# Approximation Algorithms for Allocation and Network Design 

Présentée le 9 juin 2023
Faculté informatique et communications
Laboratoire de théorie du calcul 2
Programme doctoral en informatique et communications
pour l'obtention du grade de Docteur ès Sciences
par

## Etienne Michel François BAMAS

Acceptée sur proposition du jury
Prof. E. Telatar, président du jury
Prof. O. N. A. Svensson, directeur de thèse
Prof. A. Gupta, rapporteur
Prof. F. Grandoni, rapporteur
Prof. F. Eisenbrand, rapporteur

To my parents,
and Johanna.

I am indebted to many people whose support and encouragement were crucial during my journey towards completing my PhD thesis. The years I spent at EPFL have been a fantastic experience, and I sincerely think that I have been lucky in many ways until this moment. Now is the time to acknowledge it.

I would like to express my deepest gratitude to my advisor, Ola Svensson, for being an exceptional advisor and a role model as a researcher. His neverending enthusiasm is truly contagious, and it is one of the reasons why working in the Theory Lab has been so enjoyable to me. He often aims at solving problems known to be very difficult, and seeing his success with that strategy has been very inspiring. Even if he is not a co-author on several papers in this thesis, he definitely had a big influence on all of them. In addition to being a fantastic researcher, Ola always showed support in various situations, from Covid times to life after the PhD . For this, I would like to thank him.

I am also grateful to the members of my thesis committee Fritz Eisenbrand, Fabrizio Grandoni, Anupam Gupta, and Emre Telatar, for the careful reading of this thesis and the interesting discussion during the defense.

I have been very lucky to work with amazing co-authors during my time at EPFL and the results of this thesis would not have existed without them. For this, I would like to thank Marina Drygala, Paritosh Garg, Andreas Maggiori, Lars Rohwedder, and my advisor Ola. I would also like to give a special thanks to Lars, with whom I learned a lot about the Santa Claus problem, which is now one of my favorite problems. I should also thank Louis Esperet, who was my advisor during my master thesis, and is one of the people that inspired me to pursue a PhD .

I am very grateful to Adam, Andreas, Chantal, Muriel, and Pauline with whom we organized the workshop ALPS in 2022, inviting almost 50 people from all over the world to come to EPFL for one week. It was a pleasure organizing this with you. Special thanks to Chantal for also being a great help in many other administrative or organizational aspects of my life, and for always being friendly and welcoming.

During my PhD, I also had the opportunity to travel a lot for academic visits. For this, I must thank Lars for inviting me in his lab in Maastricht, the Simons Institute for giving me the opportunity to visit UC Berkeley, and Rico Zenklusen for inviting me during an entire month at ETH Zurich.

I enjoyed a lot being part of the Theory Group and I would like to thank all my friends and colleagues there for the friendly atmosphere that they created, the very nice activities we did together, and the fun discussions on many topics. Moreover, I am especially grateful to Xinrui for proofreading the introduction of this thesis.

Although I will not attempt to name all of them, I want to thank all my friends in France and Switzerland for making these years so enjoyable. Mention spéciale à mes colocs et amis Arnout, Quentin, et Sylvain pour ces belles
années de doctorat ponctuées d'épisodes de Top Chef, de moments difficiles avec notre ami Shaun T, de randos, d'innombrables virées à KFC, et j'en passe.
Je tiens à exprimer ma reconnaissance particulière envers Géraldine et (un autre) Sylvain, qui m'ont accueilli chez eux pendant plusieurs mois durant l'épidémie de Covid. Je me suis toujours senti le bienvenu chez eux.
Mes plus sincères remerciements vont à ma famille, en particulier mes parents Nathalie et Philippe, mon frère Arthur, et ma soeur Eliette. Sans leur soutien, leur amour, et leurs encouragements sans faille depuis 28 ans, je ne serais sans doute pas là où je suis aujourd'hui.
Je voudrais dédier ces dernières lignes à ma compagne Johanna. Johanna, j'aurais sans doute dû écrire ton nom dans chacun des 3 derniers paragraphes tellement tu es importante à mes yeux. Je veux te remercier pour ces belles années à tes côtés, et j'espère en vivre beaucoup d'autres avec toi.

## ABSTRACT

In this thesis, we give new approximation algorithms for some NP-hard problems arising in resource allocation and network design. As a resource allocation problem, we study the Santa Claus problem (also known as the MaxMin Fair Allocation problem) in which we are given $m$ agents, $n$ indivisible resources, and we have to allocate resources in order to maximize the happiness of the least happy agent. For the network design part, we study two key problems: the Steiner Forest problem and the Matching Augmentation problem.

In the first part of this thesis, we give improved guarantees for the Santa Claus problem in two prominent settings. First, we consider the Santa Claus problem in the setting where the happiness of each agent $i$ is an arbitrary linear function $f_{i}$. Obtaining a constant factor approximation in that setting is a major open problem in the area of approximation algorithms. In 2009, the MaxMin Arborescences problem was identified as a key special case that appears to capture most of the difficulty of the general problem. Even in that special case, a constant factor approximation has remained elusive, and the current best algorithm only guarantees a polylogarithmic approximation in quasi-polynomial time. We give an exponential improvement to this, an $O(\operatorname{poly}(\log \log (n)))$-approximation in quasi-polynomial time for the MaxMin Arborescences problem. Our second result in this part considers the Santa Claus problem in the restricted assignment case: all the agents have the same happiness function $f$, but each agent $i$ is only interested in a subset $\Gamma_{i}$ of the resources. When $f$ is a monotone submodular function, we show that we can obtain an $O(\log \log n)$-approximation in polynomial time. Before our work, comparable results in this setting were only known for the case where $f$ is a linear function.

In the second part of this thesis, we start with the Steiner Forest problem which is defined as follows. Given a graph $G$ and an arbitrary set of $k$ terminal pairs $\left\{\left\{s_{1}, t_{1}\right\}, \ldots,\left\{s_{k}, t_{k}\right\}\right\}$, the goal is to return a minimum-weight subgraph that connects all the pairs. In 1996, Awerbuch, Azar, and Bartal showed that an intuitive greedy algorithm guarantees an $O\left(\log ^{2} k\right)$-approximation. Our first result is to show that, in fact, the greedy algorithm guarantees an $O(\log (k) \cdot \log \log (k))$-approximation, which is nearly tight in light of a known lower bound of order $\Omega(\log k)$. Interestingly, our analysis also gives important insights in the online setting of the problem, in which the pairs are revealed one by one in an adversarial order. Our last result is on the Matching Augmentation problem, a key problem to compute cheap 2-edge connected subgraphs. We give a simple polynomial-time algorithm that guarantees a better-than-2 approximation when compared to the standard relaxation of the problem known as the cut LP. In contrast, previous better-than-2 approximations were much more complicated and did not compare to this simple relaxation.

Keywords-approximation algorithms, Santa Claus, maxmin allocation, network design, Steiner forest, matching augmentation.

## RÉSUMÉ

Dans cette thèse de doctorat, nous proposons de nouveaux algorithmes d'approximation pour des problèmes NP-difficiles d'allocation de ressources et de conception de réseaux. Dans le domaine d'allocation de ressources, nous étudions le problème du Père Noël (aussi connu sous le nom de problème d'allocation MaxMin) : étant donné $m$ individus et $n$ ressources indivisibles, le but est d'allouer les ressources afin de maximiser le bonheur de l'individu le moins heureux. Dans le domaine de la conception de réseaux, nous étudions deux problèmes importants : le problème de la forêt de Steiner et le problème d'augmentation de couplage.

Dans la première partie de cette thèse, nous donnons de meilleures garanties pour le problème du Père Noël dans deux cas importants. En premier lieu, nous considérons le problème du Père Noël dans le cas où le bonheur de chaque individu $i$ est défini par une fonction linéaire quelconque $f_{i}$. Obtenir une approximation de facteur constant dans ce cas est un problème ouvert majeur dans le domaine des algorithmes d'approximation. En 2009, le problème d'arborescences MaxMin a été identifié comme un cas particulier qui semble capturer l'essence du problème général. Même dans ce cas particulier, une approximation de facteur constant échappe encore largement aux techniques actuelles, et l'état de l'art ne peut garantir qu'une approximation de facteur polylogarithmique en temps quasi-polynomial. Nous proposons une approximation de facteur $O(\operatorname{poly}(\log \log (n)))$ en temps quasi-polynomial pour le problème d'arborescences MaxMin, ce qui constitue une amélioration exponentielle de l'état de l'art. Notre deuxième résultat traite du problème du Père Noël avec des restrictions d'allocation : le bonheur de tous les individus est décrit par la même fonction $f$, mais chaque individu $i$ n'accepte que les ressources appartenant à un sous-ensemble $\Gamma_{i}$. Dans le cas où $f$ est une fonction sous-modulaire et monotone, nous obtenons une approximation de facteur $O(\log \log n)$ en temps polynomial. Avant nos travaux, des résultats comparables n'étaient connus que dans le cas où $f$ est une fonction linéaire.

Dans la deuxième partie de cette thèse, nous commençons par étudier le problème de la forêt de Steiner : étant donné un graphe $G$ et un ensemble de $k$ paires de terminaux $\left\{\left\{s_{1}, t_{1}\right\}, \ldots,\left\{s_{k}, t_{k}\right\}\right\}$, nous devons trouver le sousgraphe de poids minimum qui connecte toutes les paires. En 1996, Awerbuch, Azar, et Bartal ont montré qu'un simple algorithme glouton garantit une approximation de facteur $O\left(\log ^{2} k\right)$. Notre premier résultat est de montrer que l'algorithme glouton retourne en fait une approximation de facteur $O(\log (k) \cdot \log \log (k))$, ce qui est presque optimal en raison d'une borne inférieure connue d'ordre $\Omega(\log k)$. De plus, nos résultats apportent une compréhension nouvelle de la version en-ligne du problème, lorsque les paires sont révélées les unes après les autres par un adversaire. Notre dernier résultat concerne le problème d'augmentation de couplage, un problème important pour trouver des sous-graphes 2 -connectés de poids minimum. Nous proposons un algorithme simple (et de complexité polynomiale) qui garantit une approximation de facteur strictement inférieur à 2 par rapport à
la valeur de la relaxation linéaire du problème. En comparaison, les autres algorithmes connus pour garantir une approximation de facteur strictement inférieur à 2 sont bien plus complexes et ne peuvent pas être comparés à cette relaxation linéaire.

Mots-clefs- algorithmes d'approximation, Père Noël, allocation maxmin, conception de réseaux, forêt de Steiner, augmentation de couplage.

## CONTENTS

1 Introduction ..... 1
1.1 Our contributions ..... 3
1.1.1 The Santa Claus problem ..... 3
1.1.2 Simple Algorithms for Network Design ..... 4
1.2 How to read this thesis ..... 6
1.3 Linear Programming ..... 6
1.3.1 The use of LPs in this thesis ..... 10
I The Santa Claus Problem
2 Introduction ..... 15
2.1 MaxMin Arborescences ..... 15
2.1.1 Our results for MaxMin Arborescences ..... 17
2.2 The Restricted Submodular Santa Claus ..... 18
2.2.1 Our results for the Restricted Submodular Santa Claus ..... 19
2.3 Probabilistic lemmas ..... 19
3 MaxMin Arborescences ..... 21
3.1 Our techniques and intuition ..... 21
3.2 Formal proof structure ..... 23
3.2.1 Bounded depth solution ..... 24
3.2.2 Local congestion and layered instances ..... 25
3.2.3 Connecting the dots ..... 26
3.2.4 Bottom-to-top pruning ..... 27
3.3 Local to global congestion ..... 28
3.4 Computing a solution with locally low congestion ..... 34
3.4.1 Preprocessing the LP solution ..... 35
3.4.2 The main rounding ..... 35
3.5 From single source to multiple sources ..... 45
3.6 APX-hardness of MaxMin Arborescences ..... 48
4 The Restricted Submodular Santa Claus ..... 51
4.1 Overview of previous techniques and our new ideas ..... 51
4.2 Reduction to hypergraph matching problem ..... 54
4.2.1 Reduction to unweighted hypergraph matching ..... 56
4.3 Matchings in regular hypergraphs ..... 57
4.3.1 Overview and notations ..... 58
4.3.2 Properties of resource sets ..... 59
4.3.3 Selection of configurations ..... 60
4.3.4 Assignment of resources to configurations ..... 64
iI Network Design Problems
5 Introduction ..... 73
5.1 The Steiner Forest Problem ..... 73
5.1.1 Our results ..... 75
5.2 The Matching Augmentation Problem ..... 78
5.2.1 Our results ..... 80
6 Improved bounds for Greedy Steiner Forest ..... 83
6.1 The idea behind the proof ..... 83
6.2 Proof of Theorem 5.1 ..... 88
6.2.1 Problem definition and notation ..... 88
6.2.2 Preliminary results and preprocessing of the instance ..... 89
6.2.3 Overview of the proof ..... 94
6.2.4 Building a balanced dual solution ..... 97
6.2.5 Inductive proof using balanced dual solutions ..... 102
6.3 Proof of Theorem 5.2 and Theorem 5.3 ..... 108
7 A Simple Approximation Algorithm for MAP ..... 113
7.1 Our proof technique ..... 113
7.2 The analysis of the LP-based Algorithm ..... 114
7.2.1 Preliminaries ..... 114
7.2.2 The analysis of the algorithm ..... 115
8 Conclusion and Open Problems ..... 121
III Appendix
A Deferred proofs for MaxMin Arborescences ..... 125
A. 1 A challenging instance for randomized rounding ..... 125
A. 2 Preprocessing the Path LP ..... 129
в Deferred proofs for the Restricted Submodular Santa Claus ..... 133
B. 1 Reduction to hypergraph matching problem ..... 133
B.1.1 Solving the Configuration LP ..... 133
B.1.2 Clusters ..... 134
B. 2 Properties of resource sets ..... 137
Bibliography ..... 143

Figure 5.1 An integrality gap example. 81
Figure 6.1 An example with two cost classes before charging or clustering. 85
Figure 6.2 An example with two cost classes after charging or clustering. In the top left corner, smaller pairs are much more expensive than the pairs that created the big dual ball while we have the opposite situation in the bottom right corner. 85
Figure 6.3 The girth argument. 93
Figure 7.1 An example of a bad DFS tree. 114
Figure 7.2 On the left side, the case when the first edge selected out of $u$ is heavy. On the right the case when the first edge selected out of $u$ is light. 119
Figure B. $1 \quad$ The directed network and an $s$ - $t$ cut. 139

Many problems faced every day can be phrased as a discrete optimization problem, in which there is a finite set of possible solutions. Some examples (among many) include finding the closest supermarket, finding the best assignment of students to universities, or even how to design the energy network of a country. It is quite clear that some of these questions are of significant and global importance for a big part of the population. Many of these problems are now solved automatically by a computer that runs an algorithm, and it is therefore natural that these algorithms have a high importance in our lives. Even seemingly minor decisions can have big influence on society at large when millions of individuals rely on them, and this naturally creates the need for a thorough and rigorous understanding of these algorithms.

Regrettably, many of these problems can have a complex structure which often makes it very difficult (if not impossible) to find the best solution in a reasonable amount of time. This obstacle motivated the study of approximation algorithms, which are algorithms designed to find approximate solutions to optimization problems. Approximation algorithms are often used when it is impractical or impossible to find the best solution, because the time required to find that solution is prohibitive. The output of these algorithms may not be perfect, but can still provide valuable insights and help us make better decisions. This is the main topic of this thesis, which focuses on the design and analysis of approximation algorithms with provable, worst-case guarantees.

As a concrete illustration, imagine the following situation. You are managing a factory in which there is a set $M$ of $m$ identical machines that can perform jobs for you. Each day you are given a new set $J$ of $n$ jobs to be distributed among the $m$ machines, where job $i$ takes time $p_{i}$ to be completed by a machine (the machines can perform several types of jobs, which may not need the same time). Because you want to save electricity and go home earlier, you are interested in how to assign these tasks to the machines so as to minimize the overall completion time (also called the makespan). In future references, we will call this problem the Makespan problem. A simple solution would be to try all possible assignments of jobs to machine and keep the best one; this is what we call an exhaustive search. But how long would your computer need to perform this exhaustive search? Let us perform a quick calculation, imagine you have 40 tasks and 10 machines (this is a very small factory): then there are $10^{40}$ possible assignments. Unfortunately, this number is already astronomically big. Assuming we had all the computational power of the world (often estimated to be around $10^{20}$ basic operations per second, see [57] and references therein), then we would need at least 317 years to find the best assignment. With 44 jobs we would need more than a million years.

This exponential growth of the number of solutions is sometimes referred to as combinatorial explosion. Very often, this makes it prohibitive to try an
exhaustive search and necessitates the design of a faster algorithm. Traditionally, an algorithm is considered to be efficient if it runs in time that is a polynomial function of the input size, and the class of problems for which there is such an algorithm is denoted by P. Unfortunately, the MakeSpan problem belongs to a wide and significant class of problems called NP-hard problems, for which it is widely believed that there is no such efficient algorithm ${ }^{1}$.

Fortunately, this is not the end of the story. Here, we were looking for an exact solution to the Makespan problem but an approximate solution might be satisfying as well. For instance, for the Makespan problem, there exists an algorithm that runs in polynomial time and returns a solution that is always within $1 \%$ of the optimal solution (see [54]). This result shows that it is possible to sacrifice a tiny bit of accuracy to save one of the most valuable resources, time. Clearly, it is not worth waiting for millions of years in order to improve the solution by $1 \%$.

