorem 7.1

Recall the assumptions of Theorem 7.1: let $\varepsilon < \frac{1}{2}$ and let $\mathcal{I} = (P, R, v, \{L_p : p \in P\})$ be an instance of the $(1, \varepsilon)$-restricted Santa Claus problem. Let $T \in \mathbb{R}$, $1 \leq T < 2$ be such that CLP(T) has a feasible solution and $c = \lceil T/\varepsilon \rceil \geq 4$. Suppose that for an integer $r \geq 2$ we have $\sum_{X=r}^{c} \frac{1}{a(X)} \geq 1$.

As in Theorem 2.1, we will apply Theorem 2.5 to infer the existence of the independent transversal in $H(\alpha)$. To that end we fix a subset $U \subseteq P$, assume there is no KO-sequence starting with $J(\alpha)|_U$, and seek a DE-sequence τ of length at least $|U| - |F_U|$ starting with $J(\alpha)|_U$. Our strategy will be as follows.

> INITIALIZATION. Let τ be a sequence of deletions starting with $J(\alpha)|_U$ until no further deletion is possible and let G^* be the resulting subgraph. Let $W = \emptyset$. For each X, $c \geq X \geq r$, in decreasing order, execute the following Phase X;
>
> PHASE X: WHILE there is a configuration C with at least X resources remaining in $C \setminus W$, DO perform a DE-sequence σ starting with G^* based in C (as given by Lemma 7.4 corresponding to the value of X) and perform all possible deletions afterwards. Update G^* to be the resulting current graph. Append σ to the end of τ. Let W_σ denote the cover of σ and set $W := W \cup W_\sigma$.

Next we verify that this process is well-defined, that is, whenever Lemma 7.4 is called upon in some Phase X, the graph G^* and cover W satisfy the three conditions of the setup for Section 7.1. These conditions are that W satisfies $(\star)$ with $G_{start} = J(\alpha)|_U$ and $G_{end} = G^*$, that there is no KO-sequence starting with G^*, and that all edges of G^* are explodable. Observe that each iteration of Phase X is immediately preceded by an iteration of Phase X, of Phase $X+1$, or the initialization phase.

The last property holds since we end the initialization phase and each iteration of Phase X or $X + 1$ by performing deletions until no more were possible. Consequently all edges of G^* are explodable. The next to last property holds throughout since appending a KO-sequence to the end of the current τ would form a KO-sequence starting with $J(\alpha)|_U$, contrary to our assumption. Finally, the $(\star)$ property is maintained after each execution of an iteration of Phase X or $X + 1$, as the cover W_σ of the new segment of τ is added to W, and it holds trivially after initialization, since no explosions have yet occurred so $W = \emptyset$ is a cover.

Let W_X be the union of the covers and n_X be the number of explosions done in Phase X. By Lemma 7.4, for each X with $c \geq X \geq r$ we have $|W_X| \leq n_X a(X)$.

For $c \geq X \geq r$ we consider the moment after the last step in Phase X is executed and set $W = W_c \cup \ldots \cup W_X$. For each thin configuration $S \in \mathcal{C}_p(T)$ with $p \in U$, we know that $|S \cap W| \geq c - X + 1$, otherwise we could have continued with another step of Phase X. Recalling that ε is the common value of all thin resources, we conclude that $v(S \cap W) \geq \varepsilon(c - X + 1)$ for each such S. Hence we may apply Proposition 3.2 with W in place of Y and $\varepsilon(c - X + 1)$ in place of c to obtain

$$\varepsilon|W| = v(W) \geq \varepsilon(c - X + 1)(|U| - |F_U|).$$

Comparing the upper and lower bounds on $| \cup_{j=X}^c W_j|$ for each $X = c, c-1, \ldots, r$, we obtain

$$\sum_{j=X}^{c} a(j) n_j \geq (c - X + 1)(|U| - |F_U|).$$

Since the coefficient function a is non-increasing in j, in order to minimize the objective function $\sum_{j=r}^c n_j$, we have to choose $n_c, n_{c-1}, \ldots, n_r$ in reverse order such that all the inequalities are equalities

$$\sum_{j=X}^{c} a(j) n_j = (c - X + 1)(|U| - |F_U|).$$

This implies that the length $\ell(\tau) = \sum_{j=r}^c n_j$ of the DE-sequence τ our process creates is minimized when $n_j = \frac{|U|-|F_U|}{a(j)}$. Consequently $\ell(\tau) \geq \sum_{j=r}^c \frac{1}{a(j)}(|U| - |F_U|)$, which is at least $|U| - |F_U|$, as required for Theorem 2.5. This completes the proof of Theorem 7.1.