The above example falls into the realm of approximation algorithms, which are widely used when facing NP-hard problems. We say that an algorithm has approximation ratio $\alpha$ if it always returns a solution that is not worse than $\alpha$ times the cost of the best solution (for instance, the above algorithm has approximation ratio 1.01). Now, a fundamental question that comes to mind is how much accuracy do we need to lose in order to be able to find an efficient algorithm? To answer this question, researchers have designed increasingly elaborate algorithms to get better and better approximation ratios. These results often came with very interesting insights and new algorithmic techniques that turned out to be useful for other problems. Unfortunately, during this quest of the best approximation guarantee, another crucial property often got lost on the way: simplicity. This is also very important, because extremely complicated algorithms are difficult to use in practice and often remain in the world of theory. For many problems, it is therefore tempting to design a simple and intuitive algorithm. In the Makespan problem, we can consider the following very simple algorithm: schedule the jobs one by one (in any order), so that each job gets assigned to the machine that has been assigned the least amount of work so far. This is very natural as we intuitively want to balance the loads on the machines. It turns out that this algorithm has an approximation ratio of 2 (see [85]). It is extremely simple and runs very fast in practice, which might be an advantage over a more complicated algorithm with a better approximation ratio.

In this thesis, we focus on problems that are known to be NP-hard, which means that there is probably no efficient algorithm to solve the problem (like the Makespan problem defined above). We start by focusing on the Santa Claus problem, a fundamental problem in resource allocation. Second, we study two key network design problems: the Steiner Forest problem and the Matching Augmentation problem. For the Santa Claus problem, we aim at finding the best approximation ratio that an efficient algorithm can guarantee (without worrying about the simplicity of the algorithm). For the network design problems, we focus on analyzing the approximation ratio guaranteed

1 In fact, NP-hard problems can have polynomial time algorithms only if $\mathrm{P}=\mathrm{NP}$, an assertion widely believed to be false. The $\mathrm{P} \neq \mathrm{NP}$ conjecture is one of the Millenium Prize Problems, a list of seven of the most important problems in mathematics.
by a simple and intuitive heuristic. Interestingly, all our results leverage some linear programming relaxation of the problem at hand. We will elaborate further on this concept at the end of this introduction (see Section 1.3).

### 1.1 OUR CONTRIBUTIONS

We give new approximation algorithms for various settings of the Santa Claus problem, a fundamental problem in resource allocation, and two interesting problems in the area of network design. We detail here these contributions.

### 1.1.1 The Santa Claus problem

First, we study the Santa Claus problem, also known as MaxMin Fair Allocation. In this problem, we have to assign resources to players in order to make the least happy player as happy as possible. Formally, we have a set $P$ of $m$ players, a set $R$ of $n$ resources, and each player $i$ is equipped with an arbitrary valuation function $f_{i}: S \subseteq R \mapsto \mathbb{R}_{+}$. We have to find a partition $\left(S_{1}, \ldots, S_{i}, \ldots, S_{m}\right)$ of the resource set such that

$$
\min _{i \in P} f_{i}\left(S_{i}\right)
$$

is maximized (note that a resource can be assigned to at most one player).
The most classic version of this problem is when each $f_{i}$ is an arbitrary linear valuation function. We can rephrase this case as the setting where player $i$ assigns value $v_{i j}$ to resource $j$ and the goal is to find an assignment of resources to players $\sigma: R \mapsto P$ such that

$$
\min _{i \in P} \sum_{j: \sigma(j)=i} v_{i j}
$$

is maximized. Proving if a constant factor approximation can be obtained efficiently or not is a major open problem in the area of approximation algorithms [13, 78, 85]. In fact, there is a huge gap in our understanding of this fundamental problem. The state-of-the-art gives only an $O\left(n^{\varepsilon}\right)$-approximation in time $n^{O(1 / \varepsilon)}$ for any fixed $\varepsilon>0$, and an $O(\operatorname{poly} \log (n))$-approximation in quasi-polynomial time ${ }^{2}$ [21]. On the other hand, it is only known that getting a $(2-\varepsilon)$-approximation (for any fixed $\varepsilon>0$ ) is NP-hard [17, 68]. The mentioned algorithm is obtained by reducing to a problem similar to the MaxMin Arborescences problem, which is defined as follows: Given a directed graph $G=(V, E)$ with sources and sinks, our goal is to find vertex disjoint arborescences rooted in the sources such that at each non-sink vertex of an arborescence the out-degree is at least $k$, where $k$ is to be maximized. This problem can be seen as a special case of the Santa Claus problem (see the beginning of Part i for more details).

This special case is of particular interest, as it seems to capture most of the difficulty of the Santa Claus problem with linear valuation functions. Indeed, an $O\left(n^{\varepsilon}\right)$-approximation in time $n^{O(1 / \varepsilon)}$ for any fixed $\varepsilon>0$, and an

[^0]$O$ (polylog$(n))$-approximation in time $n^{O(\log n)}$ were first obtained only for the MaxMin Arborescence problem before being generalized to the general linear valuations case [15]. Our main result on this problem is to obtain an exponential improvement over previous works, an $O(\operatorname{poly}(\log \log n))$ approximation in time $n^{O(\log n)}$ for the MaxMin Arborescence problem. We also show that getting a better than $(\sqrt{e /(e-1)}-\varepsilon)$-approximation for the MaxMin Arborescence problem is NP-hard for any fixed $\varepsilon>0$.

Lastly, we turn our attention to the Santa Claus problem with valuation functions $f_{1}, f_{2}, \ldots, f_{m}$ that are submodular and not linear. This is of interest, as linear functions fail to capture properties that are often crucial. For instance, submodular function capture the diminishing returns property in economics, which is particularly relevant in the context of the Santa Claus problem. In the linear case, a popular line of research (started by Bansal and Srividenko [14]) has considered the restricted assignment case. This is the special case where each player $i$ is given a set of desired resources $\Gamma_{i}$ and the individual valuation functions are defined as $f_{i}(S)=f\left(S \cap \Gamma_{i}\right)$ for a global linear function $f$. This can also be interpreted as maximizing $\min _{i} f\left(S_{i}\right)$ with additional assignment restrictions, i.e., resources can only be assigned to certain players. In that case, it has been known for a long time how to obtain an $O(1)$-approximation in polynomial time [7, 39]. Our second result in Part i makes comparable progress for the submodular variant: If $f$ is a monotone submodular function, we can in polynomial time compute an $O(\log \log (n))$-approximate solution.

### 1.1.2 Simple Algorithms for Network Design

In the second part of this thesis, we are interested in the design and analysis of simple algorithms for problems in the wide area of network design, which contains problems about designing cheap networks that are robust to edge failures. Being able to design such networks is a fundamental question, both in practice and in theory. The first problem we study is the Steiner Forest problem: we are given a weighted graph $G=(V, E)$ with a weight function $w$ : $E \mapsto \mathbb{R}_{+}$and a set of $k$ pairs of vertices $\mathcal{P}:=\left\{\left(s_{1}, t_{1}\right),\left(s_{2}, t_{2}\right), \ldots,\left(s_{k}, t_{k}\right)\right\}$. The goal is to compute the cheapest set of edges such that any two vertices belonging to the same pair are connected (note that some vertices might not appear in any pair). This is a generalization of two famous optimization problems: Minimum Spanning Tree, which is the special case where $\mathcal{P}$ contains all possible pairs, and Steiner Tree, where there is a special vertex $s$ that belongs to all the pairs in $\mathcal{P}$. The Steiner Forest problem was at the heart of many new algorithmic techniques and several polynomial-time algorithms are known to guarantee a factor 2 approximation for this problem [47, 58]. However, these algorithms are fairly involved and are not necessarily the most intuitive algorithms. The simplest algorithm is arguably the greedy algorithm which can be informally defined as follows:

1. Order the pairs in $\mathcal{P}$ in an arbitrary order.
2. For $i=1$ to $k$, connect $\left\{s_{i}, t_{i}\right\}$ with the shortest path in the current metric, then contract the metric along the chosen path.

Although simple an intuitive, greedy proved itself challenging to analyze. Awerbuch, Azar, and Bartal [8] showed with a beautiful argument that this algorithm guarantees a factor $O\left(\log ^{2}(k)\right)$ approximation for any ordering of $\mathcal{P}$. They also showed that this algorithm is no better than $\Omega(\log (k))$ (even if we select a specific ordering of $\mathcal{P}$ ), and they conjectured that greedy gives in fact an $O(\log (k))$-approximation. Whether this conjecture is true is particularly interesting, because it would also mean that greedy is the optimal algorithm in the online setting, where the pairs are given in an adversarial order. Unfortunately, this conjecture has not seen any progress since its formulation. Our first result in this part is to show that if we order the pairs in non-increasing order of distance in the graph $G$ (i.e. such that $d_{G}\left(s_{i}, t_{i}\right) \geqslant$ $d_{G}\left(s_{j}, t_{j}\right)$ for all $\left.j>i\right)$, then greedy guarantees an $O(\log (k) \cdot \log \log (k))$ approximation. While it does not imply an improved bound for the online setting, our proof has important consequences even in that case.

Then, we turn our attention to the Matching Augmentation problem (MAP), which is defined as follows. We are given a graph $G=(V, E)$, a weight function $w$ such that $w(e) \in\{0,1\}$ for all $e \in E$ and the set of edges that are given weight 0 forms a matching $M$. The goal is to compute the cheapest set $E^{\prime} \subseteq E$ of edges such that $G^{\prime}=\left(V, E^{\prime}\right)$ is 2-edge connected (note that it is forbidden to take two copies of the same edge $e \in E$ ). The name "matching augmentation" stems from the fact that the matching $M$ can be included in the solution for free, and the challenge only remains in how to "augment" this matching to make it 2-edge connected. This problem has recently received significant attention as an important step towards better approximation algorithms for finding cheap 2-edge connected subgraphs. A 2-approximation can be obtained for this problem via many standard techniques, and the challenge is to overcome this barrier to obtain better-than-2 approximations. Previous works on this problem culminated in a $\frac{13}{8}$ approximation algorithm [45]. However, these algorithms and their analysis are very involved (the mentioned result is proven in a paper of 60 pages) and do not compare against the problem's well-known LP relaxation called the cut LP. Our second result in this part is a simple algorithm that, guided by an optimal solution to the cut LP, first selects a depth-first search (DFS) tree and then finds a solution to MAP by computing an optimum augmentation of this tree. Using properties of extreme point solutions, we show that our algorithm always returns (in polynomial time) a better-than-2 approximation when compared to the cut LP. We thereby also obtain an improved upper bound on the integrality gap of this natural relaxation. The algorithm is simple enough to be described in 3 basic steps, and the analysis takes roughly 5 pages.

This thesis is split into two main parts, each part presenting a detailed motivation of the problems studied (including a fairly complete literature review), followed by the proof of our results. Part i contains our results on the Santa Claus problem, while Part ii contains our results on the Steiner Forest problem and the Matching Augmentation problem. In Chapter 8, we conclude the thesis with some interesting open problems arising out of our work. For clarity and better readability, some proofs are deferred to the Appendix in Part iii. Before moving on to the main body of the thesis, we finish this chapter by Section 1.3, which contains an introduction to Linear Programming, a fundamental tool in all our results.

### 1.3 LINEAR PROGRAMMING

A major challenge on the way to our results is to compare the cost of the solution returned by an algorithm to the cost of the optimum solution, which is unknown. To make matters worse, the optimum solution is often very unstable to small changes in the instance and very difficult to reason about in general. Fortunately, linear programming relaxations can be used to cope with these issues. In this section, we introduce this powerful tool which we will use extensively in this thesis.

Linear programs play a crucial role in the design and analysis of approximation algorithms. We illustrate its role on the Makespan problem we defined earlier. In the previous version of MAKESPAN we considered, all the machines were identical. But in general, it might be that the factory has several types of machines, which can perform different sets of jobs. To account for this, we define $p_{i j}$ as the time that machine $i$ needs to finish job $j$. For instance, if a machine cannot perform a specific job, one can model this by setting $p_{i j}=\infty$ for the corresponding $i, j$. This more general version is called makespan minimization on unrelated machines, and we will refer to this version as Unrelated Makespan. It is significantly more challenging than the version on identical machines. In particular, all the results we cited for the Makespan problem do not hold anymore for the Unrelated Makespan problem. For instance, it is easy to check that the simple factor 2 approximation for MaKespan does not guarantee any finite approximation ratio for Unrelated Makespan.

A natural question arises: what is the best approximation ratio an efficient algorithm can guarantee for the Unrelated Makespan problem? To tackle this question, we will use a linear program, a concept that we define next.

Definition 1.1 (Linear Program (LP) [85]). A linear program is formulated in terms of some number of decision variables that represent some sort of decision that needs to be made. The variables are constrained by a number of linear inequalities and equalities called constraints. Any assignment of real numbers to the variables such that all of the constraints are satisfied is called a feasible solution. A linear program is said to be feasible if it admits a feasible solution. In addition to the constraints, linear programs are defined by a linear function of the decision variables called the objective function. The linear program seeks
to find a feasible solution that either maximizes or minimizes this objective function. Such a solution is called an optimal solution.

A breakthrough result by Khachiyan [60] shows that there is an efficient (i.e. polynomial time) algorithm that decides, for any given linear program, if there exists an optimal solution or not. Moreover, if an optimal solution exists, then this algorithm returns such a solution in polynomial time.

We are now ready to use this concept in the context of approximation algorithms. To this end, let us consider a slightly easier version of the Unrelated Makespan problem where instead of building the assignment we only want to decide if there is an assignment of value less than $T$ for $T>0$ that is given in input. This is the decision version of the problem where the algorithm simply has to answer YES or NO. While this may seem restrictive, we can combine an algorithm for the decision version with any standard binary search framework to find the smallest $T$ such that there exists an assignment of value at most $T$ (let us call this optimum value $T^{*}$ ). This is often called an estimation algorithm, in contrast to a constructive algorithm which also returns a schedule that achieves the claimed objective $T^{*}$.

However, even the decision version of Unrelated Makespan is NP-hard, hence we need to relax a little bit the problem. Fix an $\alpha \geqslant 1$, we say that a decision algorithm is an $\alpha$-approximation to the decision version of the problem if:

1. For any $T \geqslant T^{*}$, the algorithm returns YES.
2. If $T<T^{*} / \alpha$, the algorithm returns NO.
3. If $T^{*} / \alpha \leqslant T<T^{*}$, the algorithm can output YES or NO indifferently.

We recall here that $T$ is given as input to the algorithm. Combining an $\alpha$-approximation to the decision problem with a binary search gives an $\alpha$ approximation to the estimation problem, i.e. an algorithm that returns in polynomial time a value $\hat{T}$ such that $T^{*} \leqslant \hat{T} \leqslant \alpha T^{*}$. In the rest of this section, we focus on obtaining an $\alpha$-approximation to the decision problem, with $\alpha$ as close to 1 as possible (note that $\alpha$ is always at least 1 ). A landmark result by Lenstra, Shmoys, and Tardos [68] shows that this is possible for $\alpha=2$. The rest of the section is devoted to provide some brief introduction to this result. To this end, we formulate a linear program and our estimation algorithm will simply return YES if the linear program is feasible, and NO otherwise. To decide if the linear program is feasible we can use Khachiyan's algorithm mentioned above. Now, we are ready to formulate the decision variables and constraints of the linear program.

For any machine $i$, job $j$ we define the decision variable $x_{i j}$ that takes value 1 if the job $j$ is assigned to machine $i$ and 0 otherwise. A solution to our problem can now be described as the binary vector $x:=\left(x_{i j}\right)_{i \in M, j \in J}$. In order for the solution to be feasible for our problem, it needs to respect a few constraints, that we can enforce using linear inequalities. First, any job $j$ needs to be assigned to one machine to be processed therefore it must be that

$$
\sum_{i \in M} x_{i j}=1
$$

for all $j \in J$. Finally, it must be that the load on any machine is no more than $T$ hence we can write

$$
\sum_{j \in J} x_{i j} p_{i j} \leqslant T
$$

for any machine $i \in M$. Here we are simply designing a decision algorithm that outputs YES if and only if there is a feasible solution to the linear program, hence the objective function does not matter and we can simply define it to be the constant function equal to 0 , which we omit for clarity. In summary, we have rephrased the decision version of the Unrelated Makespan problem into the following set of constraints:

$$
\begin{array}{rr}
\sum_{j \in J} x_{i j} p_{i j} \leqslant T & \forall i \in M \\
\sum_{i \in M} x_{i j}=1 & \forall j \in J \\
x_{i j} \in\{0,1\} & \forall i \in M, \forall j \in J . \tag{1.3}
\end{array}
$$

Unfortunately, this is not a linear program, because of the last constraint $x_{i j} \in\{0,1\}$ (which is not linear). The above program is often called an integer linear program, because the decision variables are forced to be integers. Solving integer linear programs in general is NP-hard, and this is not surprising since the above integer linear program is equivalent to the exact decision version of Unrelated Makespan, which is also NP-hard. Hence a common approach in algorithm design is to relax the integer constraints. The fact of relaxing the constraint naturally gave the name of linear programming relaxation. Here, we relax the integer constraints into $x_{i j} \geqslant 0$ to obtain the following linear program.

$$
\begin{array}{rr}
\sum_{j \in J} x_{i j} p_{i j} \leqslant T & \forall i \in M \\
\sum_{i \in M} x_{i j}=1 & \forall j \in J \\
x_{i j} \geqslant 0 & \forall i \in M, j \in J .
\end{array}
$$

This is the linear programming relaxation of our decision problem. Because it is now a linear program, we can use Khachiyan's algorithm to decide in polynomial time if there exists a feasible solution. Let us denote by $T_{\mathrm{LP}}$ the smallest $T$ such that the above LP is feasible. Then clearly $T_{\mathrm{LP}} \leqslant T^{*}$, because Constraint (1.6) is less restrictive than Constraint (1.3). Recall that our decision algorithm simply returns YES if and only if the LP is feasible. Therefore if $T \geqslant T^{*}$, then the algorithm always returns YES. To guarantee an $\alpha$-approximation, it only remains to prove that

$$
\frac{T^{*}}{T_{\mathrm{LP}}} \leqslant \alpha
$$

always holds. This ratio is often referred to as the integrality gap the relaxation, and it measures how strong the relaxation is. The bigger $\frac{T^{*}}{T_{\mathrm{LP}}}$ is, the less useful the relaxation is.