7.3 Proof of Theorem 1.2

Proof of Theorem 1.2. We define f by

$$f(x) = \frac{1}{x r_{\lceil \frac{1}{x} \rceil}}.$$

Here, recall, that for an integer $c \in \mathbb{N}$, we denote by r_c the largest integer $r \in \mathbb{N}$ such that $\sum_{X=r}^c \frac{1}{a(X)} \geq 1$. We show that $f(x)$ bounds the integrality gap for $\mathcal{I}$, where $x = \frac{\varepsilon}{T^*}$.

First observe that for any instance with $T^* > 0$, it is easy to check Hall's condition on the bipartite graph of players and coveted resources to demonstrate that there is a valid allocation of one resource to each player. Hence, in the two-values case $OPT \geq \varepsilon$. Since $T^* \geq OPT$, we always have that the integrality gap is at most $\frac{T^*}{\varepsilon} \geq 1$. This shows that $f(x) = \frac{1}{x}$ is an appropriate choice to bound the integrality gap for every $x \in (0, 1]$. Since $r_1 = r_2 = r_3 = 1$, this verifies the statement when $x \geq \frac{1}{3}$.

We proceed now with the case $x := \frac{\varepsilon}{T^*} < \frac{1}{3}$.

Analogously to [17] we start by reducing the problem to the case when $1 \leq T^* < 2$. If $T^* \geq 2$ recall that the additive approximation result of Bezáková and Dani [15], mentioned in the Introduction, gives a polynomial-time algorithm to find an allocation with min-value $T_{ALP} - \max v_r$, where T_{ALP} is the optimum of the standard assignment LP. Hence $OPT \geq T_{ALP} - 1 \geq T^* - 1 \geq \frac{T^*}{2}$ as the CLP is stronger than the ALP and in our case $\max v_r = 1$. So the integrality gap is at most 2, which is less than $f(x)$ for every x.

If $T^* < 1$, then we create another instance $\mathcal{I}'$ where for each resource $r \in R$ with $v_r = \varepsilon$, we change the value of r to $v'_r = \varepsilon' := \frac{\varepsilon}{T^*}$, and keep $v'_r = v_r = 1$ otherwise. It is easy check that $T^*(\mathcal{I}') = 1$ and $OPT(\mathcal{I}) \geq OPT(\mathcal{I}')T^*$. Note that $\varepsilon' < \frac{1}{3}$. Hence, applying our theorem to the $(1, \varepsilon')$-instance $\mathcal{I}'$, there is an allocation with min-value at least $\frac{T^*(\mathcal{I}')}{f(\varepsilon'/T^*(\mathcal{I}'))}$. The very same allocation in $\mathcal{I}$ has min-value at least $\frac{T^*}{f(\varepsilon')}$ and hence exhibits an integrality gap of at most $f(\varepsilon/T^*)$ for $\mathcal{I}$.

From now on we assume $1 \leq T^* < 2$ and apply Theorem 7.1. For this we note that $c := \lceil \frac{T^*}{\varepsilon} \rceil = \lceil \frac{1}{x} \rceil > 3$ and $\varepsilon < \frac{1}{2}$. Theorem 7.1 then implies that we have an allocation for $\mathcal{I}$ with min-value at least $r_c\varepsilon$, thus $OPT \geq r_c\varepsilon$. Hence the integrality gap for $\mathcal{I}$ is at most

$$\frac{T^*}{r_c\varepsilon} = \frac{1}{x r_{\lceil \frac{1}{x} \rceil}} = f(x).$$

as promised.