Here, it turns out that $\frac{T^{*}}{T_{\mathrm{LP}}}$ can be as big as $m$ (the number of machines), hence the approximation ratio $\alpha$ of our algorithm is at least $m$, far from our goal of a 2-approximation. Consider the following example. There are $m$ machines, and a unique job $\{1\}$ that has load $m$ on all the machines (i.e. $p_{i 1}=m$ for any machine $i$ ). Then the linear program is feasible for $T=1$, by setting $x_{i 1}=1 / m$ for all machines $i$. Clearly the job is assigned once in total, while every machine receives only a $1 / m$ fraction of the job, for a total load of 1 . Hence $T_{L P} \leqslant 1$. However, $T^{*} \geqslant m$ because no matter where we assign the job, it takes space at least $m$.

To cope with this, a useful idea is to introduce more restrictive constraints to the linear program. While doing this, we must ensure that an integral solution also satisfies those constraints, as otherwise it would not be a relaxation anymore. In order to find a suitable candidate constraint, we must understand intuitively what went wrong in the previous relaxation. Here, what happened is that the linear programming relaxation is allowed to spread a huge job on many machines, while the integral solution has to assign it entirely on a single machine. But clearly, if the objective is $T$, then any integral solution cannot assign job $j$ on machine $j$ if $p_{i j}>T$. This additional insight can be incorporated into the relaxation in the following sense. For any job $j$, we denote by $B_{T}(j)$ the set of machines such that $p_{i j}>T$. We then obtain the following new relaxation.

$$
\begin{array}{rrr}
\sum_{j \in J} x_{i j} p_{i j} & \leqslant T & \forall i \in M \\
\sum_{i \in M \backslash B_{T}(j)} x_{i j}=1 & \forall j \in J \\
x_{i j} \geqslant 0 & \forall i \in M, j \in J . \tag{1.8}
\end{array}
$$

(1.5) was simply replaced by Constraint (1.8), in which we allow the LP to assign job $j$ to machine $i$ only if $p_{i j} \leqslant T$. Note that the previous example is not feasible for $T=1$ on this new relaxation. It turns out that this seemingly benign modification of the LP now guarantees an integrality gap of at most 2. This was showed in [68]. We do not give here the proof here as it is beyond the scope of this introduction. To obtain such a result, [68] used a procedure that transforms the feasible solution $x$ into an integral solution $\hat{x}$ (i.e. a solution such that $\hat{x}_{i j} \in\{0,1\}$ ). Such a procedure is often referred to as a rounding scheme, a crucial concept in many results in approximation algorithms. Indeed, after solving the LP, we obtain a feasible solution $x$ where $x_{i j} \in[0,1]$ and we would like to show that there exists an integral solution that is not much worse than $x$. The most intuitive way to proceed is to transform $x$ into an integral vector $\hat{x}$ where $\hat{x}_{i j} \in\{0,1\}$. Along the way, if we manage to show that

$$
\begin{equation*}
\sum_{j \in J} \hat{x}_{i j} p_{i j} \leqslant \alpha \cdot\left(\sum_{j \in J} x_{i j} p_{i j}\right) \tag{1.10}
\end{equation*}
$$

for all $i \in M$, then we will have shown that the integrality gap of the relaxation is at most $\alpha$, and that our algorithm is an $\alpha$-approximation. This is precisely what is shown in [68] for $\alpha=2$, hence obtaining a 2-approximation
for our problem. Finding the good rounding scheme is however non-trivial, as some intuitive rules do not work. For instance, one might be tempted to simply assign each job $j$ to the machine $i$ such that $x_{i j}$ is maximum. This intuitively follows the decisions made by the LP, but unfortunately does not guarantee a factor 2 . Another intuitive rounding is simply to assign randomly each job $j$ randomly to machine $i$ with probability $x_{i j}$, but again this does not guarantee a factor 2 . In fact, the rounding scheme of [68] is non-trivial and crucially exploits an additional property of the feasible solution $x$.
Indeed, [68] assumes that the feasible point $x$ that we need to round is an extreme point solution to the LP. A feasible solution $x$ is said to be an extreme point if it cannot be expressed as a convex combination of other feasible points. It turns out that if the LP has a bounded feasible region (which is the case in many relevant settings in approximation algorithms), then there always exists an extreme point that also achieves the optimum value, and an optimum extreme point can be found in polynomial time. These extreme points have some additional properties that are often useful and that we do not develop here. For intuition in our setting, the reader might just think that the extreme point $x$ that we obtain has only $n+m$ coordinates that are non-zero, i.e. the solution is very sparse (note that in total there are $n \cdot m$ coordinates). Using the special structure of extreme points, [68] design a clever rounding scheme that satisfies Equation (1.10) with $\alpha=2$. Interestingly, this rounding scheme runs in polynomial time, which renders the whole proof constructive and although we were focusing only on estimating the optimum value, we get a constructive algorithm that runs in polynomial time.

While the techniques here are specific to the Unrelated Makespan problem, the general structure of the example illustrates a very successful paradigm in the area approximation algorithms: (1) write a relaxation of the problem, then (2) devise a rounding scheme of the fractional solution. Both steps are crucial and challenging: Step (1) might necessitate to introduce stronger constraints, and carrying out Step (2) successfully involves designing a new algorithm. There is no magic recipe that works for all problems, and overcoming each step is often an interesting mathematical problem on its own.

### 1.3.1 The use of LPs in this thesis

In all the results presented in this thesis, we will use a linear programming relaxation of the problem considered. Our results on the Santa Claus problem and the Matching Augmentation problem essentially fit into the classic framework described above: first write a relaxation of the problem, then round a fractional solution into an integral one. Interestingly, for the Matching Augmentation problem we crucially need that the LP solution has a sparse support, which is satisfied in general by extreme points of the linear program. Lastly, for our result on the Steiner Forest problem, the use of linear programming techniques is less obvious, but still important. Indeed, to compare the cost of the greedy solution to the optimum solution, we need to lower bound the cost of the latter. In order to achieve this, we rely on the weak duality theorem for linear programs, which states that the value of any feasible solution
to the dual linear program is a lower bound on the objective of the primal linear program. In our setting, the primal will be a relaxation of the Steiner Forest problem, hence a feasible dual solution gives us a lower bound on the value of any feasible solution to the Steiner Forest problem. This last example is interesting, as it shows that even if the algorithm itself does not use the LP relaxation, this relaxation is still useful in the analysis of the algorithm.

Part I
THE SANTA CLAUS PROBLEM

In this part, we focus on the Santa Claus problem, a central problem in scheduling and algorithmic fairness. In its full generality, we have a set $P$ of $m$ players, a set $R$ of $n$ resources, and each player $i$ is equipped with an arbitrary valuation function $f_{i}: S \subseteq R \mapsto \mathbb{R}_{+}$. We have to find a partition $\left(S_{1}, \ldots, S_{i}, \ldots, S_{m}\right)$ of the resource set such that

$$
\min _{i \in P} f_{i}\left(S_{i}\right)
$$

is maximized (note that a resource can be assigned to at most one player). This problem is often referred to as MaxMin Fair Allocation in the sense that the objective is to find the fairest solution, unlike other objectives like total welfare which maximizes the average happiness. Unfortunately, without any restrictions on the utility functions the Santa Claus problem becomes hopelessly difficult. Consider the following reduction from set packing. There are sets of resources $\left\{S_{1}, \ldots, S_{k}\right\}$ and all utility functions are equal and defined by $f_{i}(S)=1$ if $S_{j} \subseteq S$ for some $j$ and $f_{i}(S)=0$ otherwise. Deciding whether there are $m$ disjoint sets in $S_{1}, \ldots, S_{k}$ (a classical NP-hard problem, see [59]) is equivalent to deciding whether the optimum of the Santa Claus problem is non-zero. In particular, obtaining any bounded approximation ratio for Santa Claus in this case is NP-hard. This motivates natural restrictions on the valuation functions. This part gives new results for two special cases of the Santa Claus problem that we call MaxMin Arborescences, and the Restricted Submodular Santa Claus.

### 2.1 MAXMIN ARBORESCENCES

The most natural restriction is to assume that all the $f_{i}$ are linear valuation functions. We can rephrase this as follows: each resource $j$ has unrelated values $v_{i j}$ for each of the player $i$. The goal is to assign each resource $j$ to a player $\sigma(j)$ such that we maximize the utility of the least happy player, that is,

$$
\min _{i \in P} \sum_{j: \sigma(j)=i} v_{i j}
$$

The dual of the problem, where one has to minimize the maximum instead of maximizing the minimum is the problem of makespan minimization on unrelated parallel machines (this is exactly the Unrelated Makespan problem from the introduction of this thesis). Both variants form notoriously difficult open problems in approximation algorithms [13, 78, 85] and there is a common believe that the Santa Claus problem admits a constant approximation if and only if makespan minimization on unrelated machines admits a better-than-2 approximation [13]. Although formally no such reduction
is known, techniques often seem to transfer from one problem to the other, which gives an additional motivation to study this problem.

Bateni, Charikar, and Guruswami [15] identified as a central special case of the Santa Claus problem the restriction that for all values we have $v_{i j} \in$ $\{0,1, \infty\}$ and for each resource there is at most one player with $v_{i j}=\infty$ and for each player there is at most one resource with $v_{i j}=\infty$. This special case can be rephrased in a graph problem as follows. For each resource $j$ such that there is no player $i$ with $v_{i j}=\infty$, we create a sink vertex. For each player $i$ such that there is no resource $j$ with $v_{i j}=\infty$, we create a source vertex. For any pair of a player $i$ and resource $j$ such that $v_{i j}=\infty$, we create a vertex in our graph. Note that in this construction, source vertices correspond to a single player in the original instance, sink vertices correspond to a single resource, and other vertices correspond to a single player and a single resource. Next we create directed edges in our graph as follows. For every vertex $v$ that corresponds to a player $i$ in the original instance, we add a directed edge from $v$ to all other vertices in the graph that corresponds to a resource $j$ with $v_{i j}=1$. The Santa Claus problem in this special case can now be rephrased as follows. We are given a directed graph $G=(V, E)$, a set of sources $S \subseteq V$, and a set of sinks $T \subseteq V$. Our goal is to compute a set of vertex disjoint arborescences rooted in each of the sources $S$ (one for each source). The leaves of these arborescences must be at the sinks (i.e. only vertices in $T$ can be included as a leaf of an arborescence). For all inner vertices we must have an out-degree of $k$, where $k$ is to be maximized. This problem is called the MaxMin Arborescences problem.

Bateni et al. gave a $\max \left\{\operatorname{poly}(\log n), n^{\varepsilon}\right\}$-approximation in time $n^{O(1 / \varepsilon)}$, which is still the state-of-the-art for this problem. In particular, they obtain a quasi-polynomial time polylogarithmic approximation by setting $\varepsilon=\log \log n / \log n$. In a highly non-trivial way, the same approach was then generalized by Chakrabarty, Chuzhoy, and Khanna [21] to obtain the same result also for the Santa Claus problem. The only hardness known for the Santa Claus problem is that there is no better-than-2 approximation (see $[17,68])$. This still leaves a large gap in the understanding of both problems. In particular, the two most pressing questions are whether the polylogarithmic guarantee can also be achieved in polynomial time and whether a sublogarithmic approximation guarantee can be achieved.

While much progress on sublogarithmic approximations has been made on other special cases of the Santa Claus problem, most notable the restricted assignment case (on which we will elaborate later), these results are all based on a popular relaxation of the problem named the configuration $L P$. However, it is known that for the general Santa Claus problem the configuration LP has an polynomial integrality gap and hence these methods seem very unlikely to generalize. Similarly, the MaxMin Arborescences problem is a case where the integrality gap of the configuration LP is already high [15] and as such the algorithmic techniques need to be rethought. Given this and the fact that previous progress on the problem was quickly extended to the general Santa Claus problem, the MaxMin Arborescences problem seems to be an important piece of the puzzle towards the goal of understanding the approximability of the Santa Claus problem.

The rephrasing of this special case of the Santa Claus problem as an arborescence problem also uncovers an intriguing connexion to the Directed Steiner Tree problem. In this problem we are given an edge-weighted directed graph with a source and a set of sinks. The goal is to find an arborescence of minimal weight, which is rooted at the source and spans all sinks. It is quite remarkable that the state-of-the-art for this problem is very similar to ours: there is a $n^{\varepsilon}$-approximation algorithm in polynomial time for every fixed $\varepsilon>0$ and a polylogarithmic approximation algorithm in quasi-polynomial time [22]. Unlike our problem, it was shown that no sublogarithmic (in fact, no $\log ^{2-\varepsilon} n$ ) approximation exists [52].

### 2.1.1 Our results for MaxMin Arborescences

Our main result for MaxMin Arborescences is the following.
Theorem 2.1. There exists an $O\left((\log \log n)^{20}\right)$-approximation running in time $n^{O(\log n)}$ for the MaxMin Arborescences problem.

We remark that we did not try to optimize the exponent of 20 in the approximation guarantee. The main challenge was to get a sublogarithmic guarantee (i.e. an $o(\log n)$-approximation) and approach a constant factor. Indeed, in light of the $\Omega\left(\log ^{2-\varepsilon}(n)\right)$-hardness for directed Steiner Tree [52], it was not even clear if a sublogarithmic guarantee was possible for the MaxMin Arborescences problem. In fact, if one takes a closer look at the literature on directed Steiner Tree problems, we believe our result to be rather surprising. We elaborate now on this. Related to the directed Steiner Tree problem is also the (undirected) group Steiner Tree problem, which can be shown to be a special case. In this problem, we are given an undirected weighted graph and a list of groups that are subsets of vertices. The goal is to compute the cheapest set of edges that is connected and contains at least one vertex from each group. Here, there are more subtile connections to our problem: taking a closer look at the literature one can notice that the challenging instances in the group Steiner Tree problem have a similar structure to the challenging ones for the MaxMin Arborescence problem. More precisely they are layered graphs with $O(\log n)$ layers. Halperin et al. [51] show that the integrality gap of a natural LP relaxation of group Steiner Tree could be amplified from $\Theta(\log (n))$ on $O(1)$-layered instances to $\tilde{\Theta}\left(\log ^{2}(n)\right)$ on $\Omega(\log (n))$-layered graphs. This construction was later transformed to the hardness result by Halperin and Krauthgamer [52].

Before our result, it was quite plausible that such an amplification technique could also apply in the context of the MaxMin Arborescences problem. This would have shown how to amplify a $\Omega(1)$ gap on $O(1)$-layered instances to a $\tilde{\Omega}(\log (n))$ gap on $\Theta(\log (n))$-layered instances. In Appendix A.1, we adapt the construction of [51] to our setting. At first sight, the gap indeed seems to amplify and our construction shows that the previous rounding algorithms of $[15,21]$ cannot hope to get better than a $\tilde{\Omega}(\log (n))$-approximation. Fortunately, we notice that a simple (but crucial) pruning trick seems to resolve the issue for this specific construction. We note that the group Steiner Tree problem has a rich history and it would be very interesting to see if more techniques could be transferred to the MaxMin Arborescences problem.

Finally, we also show as a side result a hardness of approximation for the MaxMin Arborescences problem.

Theorem 2.2. For any $\varepsilon>0$, there is no $(\sqrt{e /(e-1)}-\varepsilon)$-approximation algorithm for the MaxMin Arborescences problem, unless $\mathrm{P}=\mathrm{NP}$.

It is worthwhile to mention that our hardness result is obtained by a reduction from the max- $k$-cover problem. Interestingly, the hardness of directed Steiner Tree problems in [52] is obtained by a reduction from set cover combined with the amplification technique mentioned above.

### 2.2 THE RESTRICTED SUBMODULAR SANTA CLAUS

We introduce here our second result on the Santa Claus problem. As explained in the beginning in this chapter, assuming no restrictions on the valuation functions $f_{i}$ makes the problem hopelessly difficult: no finite approximation is possible. However, we notice that most of the work on the Santa Claus problem has focused on the case where all valuation functions $f_{i}$ are linear, and there are some interesting special cases of linear functions where a constant factor is already known. A prominent such example is the so-called restricted assignment case.
In the restricted assignment case, the utility functions are defined by one linear function $f$ and a set of resources $\Gamma_{i}$ for each player $i$. Intuitively, player $i$ is interested in the resources $\Gamma_{i}$, whereas the other resources are worthless for him. The individual utility functions are then implicitly defined by $f_{i}(S)=f\left(S \cap \Gamma_{i}\right)$.
In a seminal work, Bansal and Srividenko [14] provide an $\tilde{O}(\log \log (m))^{1}$ approximation algorithm for this case using randomized rounding of the configuration LP. This was improved by Feige [39] to an $O(1)$-approximation. Further progress on the constant or the running time was made since then, see e.g. $[5,6,24,25,35,53,76]$. These works culminate in a $(4+\varepsilon)$-approximation in polynomial time and an 3.534-approximation that only gives an estimation of the optimum value in polynomial time (but does not return the corresponding assignment in polynomial time).