To prove the assertions of Theorem 1.2 about the values, first note that when $x \geq \frac{1}{3}$ we have $f(x) = \frac{1}{x}$, which is less than 3 unless $x = \frac{1}{3}$, and at most 2.75 for $x \geq \frac{4}{11}$. For $x < \frac{1}{3}$ note that with $c = \lceil \frac{1}{x} \rceil$, we have $f(x) = \frac{1}{x r_c} \leq \frac{c}{r_c}$. It is easy to verify that for every $c \geq 4$ we have $c/r_c \leq 2.75$, unless $c = 6$. (In fact the ratio 2.75 is attained on the unique pair $c = 11, r = 4$.) That is, unless $\lceil \frac{1}{x} \rceil = 6$ we have $f(x) \leq 2.75$. When $\lceil \frac{1}{x} \rceil = 6$ we have $x \geq \frac{1}{6}$ and $f(x) = \frac{1}{r_6 x} = \frac{1}{2x}$, which is at most 2.75 for $x \geq \frac{2}{11}$ and strictly less than 3 unless $x = \frac{1}{6}$.

Finally, we deal with the case in which $x \to 0$. Then $c = \lceil \frac{1}{x} \rceil \to \infty$, and hence also $r_c \geq \frac{c}{4} \to \infty$ by Observation 7.2(i).

Let us assume that $x < \frac{1}{10}$, so $c \geq 10$ implying that $c \geq 2r_c + 2$ by Observation 7.2(iii). Setting $r := r_c$, we write the sum as $\sum_{X=r}^{c} \frac{1}{a(X)} = A_r + B_r + C_r$, where

$$A_r := \sum_{X=r}^{\lfloor (3r-1)/2 \rfloor} \frac{1}{3r - X - 1} = \sum_{k=\lceil (3r-1)/2 \rceil}^{2r-1} \frac{1}{k} = (H_{2r-1} - H_{\lceil (3r-1)/2 \rceil - 1}) \to \ln(4/3),$$

$$B_r := \sum_{\lceil 3r/2 \rceil}^{2r} \frac{3}{6r - X - 1} = \sum_{k=4r-1}^{\lfloor \frac{9r}{2} - 1 \rfloor} \frac{3}{k} = 3(H_{\lfloor \frac{9r}{2} - 1 \rfloor} - H_{4r-2}) \to 3\ln(9/8),$$

$$C_r := \sum_{X=2r+1}^{c} \frac{3}{4r-1} = \frac{3(c-2r)}{4r-1},$$

when $r \to \infty$. Here we use the well-known fact for the harmonic series $H_n = \sum_{k=1}^{n} \frac{1}{k}$, that $H_n - \ln n$ converges to a constant.

By the maximality of r_c and using $c \geq 2r_c + 2$, we have that

$$A_{r_c+1} + B_{r_c+1} + \frac{3(c - 2(r_c+1))}{4(r_c+1) - 1} < 1.$$

Recall that when $x \to 0$, we also have $c \to \infty$ and $r_c \to \infty$, so we obtain $\ln(4/3) + 3\ln(9/8) + \frac{3}{4}\lim_{x\to 0} f(x) - \frac{3}{2} \leq 1$, where we again use that $f(x) \leq \frac{c}{r_c}$. Hence

$$\lim_{x \to 0} f(x) \leq \frac{10}{3} - \frac{4}{3}\ln(4/3) - 4\ln(9/8) < 2.479,$$

as desired. $\square$

8 Conclusion

In this paper we give an entirely novel approach, based on topological notions, for bounding the integrality gap of the Santa Claus problem. This leads to significant improvements on the best known estimates. We believe that this approach will prove to be fruitful in addressing other algorithmic problems involving hypergraph matchings.

As mentioned in the introduction, our argument at the moment does not come with an efficient algorithm for finding an allocation with the promised min-value. This is primarily due to the fact that we do not have a good upper bound on the number of simplices in the triangulation described in the Appendix, which ultimately governs the running time of any algorithmic procedure based on our argument. It would be of great interest to develop methods to make the approach more efficient.