It is very natural to ask what is possible beyond the case of linear valuation functions. Two naturally arising properties of utility functions are monotonicity and submodularity, see for example the related submodular welfare problem [67, 83] where the goal is to maximize $\sum_{i} f_{i}\left(S_{i}\right)$. A function $f$ is monotone, if $f(S) \leqslant f(T)$ for all $S \subseteq T$. It is submodular, if $f(S \cup\{a\})-f(S) \geqslant f(T \cup\{a\})-f(T)$ for all $S \subseteq T$ and $a \notin T$. The latter is also known as the diminishing returns property in economics. A standard assumption on monotone submodular functions (used throughout this thesis) is that the value on the empty set is zero, i.e., $f(\emptyset)=0$. Goemans, Harvey, Iwata, and Mirrokni [46] first considered the Santa Claus problem with arbitrary monotone submodular utility functions as an application of their fundamental result on submodular functions. Together with the algorithm of [21] it implies an $O\left(n^{1 / 2+\varepsilon}\right)$-approximation in time $n^{O(1 / \varepsilon)}$. In the case

[^1]that the valuation functions are all equal, that is, $f_{i}(S)=f(S)$ for a monotone submodular function $f$, Krause, Rajagopal, Gupta, and Guestrin gave a constant approximation [65]. We also refer to their work for an application of this problem in sensor placement.

Our second result investigates the approximability of the Santa Claus with monotone submodular functions in the restricted assignment case. That is, all utility functions are defined by $f_{i}(S)=f\left(S \cap \Gamma_{i}\right)$, where $f$ is a monotone submodular function and $\Gamma_{i}$ is a subset of resources for each players $i$. Before our work, the state-of-the-art for this problem was the $O\left(n^{1 / 2+\varepsilon}\right)$ approximation algorithm mentioned above, since none of the previous results for the restricted assignment case with a linear utility function apply when the utility function becomes monotone submodular. In the rest of the thesis, we refer to this problem as the Restricted Submodular Santa Claus problem.

### 2.2.1 Our results for the Restricted Submodular Santa Claus

In Chapter 4, we show the following theorem, which essentially generalizes the result of Bansal and Srividenko [14] to the case of submodular functions.

Theorem 2.3. There is a randomized polynomial time $O(\log \log (n))$-approximation algorithm for the Restricted Submodular Santa Claus problem.

### 2.3 PROBABILISTIC LEMMAS

We finish the introduction of this part by stating here two crucial probabilistic lemmas, that we will use extensively for both results.

Lemma 2.4 (Chernoff's bound, see [31]). Let $X_{1}, \ldots, X_{n}$ be independent random variables that take value in $[0, a]$ for some fixed a. Let $S_{n}=\sum_{i=1}^{n} X_{i}$. Then we have, for any $\delta \geqslant 0$,

$$
\mathbb{P}\left[S_{n} \geqslant(1+\delta) \mathbb{E}\left[S_{n}\right]\right] \leqslant \exp \left(-\frac{\delta^{2} \mathbb{E}\left[S_{n}\right]}{(2+\delta) a}\right)
$$

Lemma 2.5 (Constructive Lovász Local Lemma, see [70]). Let $\mathcal{X}$ be a finite set of mutually independent random variables in a probability space. Let $\mathcal{A}$ be a finite set of events determined by these variables. For any $A \in \mathcal{A}$, let $\Gamma_{\mathcal{A}}(A)$ be the set of events $B \in \mathcal{A}$ such that $A$ and $B$ depend on at least one common variable. If there exists an assignment of reals $x: \mathcal{A} \mapsto(0,1)$ such that for all $A \in \mathcal{A}$,

$$
\mathbb{P}[A] \leqslant x(A) \cdot \prod_{B \in \Gamma_{\mathcal{A}}(A)}(1-x(B))
$$

then there exists an assignment of values to the variables $\mathcal{X}$ not triggering any of the events in $\mathcal{A}$. Moreover, there exists a randomized algorithm that finds such an assignment in expected time

$$
|\mathcal{X}| \cdot \sum_{A \in \mathcal{A}} \frac{x(A)}{1-x(A)}
$$

This chapter is devoted to the proof of Theorem 2.1 and Theorem 2.2. For Theorem 2.1, we start by a brief overview of the main ideas and intuition behind our result before giving the formal proof structure in Section 3.2, and then we provide the three crucial steps of the proof in Section 3.3, Section 3.4 , and Section 3.5. We close this chapter by a short proof of Theorem 2.2 in Section 3.6. All the results in this chapter are based on a joint work with Lars Rohwedder which has been accepted for publication at the Annual ACM Symposium on Theory of Computing (STOC '23). It is currently available on ArXiv [12].

### 3.1 OUR TECHNIQUES AND INTUITION

A crucial idea that goes back to previous work [15, 21] is to allow congestion in the solution. Generally, a vertex can only have one incoming edge in the solution, but we relax this constraint. We call the maximum number of times a vertex is used the congestion. The algorithms in [15, 21] employ randomized rounding to obtain a solution with polylogarithmic congestion. The congestion can then be translated into an approximation rate by relatively straight-forward arguments. This polylogarithmic congestion comes from a Chernoff bound that yields an inversely polynomial probability, which is then applied to all vertices with a union bound.

A new ingredient of our algorithm is the notion of local congestion. Roughly speaking, we first compute a solution, which still has polylogarithmic congestion, but when considering only a local part of the solution (say, vertices within a distance of $\ell=O(\log \log n)$ in the arborescences), then this local part needs to have much smaller congestion, i.e., poly $(\log \log n)$. In other words, if a vertex is used multiple times in the solution, then the occurrences should be far apart in the arborescences.

First, let us describe why it is plausible to be able to obtain such a guarantee. It is already known that by randomized rounding a polylogarithmic congestion can be achieved. The local congestion on the other hand is by definition a very local constraint and hence Lovász Local Lemma (LLL) is natural to employ. This is indeed our approach, although the details are challenging.

Next, we will explain how to arrive at a sublogarithmic congestion. The approach that we call top-to-bottom pruning is very blunt: slighly oversimplifying, we take a given solution (with poly $(\log n)$ congestion) and start at the sources of the arborescences. We throw away randomly a constant fraction (say, half) of their children. Then we move to the other children and recurse. Clearly, this decreases the approximation rate only insignificantly. However, in expectation the congestion at each vertex decreases drastically. If for example a vertex is at distance $d$ from the source of an arborescence, then the probability of it not being removed is only $1 / 2^{d}$. There is, however, a caveat here: suppose that the same vertex occurs in a (relaxed) arborescence
many times and all occurrences are very close to each other. Then there is a high positive correlation between the vertices not being removed, which forms a serious problem. Indeed, this is where the local congestion comes in. It essentially bounds the dependence of occurrences surviving. Again, this proof makes use of LLL, because it seems infeasible to try and make the probability of a vertex's congestion staying above poly $(\log \log n)$ small enough to apply a union bound.
Since both parts require LLL, it is crucial to bound the dependencies. However, if $k$ is large (say $\Omega(n)$ ), then even a local part of the arborescence contains many vertices. It then seems unlikely to be able to guarantee locally low congestion for every vertex. Roughly speaking, we will only guarantee the property for a large fraction of the vertices, so that we can make the probability inversely polynomial in $k$. All other vertices need to be removed from the solution and this is generally very dangerous: even if we remove only a small fraction of vertices, this can lead to other vertex removals becoming necessary, because they now have a low out-degree. If we are not careful, this can accelerate and corrupt the whole solution. We call this the bottom-to-top pruning and we formalize a condition, under which the damage to the solution can be controlled. This condition is then applied in both parts

This brings us to a discussion on our two pruning techniques. Intuitively, both approaches have complementary merits to each other. Bottom-to-top pruning makes it easy to maintain a low maximum congestion, but difficult to keep a good number of children for every vertex in the arborescence. On the contrary, top-to-bottom pruning makes it very easy to maintain a good number of children, but difficult to keep the maximum congestion under control. Our proof can be seen as a careful combination of those two techniques using LLL. The use of pruning makes the proof fairly involved and one might wonder if this could not be avoided. In particular, it is not clear if the analysis of the previous randomized rounding algorithm (see $[15,21])$ is tight or not. However, we argue that an $\Omega(\log n / \log \log n)$ factor seems unavoidable in previous works, and that it is not easily fixable. We now elaborate on this: In previous approaches as in ours, a crucial part of the algorithm is to solve the max-min degree arborescence problem on layered instances (that is, the sources are located in the first layer, and edges can exist only between vertices of consecutive layers). Previous works [15, 21] then solve these instances by rounding an LP relaxation of the problem (we use the same LP relaxation but with a different rounding). An intuitive randomized rounding that appears in previous works is roughly as follows. Assume that the LP says there exists a solution of value $k$. Then the source samples $k$ children with probability equal to the LP values. These selected children then select $k$ children each, again equal to the LP values (more precisely, values that correspond to conditioning on the previous selections). In that manner, we make progress layer by layer until reaching the last layer. This guarantees a maximum congestion that is at most polylogarithmic in the number of vertices, hence the polylogarithmic approximation ratio. At this point, one might be tempted to argue that very few vertices in our solution will have congestion $\Omega(\log n)$ and that they are not a serious problem. Unfortunately, we show in Appendix A. 1 an instance in which the above rounding results in a solution in which all the sinks selected in the arborescence have an expected
congestion of $\Omega(\log n / \log \log n)$. While this may seem counter-intuitive, recall that here we are implicitly conditioning by the fact of being selected in the solution. It is non-trivial to recover from this issue: For instance deletingin a bottom-to-top fashion-the vertices with high congestion will basically remove almost all the sinks which will corrupt the whole solution.
Lastly, there is one important issue that we have not mentioned so far: in a similar way that each vertex loses some fraction of its children in the rounding (compared to the LP relaxation), we would lose some fraction of the sources. This happens also in previous works [15, 21], who then only compute a solution for a $1 / \operatorname{poly}(\log n)$ fraction of the sources and repeat it for poly $(\log n)$ times to cover all sources. This again introduces a congestion of poly $(\log n)$, which seems difficult to avoid with the randomized rounding approach. First, we only design the randomized rounding to approximate single source instances so that we do not have to cope with this. Then we present a black-box reduction from many sources to one source, which uses a non-trivial machinery that is very different from the randomized rounding approach. Namely, this is by a local search framework, which has already seen a big application in the restricted assignment case of the Santa Claus problem, see related work. Usually the local search is analyzed against the configuration LP, which is not applicable here, so our way of applying it is quite different to previous works.

### 3.2 FORMAL PROOF STRUCTURE

We start with some simplifying assumptions. Let $k$ be the optimum of the given instance, which we guess via a binary search framework. If $k \leqslant$ poly $(\log \log n)$, then obtaining a $(1 / k)$-approximation is sufficient. This is easy to achieve: it is enough to find $|S|$ vertex-disjoint paths that connect each source to a sink, which can be done by a standard max-flow algorithm. Throughout the paper we assume without loss of generality that $k$ is at least poly $(\log \log n)$ with sufficiently large constants. Similarly, we assume that $n$ is larger than a sufficiently large constant. Whenever we divide $k$ by some term, for example, $k / \log \log n$, we would normally have to write $\lfloor k / \log \log n\rfloor$. All the divisors considered can be assumed to be much smaller than $k$, hence any loss due to rounding is only a small constant factor and therefore insignificant. We assume for simplicity that the divisors are always integral and the divisions have no remainder.

We describe solutions using a set of paths instead of directly as arborescences. This abstraction will become useful in the linear programming relaxation, but also in other parts throughout the paper. With a given arborescence, we associate the set of all paths from a source to some vertex in it. Let $p$ be a path from a source to some vertex $v$. Formally, $p$ is a tuple of vertices $\left(v_{0}, v_{1}, \ldots, v_{\ell}\right)$ where $v_{0}$ is a source, $v_{1}$ a vertex with an edge from $v_{0}$, and more generally $v_{i}$ is a vertex having an edge from $v_{\ell-1}$. We do not explicitly forbid circles, but the properties of a solution will imply that only simple paths can be used. Finally we say that $p$ is a closed path if its last vertex is a sink and an open path otherwise.

We will denote by $p \circ v$ the path $p$, to which we add the extra vertex $v$ at the end. We write $|p|$ for the length of path $p$. A path $q$ is said to the a descendant of $p$ if it contains $p$ as a prefix. In that case we call $p$ an ancestor of $q$. We also say that $p \circ v$ is a child of $p$ (who is then a parent of $p \circ v$ ). Within a given set of paths $P$, we denote the sets of children, parent, ancestors, and descendants of a path $p$ by $C_{P}(p), P_{P}(p), A_{P}(p)$, and $D_{P}(p)$ respectively. By $D_{P}(p, \ell)$ we denote the descendants $q \in D_{P}(p)$ with $|q|=|p|+\ell$. Furthermore, we write $D_{P}(p, \leqslant \ell)=\bigcup_{\ell^{\prime} \leqslant \ell} D_{P}\left(p, \ell^{\prime}\right)$. The set of paths that ends at a vertex $v$ is denoted by $I_{P}(v)$. If the set of paths $P$ is clear from the context, we may choose to omit the subscript for convenience.

The conditions for a set of paths $Q$ to form a degree- $k$ solution are the following:

1. $(s) \in Q$ for every source $s$,
2. $\left|I_{Q}(v)\right| \leqslant 1$ for every $v \in V$, and
3. $\left|C_{Q}(p)\right|=k$ for every path $p \in Q$ that is open.

We will also consider a relaxed version of (2), where we allow higher values than 1 . Then we call maximum over all $\left|I_{Q}(v)\right|$ the congestion of the solution. The motivation for looking at solutions with (low) congestion is that we can remove any congestion by reducing $k$ by the same factor.

Lemma 3.1. Let $Q$ be a degree- $k$ solution with congestion $K$. Then in polynomial time we can compute a degree- $k / K$ solution without congestion.

Proof. Consider a bipartite multigraph $\left(U \cup U^{\prime}, F\right)$ on two copies $U, U^{\prime}$ of the vertices in $V$ that appear in $Q$. The graph has an edge $(u, v) \in F$ for every path $p \in Q$ that ends in $u$ and for which $p \circ v \in C(p)$. Then the degree of a vertex $u^{\prime} \in U^{\prime}$ is exactly the congestion of this vertex. Similarly, the degree of a vertex $u \in U$ is $k$ times its congestion. We consider a fractional selection of edges $x_{e}=1 / K$ for each $e \in F$. Here, each vertex $u \in U$ has $\sum_{e \in \delta(u)} x_{e} \geqslant k / K$ and each vertex $u^{\prime} \in U^{\prime}$ has $\sum_{e \in \delta\left(u^{\prime}\right)} x_{e} \leqslant 1$, where $\delta(u)$ are the edges incident to $u$. As explained at the beginning of the section, we assume that $K$ divides $k$ and therefore $K / k$ is integer. By integrality of the bipartite matching polytope there exists also an integral vector $x^{\prime}$ that satisfies these bounds. This corresponds to a degree- $k / K$ solution without congestion.

### 3.2.1 Bounded depth solution

The following proof follows exactly the arguments of a similar statement in [15]. We repeat it here for convenience.

Lemma 3.2. Let $Q$ be a degree- $k$ solution. Then there exists a degree- $k / 2$ solution $Q^{\prime} \subseteq Q$ where $|p| \leqslant \log _{2} n$ for every $p \in Q^{\prime}$.

Proof. We iteratively derive $Q^{\prime}$ from $Q$. For every $d=1,2, \ldots, n$ we consider the paths $p \in Q$ with $|p|=d$. For $d=1$ we add all these paths to $Q^{\prime}$. Then given the paths of length $d$ in $Q^{\prime}$ we select for each of them the subset of $k / 2$ children in $Q$, which has the least number of descendants (assuming for
simplicity that 2 divides $k$ ). For every $d$ we will now bound the total number of descendants in $Q$ of paths of this length, namely

$$
n_{d}=\sum_{p \in Q^{\prime}:|p|=d}\left|D_{Q}(p)\right|
$$

Notice that the descendants are counted in set $Q$ and not $Q^{\prime}$. Clearly, we have $n_{1} \leqslant n$, since $|Q| \leqslant n$. Then since we remove the children with the largest number of descendants, we get $n_{i+1} \leqslant n_{i} / 2$ for every $i$. Thus, $n_{\log _{2} n}=0$.

### 3.2.2 Local congestion and layered instances

A crucial concept in our algorithm are solutions with locally low congestion. It will later be shown that such a solution suffices to derive a solution with (globally) low congestion.

Definition 3.3. Let $Q$ be a solution and $\ell \in \mathbb{N}$. We say that $Q$ has an $\ell$-local congestion of $L$, if for every $p \in Q \cup\{\emptyset\}$ and $v \in V$ we have

$$
\left|I_{D(p, \leqslant \ell)}(v)\right| \leqslant L
$$

For sake of clarity, let us note the special case of $p=\emptyset$ in the definition above. In this case, $D(p, \leqslant \ell)$ is simply the set of all paths in $P$ with length at most $\ell$ (potentially starting at different sources). Throughout the paper we use values of the order $\ell=O(\log \log n)$ and $L=\operatorname{poly}(\log \log n)$. The usefulness of local congestion is captured by the following lemma, which we will prove in Section 3.3.

Lemma 3.4. Let $Q$ be a degree- $k$ solution with $\ell$-local congestion of $L$ and global congestion $K$, where $\ell \geqslant \log _{2} K$ and $K \geqslant \log _{2} n$. Then we can compute in polynomial time a degree- $k /(8 \ell)$ solution $Q^{\prime} \subseteq Q$, which has global congestion at most

$$
O\left(\ell^{7} L\right)
$$

It remains to show how to compute a solution with low local congestion. The abstraction of local congestion and the lemma on bounded depth allows us to reduce at a low expense to instances in layered graphs.

Definition 3.5. A layered instance has layers $L_{0} \dot{\cup} L_{1} \dot{\cup} \ldots \dot{U} L_{h}=V$ such that $L_{0}$ consists of all sources and edges go only from one layer $L_{i}$ to the next layer $L_{i+1}$.

Lemma 3.6. In polynomial time we can construct a layered instance with $h=$ $\log _{2} n$ such that if there exists a degree- $k$ solution for the original instance, then there exists a degree- $k / 2$ solution for the layered instance. Further, any degree$k^{\prime}$ solution with $\ell$-local congestion $L$ and global congestion $K$ in the layered instance can in polynomial time be transformed to a degree- $k^{\prime}$ solution for the original instance with $\ell$-local congestion $\ell L$ and global congestion $K \log _{2} n$.

Proof. Let $L_{0}$ be the set of sources. Then for each $i=1,2, \ldots, \log n$ let $L_{i}$ be a copy of of all vertices $V$. We introduce an edge from $u \in L_{i}$ to $v \in L_{i+1}$ if $(u, v)$ is an edge in the original instance. The new set of sinks is the union of all sink vertices in all copies.

Consider now a degree- $k$ solution $Q$ for the original instance. By Lemma 3.2 there exists a degree- $k / 2$ solution $Q^{\prime}$ where each $p \in Q^{\prime}$ has $|p| \leqslant \log _{2} n$. For each such $p=\left(v_{1}, v_{2}, \ldots, v_{t}\right)$ we introduce a path $p^{\prime}=\left(v_{1}^{\prime}, v_{2}^{\prime}, \ldots, v_{t}^{\prime}\right)$ in the layered instance where $v_{i}^{\prime}$ is the copy of $v_{i}$ in $L_{i}$.

Now let $Q^{\prime}$ be a degree- $k^{\prime}$ solution with $\ell$-local congestion $L$ and global congestion $C$ in the layered instance. We transform $Q^{\prime}$ to a solution $Q$ for the original instance by replacing each path $p^{\prime}=\left(v_{1}^{\prime}, v_{2}^{\prime}, \ldots, v_{t}^{\prime}\right)$ by a path $p=\left(v_{1}, v_{2}, \ldots, v_{t}\right)$, where $v_{i}^{\prime}$ is a copy of $v_{i}$. Since there are only $\log _{2} n$ copies of each vertex, the global congestion increases by at most a factor of $\log _{2} n$. For the local congestion consider a path $p \in Q$. This path was derived from a path $p^{\prime} \in Q^{\prime}$. Notice that any path $q^{\prime} \in D\left(p^{\prime}, \leqslant \ell\right)$ ends in some vertex in $L_{|p|+1}, L_{|p|+2}, \ldots, L_{|p|+\ell}$. Thus, there are only $\ell$ copies of each vertex that $q^{\prime}$ can end in. Consequently, the $\ell$-local congestion can increase at most by a factor of $\ell$.

Lemma 3.7. Let $K=2^{11} \log ^{3} n, \ell=10 \log \log n$, and $L=2^{10} \ell^{2}$. Given a layered instance with a single source and optimum $k$, we can in time $n^{O(\log n)}$ compute a degree- $k /(64 \ell)$ solution with $\ell$-local congestion at most $L$ and global congestion at most $K$.

This lemma is proven in Section 3.4. The lemmas above would allow us already to obtain our main result for instances with a single source. To generalize to multiple sources we present a black-box reduction on layered instances. This is proved in Section 3.5.

Lemma 3.8. Suppose we have an $\alpha$-approximation for the max-min degree bounded arborescence problem on layered instances with a single source running in time $n^{O(\log n)}$. Then there is also a $256 \alpha$-approximation for layered graphs and an arbitrary number of sources running in time $n^{O(\log n)}$.

### 3.2.3 Connecting the dots

We complete the proof of the main theorem (Theorem 2.1) as follows.
Theorem 2.1. First we prove the theorem on layered instances with a single source. Let $k$ be the optimum of the given instance. Using Lemma 3.7 we can find a degree- $\Omega(k / \ell)$ solution with $\ell$-local congestion $L$ and global congestion $K$. Here $K=O\left(\log ^{3} n\right), \ell=O(\log \log n)$ with $2^{\ell} \geqslant K$, and $L=O\left(\ell^{2}\right)$. Next, we apply Lemma 3.4 to turn this into a degree- $\Omega\left(k / \ell^{2}\right)$ solution with global congestion at most $O\left(\ell^{7} L\right)=O\left(\ell^{9}\right)$. Using Lemma 3.1 we can convert this to a degree- $\Omega\left(k / \ell^{11}\right)$ solution without congestion. We therefore have an $O\left(\ell^{11}\right)$-approximation algorithm for a single source on layered graphs and Lemma 3.8 implies that we can extend this to an arbitrary number of sources.

We now turn our attention to instances that are not necessarily layered. Let again $k$ be the optimum. Using Lemma 3.6 we construct a layered instance that is guaranteed to contain a degree- $k / 2$ solution. Thus, with our algorithm
for layered instances we can obtain a degree- $\Omega\left(k / \ell^{11}\right)$ solution for it. This solution has $\ell$-local congestion at most 1 and global congestion at most 1. Using Lemma 3.6 we can construct a degree- $\Omega\left(k / \ell^{11}\right)$ solution for the original (non-layered) instance with $\ell$-local congestion at most $\ell$ and global congestion at most $\log n$. Using again Lemma 3.4 we obtain a degree $-\Omega\left(k / \ell^{12}\right)$ solution with global congestion at most $O\left(\ell^{8}\right)$. Finally, applying Lemma 3.1 we obtain a degree $-\Omega\left(k / \ell^{20}\right)$ solution without congestion. In particular, our approximation ratio is

$$
O\left(\ell^{20}\right)=O\left((\log \log n)^{20}\right)
$$

Finally, the overall running time in clearly dominated by the use of Lemmas 3.7 and 3.8 , which is $n^{O(\log n)}$.

### 3.2.4 Bottom-to-top pruning

In this subsection we will describe a method of pruning that is used in the proofs of Lemmas 3.4 and 3.7. Before stating the result we need, we illustrate why this pruning technique can be tricky to use and what are its limitations.

The name "bottom-to-top pruning" refers to the following intuitive strategy. Suppose we are given a degree- $k$ solution with some congestion. A natural approach is to remove paths that have high congestion at their endpoint in the hope to decrease the maximum congestion. However, this can be dangerous. Indeed, removing paths may remove children of some other paths and therefore forcing us to remove them as well. In general, even a very small amount of removals can lead to the whole solution getting corrupted, that is, ultimately sources may need to be removed as well. In fact, a simple calculation shows that if we allow an adversary to delete some fraction $\beta$ of the paths of length $d$ in our solution and that we try to maintain a solution of degree $k / \alpha$, then the fraction of children of the sources that we might be forced to delete becomes roughly equal to $\beta e^{d / \alpha} \leqslant \beta e^{\log n / \alpha}$ (recall that at a constant loss we can assume that the depth of the solution is at most $\log n$ ). In particular if we aim at a sublogarithmic approximation ratio (i.e. $\alpha=o(\log n)$ ), the bottom-to-top pruning becomes very sensitive to a small number of deletions.

This pruning technique will be useful nonetheless, and the following lemma states a condition under which we can perform bottom-to-top pruning while keeping the damage under control.

Lemma 3.9. Let $Q$ be a degree-k solution. Let $R \subseteq Q$ be a set of paths that is supposed to be removed. Let $\ell \geqslant 2$ such that for every $p \in Q \backslash R$ we have that at most $k^{\ell} /(8 \ell)^{2}$ many descendants $q \in D(p, \ell)$ with $q \in R$. Then we can compute in polynomial time a solution $Q^{\prime} \subseteq Q \backslash R$ such that

1. $Q^{\prime}$ is a degree- $k / 2 \ell$ solution and
2. we have $(s) \in Q^{\prime}$ for every source $s$, such that for any distance $\ell^{\prime} \leqslant \ell$ there are at most $k^{\ell^{\prime}} /(8 \ell)$ many $q \in D\left((s), \ell^{\prime}\right)$ with $q \in R$.

Proof. We assume that no path $p \in R$ has a descendant also in $R$. This is without loss of generality, since removing the former implies that the latter
will be removed, and omitting the latter from $R$ still keeps the premise of the lemma valid. In particular, this assumption allows us to assert that also paths in $R$ satisfy the bound on the number of descendants in $R$.
We prune the solution from longest paths to shortest paths: We remove a path if it is in $R$ or if more than $(1-1 / 2 \ell) k$ many of its children were removed. Then we prove a stronger variant of (2) inductively, namely, that any path of length $1,1+\ell, 1+2 \ell$, etc. satisfies the implication (or an ancestor of it is removed). Let $p$ be a path with $|p|=1+t \cdot \ell$ that satisfies the premise of (2), but is not necessarily a singleton. Further, assume that all paths of length $1+(t+1) \ell$ satisfy the implication of (2). Let $\ell^{\prime} \leqslant \ell$. Each of the distance- $\ell^{\prime}$ descendants of $p$ has at most $(8 \ell)^{-2} k^{\ell}$ many distance- $\ell$ descendants in $R$. Consequently, $p$ has at most

$$
k^{\ell^{\prime}} \cdot(8 \ell)^{-2} k^{\ell}
$$

distance $\left(\ell+\ell^{\prime}\right)$-descendants in $R$. Thus, at most $(8 \ell)^{-1} k^{\ell}$ many of $p$ 's distance- $\ell$ descendants have more than $(8 \ell)^{-1} k^{\ell^{\prime}}$ distance- $\ell^{\prime}$ descendants belonging to $R$. Summing over all values of $\ell^{\prime}$, we have that at most $1 / 8 \cdot k^{\ell}$ many distance- $\ell$ descendants of $p$ do not satisfy the premise of (2). Next, let us show in a second induction that for every $\ell^{\prime} \leqslant \ell$, of the distance- $\ell^{\prime}$ descendants of $p$ at most

$$
\frac{1}{8}\left(1+\frac{2}{\ell}\right)^{\ell-\ell^{\prime}+1} k^{\ell^{\prime}}
$$

many are removed. For the base case we sum the distance- $\ell$ descendants that do not satisfy the premise of (2) and the descendants that are themselves in $R$, which together are at most

$$
\frac{1}{8} k^{\ell}+\frac{1}{(8 \ell)^{2}} k^{\ell} \leqslant \frac{1}{8}\left(1+\frac{2}{\ell}\right)^{\ell-\ell+1} k^{\ell}
$$

Now assume that we removed at most $1 / 8 \cdot(1+2 / \ell)^{\ell-\ell^{\prime}} k^{\ell^{\prime}+1}$ paths from the distance- $\left(\ell^{\prime}+1\right)$ descendants. For each distance- $\ell^{\prime}$ descendants that we remove because of few remaining children, there are $(1-1 /(2 \ell)) k$ many distance- $\left(\ell^{\prime}+1\right)$ descendants that we removed. The bound from this and the number of distance- $\ell^{\prime}$ descendants in $R$ lets us bound the number of distance- $\ell^{\prime}$ descendants that we remove by

$$
\frac{1}{k}\left(1-\frac{1}{2 \ell}\right)^{-1} \cdot \frac{1}{8}\left(1+\frac{2}{\ell}\right)^{\ell-\ell^{\prime}} k^{\ell^{\prime}+1}+\frac{1}{8 \ell} k^{\ell^{\prime}} \leqslant \frac{1}{8}\left(1+\frac{2}{\ell}\right)^{\ell-\ell^{\prime}+1} k^{\ell^{\prime}}
$$

It follows with $\ell^{\prime}=1$ that we remove at most $1 / 8 \cdot(1+2 / \ell)^{\ell} k \leqslant e^{2} / 8 \cdot$ $k<(1-1 /(2 \ell)) k$ children of $p$. Hence, $p$ is not removed itself by the procedure.

### 3.3 LOCAL TO GLOBAL CONGESTION

This section is to prove Lemma 3.4. Let $Q$ be a degree- $k$ solution with $\ell$-local congestion $L$ and global congestion $K$. We partition $Q$ by length: let $Q_{i}$
be the set of paths $p \in Q$ with $|p|=i$. Further, we split the paths into $\ell$ groups $G_{1}, G_{2}, \ldots, G_{\ell}$, where $G_{j}=Q_{j} \cup Q_{j+\ell} \cup Q_{j+2 \ell} \cup \cdots$. For some $p \in G_{j}$ we write $G(p)=G_{j}$. Roughly speaking, we proceed as follows. We sample from each $G_{\ell}$ half of the paths and throw away all others (including their descendants). Then we move to $G_{\ell-1}$ and do the same. We continue until $G_{1}$ and then repeat the same a second time, stopping afterwards. The sampling is done in a way that guarantees that each path retains a quarter of its children at the end. We prove with Lovász Local Lemma that in each step we can reduce the congestion significantly. Since it seems unclear how to argue directly about the worst case congestion, we will argue about the congestion aggregated over many paths, which we will formalize next.

Consider a path $p \in Q$ and its close descendants in $D(p, \leqslant \ell)$. Recall, that $D(p, \leqslant \ell)$ contains all descendants of length at most $|p|+\ell$. Intuitively, if many of the direct children of $p$ have high congestion (more precisely, the vertex that they end in), this is bad for $p$ as well: if they have high congestion, we may not be able to keep many of these children for $p$, which means we might not be able to include $p$ itself in the solution. Let $\operatorname{cong}_{G}(p)$ be the congestion of the last vertex in $p$, but restricted to paths in the group $G(p)$. In other words,

$$
\operatorname{cong}_{G}(p)=\mid\{q \in G(p) \mid q \text { ends in the same vertex as } p\} \mid .
$$

The restriction to other paths in $G(p)$ is only for technical reasons and almost at no cost: if we can achieve that every vertex is used by only few paths in each $G_{j}$ (i.e., $\operatorname{cong}_{G}(p)$ is small for all $p \in Q$ ), the overall congestion can only be worse by a factor $\ell$. An important quantity in the following will be the total congestion of descendants $D\left(p, \ell^{\prime}\right)$ of $p$ at some distance $\ell^{\prime}<\ell$, that is,

$$
\begin{equation*}
\operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right)\right)=\sum_{q \in D\left(p, \ell^{\prime}\right)} \operatorname{cong}_{G}(q) \tag{3.1}
\end{equation*}
$$

Since during the procedure the number of children may differ between groups $G_{j}$, we will use $k\left(G_{j}\right)$ to describe the current number of children for every open path in $G_{j}$. Our intermediate goal will be to bound the totals (3.1) for some $p \in G_{j}$ in terms of $k\left(p, \ell^{\prime}\right)=\prod_{j^{\prime}=j}^{j+\ell^{\prime}-1} k\left(G_{j}\right)$ (an upper bound on $\left.\left|D\left(p, \ell^{\prime}\right)\right|\right)$.

Notice that intially (3.1) can be at most $K \cdot k\left(p, \ell^{\prime}\right)$. When sampling down the paths in $G_{j}$, we want to show that this reduces (3.1) significantly for paths $p \in G_{j}$ and all $\ell^{\prime}<\ell$. This is captured in the following lemma.