A possible ray of hope comes from recalling the eventual success of turning the initially highly ineffective combinatorial procedure of [10], based on [26], into an efficient algorithm with the same constant factor appriximation. This was achieved through a series of important contributions of several authors, as described in the introduction. Even a quasipolynomial-time algorithm based on our approach that provides *any* constant factor approximation would seem to require new ideas. Such an algorithm would be a first step towards an efficient approximation algorithm that breaks the factor 4 barrier.

Finally, we would like to recall from the introduction that our work on the $(1,\varepsilon)$-restricted case identifies certain parameter choices that seem to capture a key difficulty for the CLP-approach. Specifically, we would like to see a $(1,1/3)$-restricted problem instance that has optimal CLP-target $T^* = 1$, and no allocation of min-value $2/3$.

Acknowledgements: The authors are grateful to Lothar Narins for helpful discussions in the early stages of this work, and to Olaf Parczyk for computing the data in Section 7.

References

[1] M. Adamaszek and J. A. Barmak, *On a lower bound for the connectivity of the independence complex of a graph*, Discrete Mathematics **311** (2011), 2566-2569.

[2] R. Aharoni and E. Berger, *The intersection of a matroid and a simplicial complex*, Trans. Amer. Math. Soc. **358** (2006), 4895-4917.

[3] R. Aharoni, E. Berger, and O. Kfir, *Acyclic systems of representatives and acyclic colorings of digraphs*, J. Graph Theory **59** (2008), 177-189.

[4] R. Aharoni, E. Berger, D. Kotlar, and R. Ziv, *Degree conditions for matchability in 3-partite hypergraphs*, J. Graph Theory **87** (2018), 61-71.

[5] R. Aharoni, E. Berger, and P. Sprüssel, *Two disjoint independent bases in matroid-graph pairs*, Graphs and Combinatorics **31** (2015), 1107-1116.

[6] R. Aharoni, E. Berger, and R. Ziv, *Independent systems of representatives in weighted graphs*, Combinatorica **27** (2007), no. 3, 253-267.

[7] R. Aharoni and P. Haxell, *Hall's theorem for hypergraphs*, J. Graph Theory **35** (2000), no. 2, 83-88.

[8] R. Aharoni, R. Holzman, D. Howard, and P. Sprüssel, *Cooperative Colorings and independent sets of representatives*, Electronic J. Comb. **22(2)** (2015), paper 2.27.

[9] C. Annamalai, C. Kalaitzis, and O. Svensson, *Combinatorial algorithm for restricted max-min fair allocation*, ACM Transactions on Algorithms, **13(3)**:37:137:28, (2017).

[10] A. Asadpour, U. Feige, and A. Saberi. *Santa Claus meets hypergraph matchings*, ACM Transactions on Algorithms, **8(3)**:24:124:9, 2012.

[11] A. Asadpour and A. Saberi, *An approximation algorithm for max-min fair allocation of individual goods*, Proceedings of ACM Symposium on the Theory of Computation (STOC), 2007, 114-121.

[12] N. Bansal and M. Sviridenko, *The Santa Claus problem*, Proceedings of the 38th ACM Symposium on Theory of Computing, pages 3140, 2006.

[13] M. H. Bateni, M. Charikar, and V. Guruswami, *New approximation algorithms for degree lower bounded arborescences and max-min allocation.* Tech Report TR-848-09. Computer Science Department, Princeton University. 2009

[14] E. Berger and R. Ziv, *A note on the cover number and independence number in hypergraphs*, Discrete Mathematics **308** (2008) 2649-2654.

[15] I. Bezáková and Varsha Dani, *Allocating indivisible goods*, SIGecom Exchanges, **5(3)** (2005), 1118.

[16] D. Chakrabarty, J. Chuzhoy, and S. Khanna, *On allocating goods to maximize fairness*, Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (2009), 107116.

[17] T.-H. Chan, Z. Tang, and X. Wu, *On* $(1, \varepsilon)$*-restricted max-min fair allocation problem*, Algorithmica **80** (2018), 2181-2200.

[18] S.-W. Cheng and Y. Mao. *Restricted max-min fair allocation* , Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (2018), 37:137:13.