Lemma 3.10. Assume we are given a solution that has an $\ell$-local congestion of at most $L$ and global congestion of at most $K$, where $\ell \geqslant \log K$ and $K \geqslant \log n$. From the paths in $G_{j}$ form pairs where each pair shares the same parent. Then select i.i.d. one path from each pair and remove it and all its descendants. Let $\operatorname{cong}_{G}$ and $\operatorname{cong}_{G}^{\prime}$ be the congestion count before and after the removal and similarly $k$ and $k^{\prime}$ the children count. Then we have with positive probability for every remaining $p \in G_{j}$ and $\ell^{\prime}<\ell$ that

$$
\frac{\operatorname{cong}_{G}^{\prime}\left(D\left(p, \ell^{\prime}\right)\right)}{k^{\prime}\left(p, \ell^{\prime}\right)} \leqslant c \ell^{4} L+\frac{1}{2}\left(1+\frac{1}{\ell}\right) \frac{\operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right)\right)}{k\left(p, \ell^{\prime}\right)}
$$

where $c$ is a fixed constant. Furthermore, we can obtain such a sampling in expected polynomial time.
Proof. We can rewrite

$$
\operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right)\right)=\sum_{q \in G(p)} \operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right),
$$

where $\operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right)$ is the number of pairs $p^{\prime} \in D\left(p, \ell^{\prime}\right), q^{\prime} \in$ $D\left(q, \ell^{\prime}\right)$ that end in the same vertex. We remove every term in the sum with probability $1 / 2$, so in expectation the sum will reduce by $1 / 2$. Furthermore, each term $\operatorname{cong}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right)$ is bounded by $L \cdot\left|D\left(p, \ell^{\prime}\right)\right| \leqslant L \cdot k\left(p, \ell^{\prime}\right)$, because we have low local congestion. This will give us good concentration. Notice also that $k\left(p, \ell^{\prime}\right)=k^{\prime}\left(p, \ell^{\prime}\right)$.
We now group the terms in the sum by their size. For $p \in G_{j}$ let $G\left(p, \ell^{\prime}, t\right)$ be the set of $q \in G(p)=G_{j}$ with $\operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right) \in\left[L \cdot k\left(p, \ell^{\prime}\right)\right.$. $\left.2^{-(t+1)}, L \cdot k\left(p, \ell^{\prime}\right) \cdot 2^{-t}\right)$. Let $Q^{\prime}$ be the set of paths remaining after sampling down (without taking into account those that are removed recursively). Let $c$ be a large constant to be specified later. Depending on whether $t$ is small or large, we define bad events $\mathcal{B}\left(p, \ell^{\prime}, t\right)$ for each $p \in G_{j}$ as

$$
\begin{aligned}
& \sum_{q \in G\left(p, \ell^{\prime}, t\right) \cap Q^{\prime}} \operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right) \\
& >24 c \ell^{3} L \cdot k\left(p, \ell^{\prime}\right)+\frac{1}{2}\left(1+\frac{1}{\ell}\right) \sum_{q \in G\left(p, \ell^{\prime}, t\right)} \operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right) \\
& \text { if } t \leqslant 2 \ell, \\
& \sum_{q \in G\left(p, \ell^{\prime}, t\right) \cap Q^{\prime}} \operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right) \\
& >24 c \ell^{3} L \cdot \frac{1}{K} k\left(p, \ell^{\prime}\right)+\frac{1}{2}\left(1+\frac{1}{\ell}\right) \sum_{q \in G\left(p, \ell^{\prime}, t\right)} \operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right) \\
& \text { if } t>2 \ell .
\end{aligned}
$$

Since from the experiment the total congestion cannot increase, we have a probability of 0 for all bad events where $\operatorname{cong}_{G\left(p, \ell^{\prime}, t\right)}\left(D\left(p, \ell^{\prime}\right)\right):=$

$$
\sum_{q \in G\left(p, \ell^{\prime}, t\right)} \operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right), D\left(q, \ell^{\prime}\right)\right) \leqslant \begin{cases}12 c \ell^{3} L \cdot k\left(p, \ell^{\prime}\right) & \text { if } t \leqslant 2 \ell \\ 12 c \ell^{3} L \cdot \frac{1}{K} k\left(p, \ell^{\prime}\right) & \text { otherwise }\end{cases}
$$

For the remaining bad events, we will now derive an upper bound on the probabilities. This holds trivially also for the zero probability events. From Chernoff's bound we get

$$
\begin{aligned}
\mathbb{P}\left[\mathcal{B}\left(p, \ell^{\prime}, t\right)\right] & \leqslant \exp \left(-\frac{\left(1 / \ell^{2}\right) \cdot 1 / 2 \cdot \operatorname{cong}_{G\left(p, \ell^{\prime}, t\right)}\left(D\left(p, \ell^{\prime}\right)\right)}{(2+1 / \ell) \cdot 2^{-t} L \cdot k\left(p, \ell^{\prime}\right)}\right) \\
& \leqslant \exp \left(-\frac{\operatorname{cong}_{G\left(p, \ell^{\prime}, t\right)}\left(D\left(p, \ell^{\prime}\right)\right)}{6 \ell^{2} \cdot 2^{-t} L \cdot k\left(p, \ell^{\prime}\right)}\right) \\
& \leqslant \begin{cases}\exp \left(-2 c \ell 2^{t}\right) & \text { if } t \leqslant 2 \ell, \\
\exp \left(-2 c \ell 2^{t} / K\right) \leqslant n^{-10} & \text { otherwise. }\end{cases}
\end{aligned}
$$

Towards applying LLL, we set the values of the bad events as

$$
x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right)= \begin{cases}\exp \left(-c \ell 2^{t}\right) & \text { if } t \leqslant 2 \ell \\ n^{-5} & \text { otherwise }\end{cases}
$$

The experiment involves binary variables $V$. An event $\mathcal{B}\left(p, \ell^{\prime}, t\right)$ depends on at most $K \cdot 2^{t+1}$ many variables. A variable $V$ influences at most $2 K \ell \cdot 2^{t+1}$ type- $t$ bad events. Notice that $2 K \ell \cdot 2^{t+1} \leqslant \exp \left(c \ell 2^{t}\right) /\left(4 K^{2}\right)$ for $c$ sufficiently large (since $\ell \geqslant \log K)$. Let $\Gamma_{t}(V)$ be the set of all these events for a specific variable $V$ and a specific value of $t$. Then

$$
\prod_{B \in \Gamma_{t}(V)}(1-x(B)) \geqslant\left(1-\exp \left(-c \ell 2^{t}\right)\right)^{2 K \ell 2^{t+1}} \geqslant 1-\frac{1}{4 K^{2}}
$$

Notice that $t$ can range only from 0 to $\log n$. Furthermore, since every vertex has congestion at most $K$, the number of paths in $Q$ is at most $K n$. Together, we can upper bound the total number of events by $K n \ell \log n \leqslant n^{5}$. Thus, for $t \leqslant 2 \ell$ and $c$ sufficiently large it holds that

$$
\begin{aligned}
\mathbb{P}\left[\mathcal{B}\left(p, \ell^{\prime}, t\right)\right] & \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot \exp \left(-c \not 2^{t}\right) \\
& \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot\left(1-\frac{1}{e}\right)^{4 \ell \cdot 2^{t}} \\
& \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot \prod_{t^{\prime}=0}^{\ell}\left(1-\frac{1}{4 K^{2}}\right)^{K \cdot 2^{t+1}} \cdot\left(1-\frac{1}{n^{5}}\right)^{n^{5}} \\
& \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot \prod_{B \in \Gamma\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right)}(1-x(B))
\end{aligned}
$$

Similarly, for $t>2 \ell$ we have

$$
\begin{aligned}
\mathbb{P}\left[\mathcal{B}\left(p, \ell^{\prime}, t\right)\right] & \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right) \cdot \exp \left(-c \not 2^{t} / K\right)\right. \\
& \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot\left(1-\frac{1}{e}\right)^{4 \ell \cdot 2^{t} / K} \\
& \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot \prod_{t^{\prime}=0}^{\ell}\left(1-\frac{1}{4 K^{2}}\right)^{K \cdot 2^{t+1}} \cdot\left(1-\frac{1}{n^{5}}\right)^{n^{5}} \\
& \leqslant x\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right) \cdot \prod_{B \in \Gamma\left(\mathcal{B}\left(p, \ell^{\prime}, t\right)\right)}(1-x(B))
\end{aligned}
$$

Hence, by LLL we have with positive probability that none of the bad events occur. If none of them occur, then by summing up bounds fixed $p$ and $\ell^{\prime}$ we get

$$
\begin{aligned}
& \operatorname{cong}^{\prime}\left(D\left(p, \ell^{\prime}\right)\right) \leqslant 2 \ell \cdot 24 c \ell^{3} L \cdot k\left(p, \ell^{\prime}\right) \\
&+\log n \cdot 24 c \ell^{3} L \cdot k\left(p, \ell^{\prime}\right) / K \\
&+\frac{1}{2}\left(1+\frac{1}{\ell}\right) \operatorname{cong}\left(D\left(p, \ell^{\prime}\right)\right) \\
& \leqslant 100 c \ell^{4} L \cdot k\left(p, \ell^{\prime}\right)+\frac{1}{2}\left(1+\frac{1}{\ell}\right) \operatorname{cong}\left(D\left(p, \ell^{\prime}\right)\right)
\end{aligned}
$$

Lemma 3.11. Consider a successful run of the random experiment in Lemma 3.10 where we sample down the paths in $G_{j}$ and satisfy the stated inequalities. For each path $p$ (potentially not in $G_{j}$ ) and every $\ell^{\prime} \in\{1,2, \ldots, \ell\}$ we have

$$
\frac{\operatorname{cong}_{G}^{\prime}\left(D\left(p, \ell^{\prime}\right)\right)}{k^{\prime}\left(p, \ell^{\prime}\right)} \leqslant c \ell^{4} L+\left(1+\frac{1}{\ell}\right) \frac{\operatorname{cong}_{G}\left(D\left(p, \ell^{\prime}\right)\right)}{k\left(p, \ell^{\prime}\right)}
$$

Here $c$ is the constant from Lemma 3.10.
Lemma 3.10 only shows that the average congestion reduces for descendants of paths in the group $G_{j}$, where we sample down. Conversely, Lemma 3.11 says that for all other groups it does not increase significantly.

Proof. Let $G_{j^{\prime}}=G(p)$. We can assume without loss of generality that $j^{\prime}<$ $j \leqslant j^{\prime}+\ell^{\prime}$ (modulo $\ell$ ), since the congestion can only decrease and $k\left(p, \ell^{\prime}\right)$ in the other case would not change. Let $\ell^{\prime \prime}=j-j^{\prime}$ (modulo $\ell$ ). Further, let $D^{\prime}\left(p, \ell^{\prime}\right)$ be the distance- $\ell^{\prime}$ descendants of $p$ after sampling down $G_{j}$ and $D\left(p, \ell^{\prime}\right)$ before it. Then $D^{\prime}\left(p, \ell^{\prime \prime}\right)$ contains half of the elements $D\left(p, \ell^{\prime \prime}\right)$. Thus,

$$
\begin{aligned}
\frac{\operatorname{cong}_{G}^{\prime}\left(p, \ell^{\prime}\right)}{k^{\prime}\left(p, \ell^{\prime}\right)} & =\frac{1}{k^{\prime}\left(p, \ell^{\prime}\right)} \sum_{q \in D^{\prime}\left(p, \ell^{\prime \prime}\right)} \operatorname{cong}_{G}^{\prime}\left(q, \ell^{\prime}-\ell^{\prime \prime}\right) \\
& =\frac{2}{k\left(p, \ell^{\prime}\right)} \sum_{q \in D^{\prime}\left(p, \ell^{\prime \prime}\right)} \operatorname{cong}_{G}^{\prime}\left(q, \ell^{\prime}-\ell^{\prime \prime}\right) \\
& \leqslant \frac{2}{k\left(p, \ell^{\prime}\right)} \sum_{q \in D^{\prime}\left(p, \ell^{\prime \prime}\right)}\left[c \ell^{4} L+\frac{1}{2}\left(1+\frac{1}{\ell}\right) \operatorname{cong}_{G}\left(q, \ell^{\prime}-\ell^{\prime \prime}\right)\right] \\
& \leqslant c \ell^{4} L+\left(1+\frac{1}{\ell}\right) \operatorname{cong}_{G}\left(p, \ell^{\prime}\right)
\end{aligned}
$$

Lemma 3.12. Given a degree- $k$ solution $Q$ with $\ell$-local congestion $L$ and global congestion $K$, we can compute a degree- $k / 4$ solution $Q^{\prime} \subseteq Q$ with

$$
\operatorname{cong}_{G}(p, \ell) \leqslant 3 c \ell^{4} L\left(\frac{k}{4}\right)^{\ell}+\left(1+\frac{1}{e}\right)^{2} \frac{K}{2^{\ell}}
$$

for all $p \in Q^{\prime}$. Here $c$ is the constant from Lemma 3.10.
Proof. We will perform the sampling from Lemma 3.10 for $G_{\ell}, G_{\ell-1}, \ldots, G_{1}$ and then again the same a second time. The reason is that we want that for every group $G_{j}$ that $G_{j}, G_{j-1}, \ldots, G_{j-\ell}$ (index modulo $\ell$ ) are down-sampled at least once in this order.

Let $G_{j}=G(p)$ and consider the first time that we sample down $G_{j+\ell-1}$ (index taken modulo $\ell$ ). Let $D\left(p, \ell^{\prime}\right)$ be the distance- $\ell^{\prime}$ descendants of $p$ before this sampling. Then for each $q \in D(p, \ell-1)$ we have

$$
\frac{\operatorname{cong}_{G}(D(q, 1))}{k(q, 1)} \leqslant K
$$

This is due to the fact that sampling down cannot increase the congestion on any vertex and initially all vertices have congestion at most $K$. After sampling according to Lemma 3.10 we have

$$
\frac{\operatorname{cong}_{G}(q, 1)}{k(q, 1)} \leqslant c \ell^{4} L+\left(1+\frac{1}{\ell}\right) \frac{K}{2} .
$$

In the next step we are sampling down $G_{j+\ell-2}$. We have for each $q \in$ $D(p, \ell-2)$ that

$$
\frac{\operatorname{cong}_{G}(D(q, 2))}{k(q, 2))}=\frac{1}{k(q, 1)} \sum_{q^{\prime} \in C(q)} \frac{\operatorname{cong}_{G}\left(q^{\prime}, 1\right)}{k\left(q^{\prime}, 1\right)} \leqslant c \ell^{4} L+\left(1+\frac{1}{\ell}\right) \frac{K}{2}
$$

Thus, after sampling

$$
\begin{aligned}
\frac{\operatorname{cong}_{G}(D(q, 2))}{k(q, 2)} \leqslant c \ell^{4} L+\frac{1}{2}\left(1+\frac{1}{\ell}\right) & c \ell^{4} L+\left(1+\frac{1}{\ell}\right)^{2} \frac{K}{4} \\
& \leqslant 2 c \ell^{4} L+\left(1+\frac{1}{\ell}\right)^{2} \frac{K}{4}
\end{aligned}
$$

Continuing this argument, after we sample down $G_{j}$ we have

$$
\frac{\operatorname{cong}_{G}(p, \ell)}{k(p, \ell)} \leqslant 3 c \ell^{4} L+\left(1+\frac{1}{\ell}\right)^{\ell} \frac{K}{2^{\ell}}
$$

After $G_{j}$ there may be at most $\ell$ more steps of sampling down, after which we finally have

$$
\frac{\operatorname{cong}_{G}(p, \ell)}{k(p, \ell)} \leqslant 3 c \ell^{4} L+\left(1+\frac{1}{\ell}\right)^{2 \ell} \frac{K}{2^{\ell}} \leqslant 3 c \ell^{4} L+\left(1+\frac{1}{e}\right)^{2} \frac{K}{2^{\ell}}
$$

We will now conclude the proof of Lemma 3.4. Using the previous lemma, we obtain a degree- $k / 4$ solution $Q^{\prime}$. Let $A:=3 c \ell^{4} L+(1+1 / e)^{2} K / 2^{\ell}$ be the upper bound on average distance- $\ell$ congestion. Assuming without loss of generality that $c$ is sufficiently large, we have that $A \leqslant 3 c \ell^{4} L$. Let $R$ be the set of all paths $p$ with $\operatorname{cong}_{G}(p)>16 A \ell^{2}$. We remove these paths using Lemma 3.9. We use the property that each path has at most a $(8 \ell)^{-2}(k / 4)^{\ell}$ many of its ancestors at distance $\ell$ in $R$. This follows directly from the bounded average congestion. The lemma implies that we can remove those high congestion paths and still keep a solution where each remaining path has $k /(4 \cdot 2 \ell)=k / 8 \ell$ children. Since we start with a $\ell$-local congestion of at most $L \leqslant 16 A \ell^{2}$ and this cannot be increased by only removing paths, we have that none of the paths of length at most $\ell$ are removed and thus all sources satisfy the premise of (2) of Lemma 3.9 and remain in the solution. Indeed, the value of cong ${ }_{G}(p)$ is now bounded by $16 A \ell^{2}$ for all remaining paths. We recall that the actual congestion is at most a factor $\ell$ higher than $\max _{p} \operatorname{cong}_{G}(p)$, that is,

$$
16 A \ell^{3} \leqslant 64 c \ell^{7} L
$$

The goal of this section is to prove Lemma 3.7. We recall that we are in a layered graph with vertices partitioned into layers $L_{0}, L_{1}, \ldots, L_{h}$ where $h=\log n$ and a single source $s$. The source $s$ belongs to layer $L_{0}$ and edges can only go from vertices in some layer $L_{i}$ to the next layer $L_{i+1}$.
In the following we will extensively argue about paths that start in the source. For the remainder of the section every path $p$ that we consider is implicitly assumed to start at the source and then traverses (a prefix of) the layers one by one. Slightly abusing notation, we sometimes use $L_{i}$ also to denote the set of paths ending in a vertex of $L_{i}$, that is,

$$
\bigcup_{v \in L_{i}} I(v),
$$

and $T$ to describe the set of closed paths (recall that those are the paths that end at a sink). Let $P$ refer to the set of all possible paths and notice that by virtue of the layers we have that $|P| \leqslant n^{h} \leqslant n^{\log n+1}$. We will now describe a linear programming relaxation, which goes back to Bateni et al. [15]. The intuition behind the linear program is to select paths similarly to the way we describe solutions, see Section 3.2. We have a variable $x(p)$ for each path $p$ that in an integral solution takes value 1 if the path $p$ is contained in an arborescence and 0 otherwise.

$$
\begin{align*}
\sum_{q \in C(p)} x(q) & =k \cdot x(p) & \forall p \in P \backslash T  \tag{3.2}\\
\sum_{q \in I(v) \cap D(p)} x(q) & \leqslant x(p) & \forall p \in P, v \in V  \tag{3.3}\\
x((s)) & =1 &  \tag{3.4}\\
x & \geqslant 0 & \tag{3.5}
\end{align*}
$$

Here we assume that $k$ is the highest value for which the linear program is feasible, obtained using a standard binary search framework. Moreover, we assume that $k \geqslant 2^{10}(\log \log n)^{8}$ in the rest of the section. This is at little cost, since a $1 / k$-approximation is easy to obtain (see the proof of Theorem 2.1) and is already sufficient for our purposes. The first two types of constraints describe that each open path has many children and each vertex has low congestion (in fact, no congestion). Constraint (3.3) comes from a lift-and-project idea. For integral solutions it would be implied by the other constraints, but without it, there could be situations with continuous variables, where for example we take a path $p$ with only value $1 / k$ and then a single child of $q$ with value 1 . Such situations easily lead to large integrality gaps, which we can avoid by this constraint.