[19] S.-W. Cheng and Y. Mao. *Integrality gap of the configuration LP for the restricted max-min fair allocation*, CoRR, abs/1807.04152 (2018).

[20] S.-W. Cheng and Y. Mao, *Restricted max-min allocation: approximation and integrality gap*, Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (2019), 38:1 - 38:13.

[21] S. Davies, T. Rothvoss, and Y. Zhang, *A tale of Santa Claus, hypergraphs and matroids*, Proceedings of the 31st ACM-SIAM Symposium on Discrete Algorithms (2020), 2748-2757.

[22] L. Epstein and J. Sgall, *Approximation schemes for scheduling on uniformly related and identical parallel machines*, Algorithmica, 39 (2004), 43-57.

[23] U. Feige, *On allocations that maximize fairness*, Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms (2008), 287293.

[24] D. Golovin, *Max-min fair allocation of indivisible goods*, Carnegie Mellon University Technical Report. CMU-CS-05-144 (2005).

[25] B. Haeupler, B. Saha, and A. Srinivasan, *New constructive aspects of the Lovász local lemma*, Journal of the ACM, **58(6)**:28:128:28 (2011).

[26] P. Haxell, *A condition for matchability in hypergraphs.* Graphs and Combinatorics, **11(3)** (1995), 245248.

[27] P. Haxell, *Independent transversals and hypergraph matchings - an elementary approach*, in Recent Trends in Combinatorics (A. Beveridge, J. Griggs, L. Hogben, G. Musiker, P. Tetali, eds), IMA Volume in Mathematics and its Applications, Springer (2016), 215-233.

[28] P. Haxell, *Topological connectedness and independent sets in graphs*, in: Surveys in Combinatorics 2019, London Mathematical Society Lecture Notes **456**, Cambridge University Press (2019), 89-113.

[29] P. Haxell and L, Narins, *A stability theorem for matchings in tripartite 3-graphs*, Combinatorics Prob. Comput, **27** (2018), 774-793.

[30] K. Jansen and L. Rohwedder, *On the configuration-LP of the restricted assignment problem*, Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), 2670-2678.

[31] K. Jansen and L. Rohwedder, *A quasi-polynomial approximation for the restricted assignment problem*, Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO 2017).

[32] K. Jansen and L. Rohwedder, *A note on the integrality gap of the configuration LP for restricted Santa Claus*. CoRR abs/1807.03626, (2018).

[33] J.K. Lenstra, D.B. Shmoys, and É. Tardos, *Approximation algorithms for scheduling unrelated parallel machines*, Proceedings of the 28th IEEE Symposium on Foundations of Computer Science (1987), 217224.

[34] R. J. Lipton, E. Markakis, E. Mossel, A. Saberi, *Of approximately fair allocation of indivisible goods*, In ACM Conference on Electric Commerce (2005) 125-131.

[35] R. Meshulam, *The clique complex and hypergraph matching*, Combinatorica 21 (2001), Ser. A 102 (2003), 89-94.

[36] R. Meshulam, *Domination numbers and homology*, J. Combin. Theory Ser. A 102 (2003), 321-330.

[37] L. Polacek and O. Svensson, *Quasi-polynomial local search for restricted max-min fair allocation*, 39th International Colloquium on Automata, Languages, and Programming (2012), 726737.

[38] B. Saha and A. Srinivasan, *A new approximation technique for resource-allocation problems*, Proceedings of the 1st Symposium on Innovations in Computer Science (2010), 342357.

[39] O. Svensson, *Santa Claus schedules jobs on unrelated machines*, SIAM Journal on Computing 41(5):1318-1341, 2012. Extended abstract in ACM Symposium on Theory of Computing (STOC) 2011.

[40] G. J. Woeginger, *A polynomial time approximation scheme for maximizing the minimum machine completion time*, Operations Research Letters, 20 (1997) 149-154.

[41] G. J. Woeginger, *When does a dynamic programming formulation guarantee the existence of a fully polynomial time approximation scheme (FPTAS)?*, INFORMS Journal on Computing, 12 (2000), 57-75.