Since the graph has $h+1$ many layers, this linear program has $n^{O(h)}$ variables and constraints and therefore can be solved in time $n^{O(h)}$. We refer to this relaxation as the path $L P$. In order to prove Lemma 3.7, we will design a rounding scheme. Before getting to the main part of the proof, we will first preprocess the fractional solution to sparsify its support.

### 3.4.1 Preprocessing the LP solution

Our first step is to sparsify the path LP solution $x$ to get another sparser solution (i.e., with a limited number of non-zero entries). For ease of notation, we might need to take several times a copy of the same path $p \in P$. We emphasize here that two copies of the same path are different objects. To make this clear, we will now have a multiset $P^{\prime}$ of paths but we will slightly change the parent/child relationship between paths. Precisely, for any path $q^{\prime} \in P^{\prime}$ is assigned as child to a unique copy $p^{\prime} \in P^{\prime}$ such that $q$ was a child of $p$ in the set $P$. With this slight twist, all the ancestors/descendants relationships extend to multisets in the natural way. For instance, we will denote by $D_{P^{\prime}}(p)$ the set of descendants of $p$ in the multiset $P^{\prime}$. Again, the set of closed path in $P^{\prime}$ will be denoted by $T^{\prime}$ and $s$ refers to the source. We assume that there is a unique copy of the trivial path $(s) \in P^{\prime}$.

In this step we will select paths such that each open path has $k \log ^{2} n$ children instead of the $k$ children one would expect. However, we use a function $y(p)$ that assigns a weight to each path and this weight decreases from layer to layer, modelling that the children are actually picked fractionally each with a $1 / 4 \log ^{2} n$ fraction of the weight of the parent. Thus, taking the weights into account we are actually picking only $k / 4$ children for each path.

Formally, the preprocessing of the LP will allow us to obtain a multiset of path $P^{\prime}$ such that

$$
\begin{align*}
\sum_{q \in C_{P^{\prime}}(p)} y(q)=\frac{k}{4} \cdot y(p) & \forall p \in P^{\prime} \backslash T^{\prime}  \tag{3.6}\\
\sum_{q \in I_{P^{\prime}}(v) \cap D_{P^{\prime}}(p)} y(q) \leqslant 2 y(p) & \forall p \in P^{\prime}, v \in V  \tag{3.7}\\
\text { where } y(p)=\frac{1}{\left(4 \log ^{2}(n)\right)^{i}} & \forall i \leqslant h, \forall p \in L_{i} \tag{3.8}
\end{align*}
$$

We obtain such a solution $P^{\prime}$ in a similar way as the randomized rounding in $[15,21$ ], which achieves polylogarithmic congestion. The fact that we select more paths, but only fractionally gives us better concentration bounds, which allows us to lose only constant congestion here. For completeness we give the proof in Appendix A.2.

### 3.4.2 The main rounding

We start this part with the sparse multiset of paths $P^{\prime}$ with the properties as above. The discount value $y$ can be thought of fractionally in the sense that each $p \in P^{\prime}$ is taken to an extent of $y(p)$. We will proceed to round this fractional solution to an integral solution $Q$ that is locally nearly good, a concept we will define formally below. Intuitively, this specifies that the number of paths that have locally high congestion can be removed without losing much with Lemma 3.9. We fix $\ell=10 \log \log n$ for the rest of this section. Let cong $(v \mid Q)$ be the global congestion of vertex $v$ induced by $Q$, that is, the number of paths in $Q$ ending in $v$. For paths $p \in Q \cup\{\emptyset\}$ and $q \in D_{Q}(p)$ We denote by $\operatorname{cong}_{p}(q \mid Q)$ the local congestion induced on the endpoint of $q$ by descendants of $p$. We consider all paths descendants of $\emptyset$.

A locally nearly good solution is a multiset of paths $Q \subseteq P^{\prime}$ (where again every path has a unique parent among the relevant copies of the same path) that has the following properties:

1. one copy of the trivial path $(s)$ belongs to $Q$;
2. every open path has $k / 32$ children;
3. no vertex has global congestion more than $2^{10} \log ^{3}(n)$;
4. for every $p \in Q \cup\{\emptyset\}$ and $\ell^{\prime} \leqslant \ell$ we have

$$
\left|\left\{q \in D_{Q}\left(p, \ell^{\prime}\right) \mid \operatorname{cong}_{p}(q \mid Q)>2^{10} \ell^{2}\right\}\right| \leqslant \frac{1}{\ell^{2}}\left(\frac{k}{32}\right)^{\ell^{\prime}}
$$

From a locally nearly good solution we will then derive a solution of low local congestion by removing all paths with high local congestion using bottom to top pruning (Lemma 3.9). Condition 4 is tailored to ensure that the number of such high local congestion paths is small enough so that the lemma succeeds.
To obtain such a nearly good solution, we will proceed layer by layer, where the top layers are already rounded integrally and the bottom layers are still fractional (as in the preprocessing). However, unlike the case of the preprocessing, we cannot argue with high probability and union bounds, since some properties we want only deviate from expectation by poly $(\log \log n)$ factors. To obtain a locally nearly good solution, we will use Lovász Local Lemma (LLL) in every iteration, where one iteration rounds one more layer. In the following, we describe the rounding procedure, the bad events and an analysis of their dependencies and finally we apply LLL.
the randomized rounding procedure. We proceed layer by layer to round the solution $P^{\prime}$ (with discount $y$ ) to an integral solution $Q$. We start by adding a single copy the trivial path $(s)$ to the partial solution $Q^{(0)}$. Assume we rounded until layer $i$, that is, we selected the final multiset of paths $Q^{(i)}$ to be used from all paths in $L_{\leqslant i}$.

To round until layer $i+1$ we proceed as follows. Every open path $p \in$ $Q^{(i)} \cap L_{i}$ selects exactly $k / 16$ children where the $i$-th child equals $q \in C_{P^{\prime}}(p)$ with probability equal to

$$
\frac{1}{k \log ^{2} n}=\frac{y(q)}{\sum_{q^{\prime} \in C_{P^{\prime}}(p)} y\left(q^{\prime}\right)}
$$

The selection of each child is independent of the choices made for other children. We then let $Q^{(i+1)}$ be the union of $Q^{(i)}$ and all newly selected paths. The reason that each path selects $k / 16$ children instead of the $k / 32$ many that were mentioned before is that we will later lose half of the children. We will repeat this procedure until reaching the last layer $h$ and we return $Q=Q^{(h)}$.
definitions related to expected congestion. In order for this iterated rounding to succeed we need to avoid that vertices get high congestion (in the local and the global sense). It is not enough to keep track only of
the congestion of vertices in the next layer that we are about to round, but we also need to maintain that the expected congestion (over the remaining iterations) remains low on all vertices in later layers. Hence, we will define quantities that help us keep track of them. To avoid confusion, we remark that the quantities we will define do not exactly correspond to the expected congestion of the vertices: notice that we sample less children for each path than $P^{\prime}$ has. Intuitively, $P^{\prime}$ has $k / 4$ children per open path (see Equation (3.6)), but we only sample $k / 16$ many. The quantities we define below would be the expectations if we would sample $k / 4$ children instead. Roughly speaking, this gives us an advantage of the form that the expectation always decreases by a factor 4 when we round one iteration. Apart from the intuition on the expectation, the definitions also incorporate a form of conditional congestion, which means that we consider the (expected) congestion based on random choices made so far. This is similar to notions that appear in Appendix A. 2 for the preprocessing step.

First, we define the fractional congestion induced by some path $p \in P^{\prime}$ on a vertex $v \in V$ as follows.

$$
\operatorname{cong}(p, v):=\sum_{q \in D_{P^{\prime}}(p) \cap I_{P^{\prime}}(v)} y(q) .
$$

Using this definition, we will define the conditional fractional congestion at step $i$ induced by $p$ on a descendant $q$ as follows (writing $v_{q}$ for the endpoint of $q$ ).

$$
\operatorname{cong}_{p}\left(q \mid Q^{(i)}\right):= \begin{cases}\left|\left\{q^{\prime} \in D_{Q^{(i)}}(p) \cap I_{Q^{(i)}}\left(v_{q}\right)\right\}\right| & \text { if } q \in L_{\leqslant i} \\ \sum_{q^{\prime} \in D_{Q^{(i)}}(p) \cap L_{i}} \frac{1}{y\left(q^{\prime}\right)} \operatorname{cong}\left(q^{\prime}, v_{q}\right) & \text { otherwise }\end{cases}
$$

We will use this definition only for $p \in Q^{(i)}$. If $v_{q}$ belongs to one of the first $i$ layers (i.e., $v_{q}$ is in the integral part corresponding to the partial solution), this is simply the number of descendants of $p \in Q^{(i)}$ in our partial solution that end at $v_{q}$. Otherwise, this is the total fractional congestion induced by all the paths that we actually selected in our partial solution $q \in Q^{(i)} \cap L_{i}$ and that are descendants of $p$. The name conditional comes from the term $1 / y\left(q^{\prime}\right)$ which is simply the fact that we condition on having already selected the path $q^{\prime}$ once integrally. The global congestion at step $i$ on a vertex $v \in V$ is defined similarly with

$$
\operatorname{cong}\left(v \mid Q^{(i)}\right):= \begin{cases}\left|\left\{q \in Q^{(i)} \cap I(v)\right\}\right| & \text { if } v \in L_{\leqslant i} \\ \sum_{q \in Q^{(i)} \cap L_{i}} \frac{1}{y(q)} \operatorname{cong}(q, v) & \text { otherwise. }\end{cases}
$$

thefirst type of bad event: global congestion. The naive way to bound the global congestion would be to simply to bound the global congestion on each vertex and make a bad event from exceeding this. However, to manage dependencies between bad events and get better concentration bounds, we partition the set of paths according to how much fractional load their ancestors are expected to put on $v$ and then bound the congestion incurred by each group on $v$. More precisely, assume we rounded until layer $i$
and are now trying to round until layer $L_{i+1}$. Consider any vertex $v \in L_{\geqslant i+1}$. The decisions made in this step of the rounding concern which paths in $L_{i+1}$ will be selected. Fix an integer $t \geqslant 0$ and and let $P_{v, t}^{(i+1)}$ to be the set of all paths $p \in L_{i+1}$ such that

$$
\frac{\operatorname{cong}(p, v)}{y(p)} \in\left(2^{-t}, 2^{-(t-1)}\right] .
$$

We define the bad event $\mathcal{B}_{1}(v, t)$ as the event that

$$
\sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)}>2 \mathbb{E}\left[\sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)}\right]+2^{10} \log n .
$$

Since $P_{v, t}^{(i+1)}$ partitions the paths in $L_{i+1}$, the bad events for all $t$ together will bound the increase in the congestion on $v$ : this is because cong $\left(v \mid Q^{(i+1)}\right)=$ $\sum_{p \in Q^{(i+1)} \cap L_{i+1}} \operatorname{cong}(p, v) / y(p)$. We argue this formally in Lemma 3.13 below.

At this point, it is worthwhile to mention that $t$ can only take values in $\left\{0,1, \ldots, \log ^{2} n\right\}$. Indeed, by Constraint (3.7) we have that

$$
\frac{\operatorname{cong}(p, v)}{y(p)} \leqslant 2
$$

is ensured for all $p \in P^{\prime}$. Moreover, we notice that $y$ has a low granularity (i.e. $y$ satisfies Equation (3.8)) which implies that either $\operatorname{cong}(p, v)=0$ or

$$
\frac{\operatorname{cong}(p, v)}{y(p)} \geqslant \frac{1}{\left(4 \log ^{2} n\right)^{h-|p|}} \geqslant \frac{1}{\left(4 \log ^{2} n\right)^{h}} .
$$

Hence bad events $\mathcal{B}_{1}(v, t)$ will be instantiated only for $t=0,1, \ldots, \log ^{2} n$ (assuming here that $n$ is sufficiently large).

Lemma 3.13. Assuming no bad event $\mathcal{B}_{1}(v, t)$ has occurred in any iteration up to $i$, we have for every $v$ that

$$
\operatorname{cong}\left(v \mid Q^{(i)}\right) \leqslant 2^{11} \log ^{3} n
$$

Proof. We argue inductively. For $i=0$ we have

$$
\begin{aligned}
\operatorname{cong}\left(v \mid Q^{(i)}\right)=\sum_{p \in Q^{(i)} \cap L_{i}} \frac{\operatorname{cong}(q, v)}{y(q)}=\sum_{s \in S} & \operatorname{cong}((s), v) \\
& =\sum_{p \in I_{P^{\prime}}(v)} y(v) \leqslant 2 .
\end{aligned}
$$

Now assume that for $i \geqslant 0$ we have

$$
\operatorname{cong}\left(v \mid Q^{(i)}\right) \leqslant 2^{11} \log ^{3} n
$$

Since the congestion of any $v \in L_{\leqslant i}$ cannot change anymore, we may assume w.l.o.g. that $v \in L_{\geqslant i+1}$. For each $t$ let

$$
\mu_{t}=\mathbb{E}\left[\sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)}\right] .
$$

Then $\sum_{t} \mu_{t}=1 / 4 \cdot \operatorname{cong}\left(v \mid Q^{(i)}\right) \leqslant 2^{9} \log ^{3} n$. To see this, denote by $N_{q}^{(i)}$ the number of copies of a path $q \in P^{\prime}$ that is selected in $Q^{(i)}$. As a shorthand, we write $P(q)$ for the parent of $q$ in $Q^{(i+1)}$. By linearity of expectation we can write

$$
\begin{aligned}
\sum_{t} \mu_{t} & =\sum_{t} \mathbb{E}\left[\sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)}\right] \\
& =\sum_{t} \mathbb{E}\left[\sum_{p \in P^{\prime} \cap P_{v, t}^{(i+1)}} N_{p}^{(i+1)} \cdot \frac{\operatorname{cong}(p, v)}{y(p)}\right] \\
& =\sum_{t} \sum_{p \in P^{\prime} \cap P_{v, t}^{(i+1)}} \mathbb{E}\left[N_{p}^{(i+1)}\right] \cdot \frac{\operatorname{cong}(p, v)}{y(p)} \\
& =\sum_{t} \sum_{p \in P^{\prime} \cap P_{v, t}^{(i+1)}} \frac{(k / 16) \cdot y(p)}{(k / 4) \cdot y(P(p))} \cdot \frac{\operatorname{cong}(p, v)}{y(p)} \\
& =\frac{1}{4} \cdot\left[\sum_{t} \sum_{p \in P^{\prime} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(P(p))}\right] \\
& =\frac{1}{4} \cdot\left[\sum_{t} \sum_{p \in P^{\prime} \cap P_{v, t}^{(i)}} \frac{\operatorname{cong}(p, v)}{y(p)}\right] \\
& =(1 / 4) \cdot \operatorname{cong}\left(v \mid Q^{(i)}\right) \leqslant 2^{9} \log ^{3} n .
\end{aligned}
$$

It follows that

$$
\begin{aligned}
\operatorname{cong}\left(v \mid Q^{(i+1)}\right) & =\sum_{p \in Q^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)} \\
& =\sum_{t} \sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)} \\
& \leqslant \sum_{t}\left[2 \mu_{t}+2^{10} \log n\right] \\
& \leqslant 2^{10} \log ^{3} n+2^{10} \log ^{3} n \\
& \leqslant 2^{11} \log ^{3} n
\end{aligned}
$$

the second type of bad event: local congestion. Recall that we want to bound the congestion induced by close descendants of some open path $p$. Let $\ell^{\prime} \leqslant \ell$ and $p \in Q^{(i)} \cap L_{i^{\prime}}$, where $i-\ell^{\prime} \leqslant i^{\prime} \leqslant i$. We define

$$
\begin{aligned}
N^{(i+1)}\left(p, \ell^{\prime}\right)=\mid\left\{q \in D_{P^{\prime}}\left(p, \ell^{\prime}\right) \mid\right. & \operatorname{cong}_{p}\left(q \mid Q^{(i)}\right) \leqslant 2^{10} \ell^{2} \\
& \text { and } \left.\operatorname{cong}_{p}\left(q \mid Q^{(i+1)}\right)>2^{10} \ell^{2}\right\} \mid
\end{aligned}
$$

These are the number of vertices with newly high local congestion (counting only the congestion induced by descendants of $p$ ). Moreover, note that by Constraint (3.6) and (3.8), we have that $\left|D_{P^{\prime}}\left(p, \ell^{\prime}\right)\right| \leqslant\left(k \log ^{2} n\right)^{\ell^{\prime}}$.

Then, for any $p \in L_{i^{\prime}} \cap Q^{(i)}$ where $i-\ell-1 \leqslant i^{\prime} \leqslant i-1$, and $t \geqslant 0$, we define the set of marked children of $p$ at step $i+1$ as the set $M^{(i+1)}(p)$ of all $q \in C_{Q^{(i)}}(p)$ such that

$$
N^{(i+1)}\left(q, \ell^{\prime}\right)>\frac{1}{\log ^{10} n}\left(\frac{k}{32}\right)^{\ell^{\prime}}
$$

for at least one $\ell^{\prime} \leqslant \ell$. Notice that the children of $p$ (in $Q^{(i)}$ ) are already determined because $p$ is in $L_{\leqslant i-1}$. We are now ready to state the second type of bad event. We define the bad event $\mathcal{B}_{2}(p)$ for any $p \in L_{i^{\prime}}$ where $i-\ell-1 \leqslant i^{\prime} \leqslant i-1$, as the event that

$$
\left|M^{(i+1)}(p)\right| \geqslant \frac{1}{\ell^{3}} \cdot \frac{k}{16}
$$

which means that a lot of children selected by $p$ become marked at step $i+1$. As we will ensure that this bad event never happens, we can guarantee that a large fraction of children $q$ of each path $p$ always satisfies that $N^{(i+1)}\left(q, \ell^{\prime}\right)$ is low for all $\ell^{\prime} \leqslant \ell$. As we explain later, we will simply remove all other children and this way indirectly prevent the existence of too many locally congested descendants for all the remaining paths. The advantage of this indirect method is that we can apply good concentration bounds on the children and thereby amplify the probability (indeed, note that the event that a child $q$ of $p$ is marked is independent of the same events for other children of $p$ ).