9 Appendix

9.1 The parameter η

As mentioned in Section 2.2 for our results we need only that there exists a graph parameter η that satisfies Fact 1 and Theorems 2.2 and 2.3. In fact such a parameter can be defined in a purely combinatorial way, without any explicit reference to topology. For completeness we begin with a precise definition of η following the treatment of [27], where the required properties are verified. However, the intuition behind the parameter η and how we use it in our proofs is very much topological, as we describe after the definition.

The definition. An *abstract simplicial complex* is a set $\mathcal{A}$ of subsets A of a finite set $V = V(\mathcal{A}) = \cup_{A\in\mathcal{A}} A$ with the property that if $A \in \mathcal{A}$ and $B \subset A$ then $B \in \mathcal{A}$. We call the sets A the *simplices* of $\mathcal{A}$, and the *dimension* of A is $|A| - 1$. The *dimension* of $\mathcal{A}$ is the maximum dimension of any $A \in \mathcal{A}$. Let $\mathcal{A}$ and Σ be abstract simplicial complexes. A function $f : V(\mathcal{A}) \to V(\Sigma)$ is called a *simplicial map from $\mathcal{A}$ to Σ* if $f(A) \in \Sigma$ for every $A \in \mathcal{A}$. We say that $\mathcal{A}$ is a *d-PSC*, i.e. a *pure* d-dimensional simplicial complex, if every maximal $A \in \mathcal{A}$ has the same dimension d. Note then that a d-PSC is the *closure* of the $(d+1)$-uniform hypergraph $\mathcal{A}^d$ consisting of the d-dimensional simplices of $\mathcal{A}$, that is, we form $\mathcal{A}$ from $\mathcal{A}^d$ by adding all subsets of the hyperedges of $\mathcal{A}^d$. For a d-PSC $\mathcal{A}$, the *boundary* $\partial\mathcal{A}$ of $\mathcal{A}$ is the $(d-1)$-PSC that is the closure of the d-uniform hypergraph

$$\{B : |B| = d, |\{A \in \mathcal{A}^d : B \subset A\}| \equiv 1 \mod 2\}.$$

If $\partial(\mathcal{A})$ is empty we say that $\mathcal{A}$ is a *d-dimensional Z_2-cycle*. The abstract simplicial complex Σ is said to be *k-connected* if for each d, $-1 \le d \le k$, for every d-dimensional Z_2-cycle $\mathcal{A}$ and every simplicial map $f : \mathcal{A} \to \Sigma$, there exists a $(d+1)$-PSC $\mathcal{B}$ and a simplicial map $f' : \mathcal{B} \to \Sigma$ such that $\partial\mathcal{B} = \mathcal{A}$ and the restriction $f'|_{\mathcal{A}}$ of f' to $\mathcal{A}$ satisfies $f'|_{\mathcal{A}} = f$.

The independence complex of a graph is an abstract simplicial complex and the value of $\eta(G)$ for a graph G is defined as the largest integer t such that the independence complex $\mathcal{J}(G)$ is $(t-2)$-connected. The parameter η is not explicitly defined in [27], but the main theorems about η are stated and proved there in terms of the above definition of k-connected. (Theorem 2.2 and 2.3 appear as Theorems 11 and 12, respectively.) We may verify Fact 1 as follows. Fact 1(1) follows directly from the definition of η, as saying that Σ is (-1)-connected is the same as saying that $V(\Sigma)$ is nonempty. For the second statement of Fact 1(2), suppose G has an isolated vertex x. Let $\mathcal{A}$ be an arbitrary Z_2-cycle, with a simplicial map f from $\mathcal{A}$ to $\mathcal{J}(G)$. Then f can be extended to a simplicial map from the closure $\mathcal{B}$ of $\{A \cup \{w\} : A \in \mathcal{A}\}$, where $w \notin V(\mathcal{A})$ is a new vertex, by mapping w to x. Since the dimension of $\mathcal{A}$ is arbitrary, this implies that $\eta(G)$ is infinite. Otherwise, if G_2 contains an edge, then by Theorem 2.3 we can keep deleting and/or exploding edges from G_2, one by one, until all edges of G_2 have disappeared. The resulting graph G_{end} still contains G_1. If G_{end} has an isolated vertex, then $\eta(G) \ge \eta(G_{end}) = \infty$ by the above. Otherwise at least one explosion was performed and $G_{end} = G_1$, hence $\eta(G) \ge \eta(G_1) + 1$ by Observation 2.4(ii).

The intuition. In what follows, we describe the topological nature of our work at an intuitive level, without getting into precise details. The topological space X is said to be *k-connected* if for every d, $-1 \leq d \leq k$, every continuous map from the d-dimensional sphere to X extends to a continuous map from the $(d+1)$-dimensional ball to X. This property indicates that X lacks a $(d+1)$-dimensional "hole".

To get a better understanding for the topological core of our arguments, it helps to think of connectedness as defined in the following way, that provides a link between the notion of connectedness for a topological space and our earlier definition for abstract simplicial complexes. This link goes via triangulations of a simplex, which are *geometric simplicial complexes*, that can be viewed both as topological spaces and as abstract simplicial complexes. We say that an abstract simplicial complex Σ is *k-connected* if for every d, $-1 \leq d \leq k$, for every triangulation $\mathcal{T}$ of the boundary of the $(d+1)$-dimensional simplex τ, and every simplicial map f from $\mathcal{T}$ to Σ, there exists a triangulation $\mathcal{T}'$ of the whole of τ that extends $\mathcal{T}$, and a simplicial map f' from $\mathcal{T}'$ to Σ that extends f.

Our argument gives a process that, given an instance $\mathcal{I}$ of the Santa Claus problem with player set P, produces an allocation with the promised min-value. Very broadly speaking, the process has two main phases. Following the proofs of Theorems 2.2 and 2.3, the first phase constructs a triangulation $\mathcal{T}$ of the $(|P|-1)$-dimensional simplex τ, and a simplicial map f from $\mathcal{T}$ to the independence complex of the graph $H(\alpha)$ (defined in Section 2.1), such that the $|P|$-coloring of the points $v \in V(\mathcal{T})$, defined by the "owner" of the α-hyperedge $f(v)$, satisfies the conditions of Sperner's Lemma. The second phase applies Sperner's Lemma to find a multicolored simplex, which corresponds to an independent transversal of $H(\alpha)$, i.e. an allocation for instance $\mathcal{I}$ with min-value at least αT as promised.

Executing the first phase is the main aim of our paper and here is where topological connectedness helps us. The triangulation $\mathcal{T}$ and the map f are built on the faces of τ one by one, in increasing order of dimension. When $\mathcal{T}$ and f on a face σ of dimension d are to be defined, triangulations and maps of all the facets of σ are already in place, forming the boundary of σ, and these need to be extended to a triangulation and a map of the whole of σ. This notion of extending a map from the boundary of σ to the interior is captured by the parameter η, so if η is sufficiently large for each σ, then this extension is possible.

9.2 Demonstrating DE-sequences

Here we demonstrate how to use DE-sequences to show that $\eta(G) \geq 2$ for the cycle $G = C_5$ of length 5. (In fact $\eta(G) = 2$, since the independence complex itself is a 5-cycle, which has a 2-dimensional hole.) For an edge e of G, the graph $G \circledast e$ consists of a single isolated vertex, and hence $\eta(G \circledast e) = \infty$ by Observation 2.4(iii). Therefore e is not explodable, and hence is deletable. Deleting e results in the path P_5 with 5 vertices, and by the definition of deletable edge $\eta(G) \geq \eta(P_5)$. Next consider an edge e' of P_5 joining two of its degree-2 vertices. Again $P_5 \circledast e'$ consists of a single isolated vertex, showing that e' is not explodable and hence deletable. The graph $P_5 - e'$ consists of two components, a P_2 and a P_3. Each of these has a positive value of η by Fact 1(1). Hence

$$\eta(G) \geq \eta(P_5) \geq \eta(P_5 - e') \geq 1 + \eta(P_3) \geq 2$$

by Fact 1(2).