THE THIRD Type of bAd EVENT: KEEPING THE SOURCE PATH $(s)$. Because our algorithm will delete in the end all the paths that have been marked in some round, we need to ensure that the source path $(s)$ is never marked. Hence for the first $\ell$ steps of rounding, we define a single bad event $\mathcal{B}_{3}(s)$ that is that the source path $(s)$ is marked during this step.
the probabilities of the bad events. In this part, we derive upper bounds for the probability of each bad event. Some bounds are sub-optimal to simplify later formulas.

Lemma 3.14. For any $v \in V$ and integer $t$ we have that

$$
\mathbb{P}\left[\mathcal{B}_{1}(v, t)\right] \leqslant n^{-10 \cdot 2^{t-1}}
$$

Proof. We need to prove that

$$
\begin{gathered}
\mathbb{P}\left[\sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)}>2 \mathbb{E}\left[\sum_{p \in Q^{(i+1)} \cap P_{v, t}^{(i+1)}} \frac{\operatorname{cong}(p, v)}{y(p)}\right]+2^{10} \log n\right] \\
\leqslant n^{-10 \cdot 2^{t-1}}
\end{gathered}
$$

Notice that cong $(p, v) / y(p)$ is upper bounded by $2^{-(t-1)}$ by definition of $P_{v, t}^{(i+1)}$ and the sum can be rewritten as a sum over the paths chosen in iteration $i+1$, where a path $p$ contributes cong $(p, v) / y(p)$ to the sum. These
paths are chosen independently and thus by a standard Chernoff bound we obtain that the probability is at most

$$
\exp \left(\frac{-2^{10} \log n}{2 \cdot 2^{-(t-1)}}\right) \leqslant n^{-10 \cdot 2^{t-1}}
$$

Lemma 3.15. Let $\ell^{\prime} \leqslant \ell$ and $p \in Q^{(i)} \cap L_{i^{\prime}}$, where $i-\ell^{\prime} \leqslant i^{\prime} \leqslant i$. Then

$$
\begin{equation*}
\mathbb{P}\left[N^{(i+1)}\left(p, \ell^{\prime}\right)>\left(1 / \log ^{10} n\right) \cdot(k / 32)^{\ell^{\prime}}\right] \leqslant(\log n)^{-40 \ell} . \tag{3.9}
\end{equation*}
$$

Moreover,

$$
\begin{equation*}
\mathbb{P}\left[\mathcal{B}_{3}(s)\right] \leqslant(\log n)^{-1} . \tag{3.10}
\end{equation*}
$$

Proof. Recall that we are bounding the number of descendants of $p$ that had low local congestion of at most $2^{10} \ell^{2}$ before, but now get high local congestion. More concretely, consider some $q \in D_{P^{\prime}}\left(p, \ell^{\prime}\right)$ with $\operatorname{cong}_{p}(q \mid$ $\left.Q^{(i)}\right) \leqslant 2^{10} \ell^{2}$. Then

$$
\mathbb{E}\left[\operatorname{cong}_{p}\left(q \mid Q^{(i+1)}\right)\right]=\frac{1}{4} \cdot \operatorname{cong}_{p}\left(q \mid Q^{(i)}\right) \leqslant 2^{8} \ell^{2} .
$$

We do not go into detail for this calculation here but similar calculations have been derived already in Lemma 3.13 or Appendix A.2. This is exactly the same type or argument.

Further, $\operatorname{cong}_{p}\left(q \mid Q^{(i+1)}\right)$ can be decomposed into a sum of independent random variables bounded by 2 . Indeed, recalling the definition of $\operatorname{cong}_{p}(q)$ $Q^{(i+1)}$ ) we see that this quantity can be written as a sum over $1 / y\left(q^{\prime}\right)$. cong $_{q^{\prime}}\left(q \mid Q^{(i+1)}\right)$, where a paths $q^{\prime}$ are taken from a certain multiset of chosen paths. Constraint (3.7) of the path LP ensures that this is always at most 2 . Those variables are independent because of the randomized rounding that selects children independently of each other. Hence, by a Chernoff bound we have

$$
\mathbb{P}\left[\operatorname{cong}_{p}\left(q \mid Q^{(i+1)}\right)>2^{10} \ell^{2}\right] \leqslant \exp \left(-2^{7} \ell^{2}\right)
$$

Therefore, by linearity of expectation we obtain

$$
\mathbb{E}\left[N^{(i+1)}\left(p, \ell^{\prime}\right)\right] \leqslant \frac{\left(k \log ^{2} n\right)^{\ell^{\prime}}}{\exp \left(2^{7} \ell^{2}\right)} \leqslant \frac{1}{(\log n)^{2^{7} \cdot 10 \ell}}\left(\frac{k}{32}\right)^{\ell^{\prime}}
$$

From Markov's inequality it follows that

$$
\mathbb{P}\left[N^{(i+1)}\left(p, \ell^{\prime}\right)>\left(1 / \log ^{10} n\right) \cdot(k / 32)^{\ell^{\prime}}\right] \leqslant(\log n)^{-40 \ell} .
$$

For the second claim, notice that $\mathcal{B}_{3}(s)$ is simply the event that the source path $(s)$ is marked, which is equivalent to $N^{(i+1)}\left((s), \ell^{\prime}\right)>\left(1 / \log ^{10} n\right) \cdot(k / 32)^{\ell^{\prime}}$ for at least one $\ell^{\prime} \leqslant \ell$. By a simple union bound, we get the desired result.

Finally, we upper-bound the probability of $\mathcal{B}_{3}(p)$.

Lemma 3.16. Let $p \in L_{i^{\prime}}$ where $i-\ell-1 \leqslant i^{\prime} \leqslant i-1$. Then we have that

$$
\mathbb{P}\left[\mathcal{B}_{3}(p)\right] \leqslant \exp (-\sqrt{k}) .
$$

Proof. To prove this, note that a child $q$ of $p$ (in the set $Q^{(i)}$ ) is marked independently of other children of $p$. This is because $q$ being marked depends only on the random choices made by descendants of $q$, which are independent of descendants of other children of $p$. Therefore using Lemma 3.15 and a standard union bound we obtain that each child of $p$ is marked independently with probability at most $\ell \cdot(\log n)^{-40 \ell}$. Recall that $k \geqslant 2^{10}(\log \log n)^{8} \geqslant$ $\left(32 \ell^{3}\right)^{2}$. We note that $p$ has $k / 16$ children in $Q^{(i)}$ and therefore the probability that more than $\left(1 / \ell^{3}\right) \cdot(k / 16)$ of them are marked is at most

$$
\exp \left(-\frac{k}{32 \ell^{3}}\right) \leqslant \exp (-\sqrt{k})
$$

the dependencies of the bad events. For any bad event $\mathcal{B}$ we define $\Gamma_{1}(\mathcal{B})$ to be the set of bad events of the first type $\mathcal{B}_{1}(v, t)$ that depend on the bad event $\mathcal{B}$. Similarly, let $\Gamma_{2}(\mathcal{B})$ the set of bad events $\mathcal{B}_{2}(p)$ that depend on $\mathcal{B}$. We will now upper bound the cardinality of these sets. Note that there is a single bad event of third type $\mathcal{B}_{3}(s)$ hence it is clear that $\Gamma_{3}(\mathcal{B}) \leqslant 1$ for any bad event $\mathcal{B}$. For the rest, we remark that the focus here is on simplicity of the terms rather than optimizing the precise bounds.

Lemma 3.17. For any bad event $\mathcal{B}$, we have that

$$
\left|\Gamma_{1}(\mathcal{B})\right| \leqslant n^{2} .
$$

Proof. The statement holds trivially, since there are at most $n \cdot\left(1+\log ^{2} n\right) \leqslant$ $n^{2}$ bad events of the first type in total ( $n$ possibilities of the choice of vertex $v$ and less than $n$ possibilities for the choice of $t$ ).

Before proving the dependencies to events of type 2, we will prove an auxiliary lemma that concerns the events affected by the children picked by one particular path $p \in Q^{(i)} \cap L_{i}$.

Lemma 3.18. Let $p \in Q^{(i)} \cap L_{i}$. Then the choice of children picked by $p$ affects in total at most $\log n$ events of type 2 .

Proof. In order to influence a bad event $B_{2}(q)$ for some $q$, it must be that $p$ is a descendant of $q$ in the multiset $Q^{(i)}$. Any path has at most $h \leqslant \log n$ ancestors.

Lemma 3.19. Assuming that no bad event has happened until the current iteration of rounding, for any $v \in V$ and integers $t \geqslant 0$, we have

$$
\begin{equation*}
\left|\Gamma_{2}\left(\mathcal{B}_{1}(v, t)\right)\right| \leqslant 2^{t} \cdot(\log n)^{10} \tag{3.11}
\end{equation*}
$$

Proof. We first notice that the bad event $\mathcal{B}_{1}(v, t)$ depends only on the random choices made by paths $q \in Q^{(i)} \cap L_{i}$ that have at least one child $q^{\prime} \in C_{P^{\prime}}(q)$ such that

$$
\frac{\operatorname{cong}\left(q^{\prime}, v\right)}{y\left(q^{\prime}\right)} \geqslant 2^{-t}
$$

Let us count how many such paths $q$ can exist in $Q^{(i)}$. For each such path we have that

$$
\begin{aligned}
\frac{\operatorname{cong}(q, v)}{y(q)}=\sum_{q^{\prime} \in C_{P^{\prime}}(q)} \frac{\operatorname{cong}\left(q^{\prime}, v\right)}{y(q)} & =\sum_{q^{\prime} \in C_{P^{\prime}}(q)} \frac{\operatorname{cong}\left(q^{\prime}, v\right)}{\left(4 \log ^{2} n\right) y\left(q^{\prime}\right)} \\
& \geqslant \frac{2^{-t}}{4 \log ^{2} n}
\end{aligned}
$$

where we used here the granularity of $y$, see Constraint (3.8). Since we assume that no bad event happened so far, we can apply Lemma 3.13 to get

$$
\sum_{q \in Q^{(i)} \cap L_{i}} \frac{\operatorname{cong}(q, v)}{y(q)}=\operatorname{cong}\left(v \mid Q^{(i)}\right) \leqslant 2^{11} \log ^{3} n
$$

Hence, there can be at most $2^{t} \cdot 2^{13} \log ^{5} n$ such paths $q$. We conclude with Lemma 3.18 to bound the number of bad events of second type influenced by each $q$ and obtain

$$
\left|\Gamma_{2}\left(\mathcal{B}_{1}(v, t)\right)\right| \leqslant 2^{t} \cdot\left(2^{13} \log ^{5} n\right) \cdot(\log n) \leqslant 2^{t} \cdot\left(\log ^{10} n\right)
$$

Lemma 3.20. For any $p \in L_{i^{\prime}}$ where $i-\ell-1 \leqslant i^{\prime} \leqslant i-1$ we have

$$
\begin{equation*}
\left|\Gamma_{2}\left(\mathcal{B}_{2}(p)\right)\right| \leqslant k^{\ell} \cdot(\log n) \tag{3.12}
\end{equation*}
$$

## Furthermore,

$$
\begin{equation*}
\left|\Gamma_{2}\left(\mathcal{B}_{3}(s)\right)\right| \leqslant k^{\ell} \cdot(\log n) \tag{3.13}
\end{equation*}
$$

Proof. The bad event $\mathcal{B}_{2}(p)$ (or $\mathcal{B}_{3}(s)$ ) is entirely determined by the random choices made by the descendants of $p$ (or $(s)$ ) in $Q^{(i)}$ that are located in layer $L_{i}$. There are at most $(k / 16)^{\ell} \leqslant k^{\ell}$ such descendants $q$. By Lemma 3.18, each descendant can influence at most $\log n$ other bad events of the second type. This concludes the proof.
instantiating lll. For the bad events we set

$$
\begin{align*}
x\left(\mathcal{B}_{1}(v, t)\right) & =n^{-2}  \tag{3.14}\\
x\left(\mathcal{B}_{2}(p)\right) & =\exp \left(-k^{1 / 3}\right)  \tag{3.15}\\
x\left(\mathcal{B}_{3}(s)\right) & =1 / 2 \tag{3.16}
\end{align*}
$$

Consider now a bad event $\mathcal{B}_{1}(v, t)$ and recall that $k \geqslant 2^{10}(\log \log n)^{8}$. Then

$$
\begin{align*}
\mathbb{P}\left[\mathcal{B}_{1}(v, t)\right] & \leqslant n^{-10 \cdot 2^{(t-1)}} \\
& \leqslant n^{-5 \cdot 2^{t}} \\
& \leqslant n^{-2} \cdot\left(1-n^{-2}\right)^{n^{2}} \cdot\left(1-\exp \left(-k^{1 / 3}\right)\right)^{2^{t} \cdot(\log n)^{10}}  \tag{1/2}\\
& \leqslant n^{-2} \cdot \prod_{\mathcal{B} \in \Gamma\left(\mathcal{B}_{1}(v, t)\right)}(1-x(\mathcal{B}))
\end{align*}
$$

Next, we get

$$
\begin{aligned}
\mathbb{P}\left[\mathcal{B}_{2}(p)\right] & \leqslant \exp (-\sqrt{k}) \\
& \leqslant \exp \left(-k^{1 / 3}\right) \cdot \frac{\left(1-n^{-2}\right)^{n^{2}}}{2} \cdot\left(1-\exp \left(-k^{1 / 3}\right)\right)^{k^{\ell} \cdot(\log n)} \\
& \leqslant x\left(\mathcal{B}_{2}(p)\right) \cdot \prod_{\mathcal{B} \in \Gamma\left(\mathcal{B}_{2}(p)\right)}(1-x(B))
\end{aligned}
$$

Similarly, for $\mathcal{B}_{3}(s)$ we have

$$
\begin{aligned}
\mathbb{P}\left[\mathcal{B}_{3}(s)\right] & \leqslant(\log n)^{-1} \\
& \leqslant(1 / 2) \cdot\left(1-n^{-2}\right)^{n^{2}} \cdot\left(1-\exp \left(-k^{1 / 3}\right)\right)^{k^{\ell} \cdot(\log n)} \\
& \leqslant x\left(\mathcal{B}_{3}(s)\right) \cdot \prod_{\mathcal{B} \in \Gamma\left(\mathcal{B}_{3}(s)\right)}(1-x(B))
\end{aligned}
$$

Hence, with positive probability none of the bad events occur and we can compute such a solution in expected quasi-polynomial time (polynomial in $\left|P^{\prime}\right|$ ).

FINISHING NOTES. We proved that we can round in (expected) quasipolynomial time from layer 0 to layer $h$ without any bad event ever happening. Next, observe that we get that any sampled path $p \in Q \cap L_{i}$ has at most

$$
\left|\bigcup_{i^{\prime}=i+1}^{i+\ell+1} M^{\left(i^{\prime}\right)}(p)\right| \leqslant \ell \cdot \frac{1}{\ell^{3}} \cdot \frac{k}{16} \leqslant \frac{k}{32}
$$

marked children. We now delete from the solution all the marked paths obtaining a solution $Q^{\prime}$. Note that the source remains as it is never marked. This means that any sampled open path retains at least $k / 32$ children. For any path $p$, we denote by $N\left(p, \ell^{\prime}\right)$ the number of its descendants at distance $\ell^{\prime}$ that have congestion induced by $p$ bigger than $2^{10} \ell^{2}$. Then if $p$ was never marked at any round, we have that

$$
N\left(p, \ell^{\prime}\right) \leqslant \frac{1}{\log n}\left(\frac{k}{32}\right)^{\ell^{\prime}}
$$

Indeed, if $p$ belongs to layer $i$ then $\operatorname{cong}_{p}\left(q \mid Q^{(i)}\right)=1 / y(p) \cdot \operatorname{cong}\left(p, v_{q}\right) \leqslant 2$ by Constraint (3.7). Hence the first time we instantiate bad events involving local congestion induced by the path $p$, it is the case that none of its descendants are bad. Then if $p$ is never marked we have that

$$
N^{\left(i^{\prime}\right)}\left(p, \ell^{\prime}\right) \leqslant \frac{1}{\log ^{10} n} \cdot\left(\frac{k}{32}\right)^{\ell^{\prime}}
$$

at every iteration $i \leqslant i^{\prime} \leqslant i+\ell$, which means that few of the good descendants become bad at each round. Hence we obtain that

$$
N\left(p, \ell^{\prime}\right) \leqslant \sum_{i^{\prime}=i}^{i+\ell} N^{\left(i^{\prime}\right)}\left(p, \ell^{\prime}\right) \leqslant \frac{1}{\log n}\left(\frac{k}{32}\right)^{\ell^{\prime}}
$$

This ensures that any remaining path has few bad descendants (in terms of local congestion).

We now define a set of paths $R \subseteq Q$, which contains all paths with high local congestion. More precisely, let $R$ contain for every path $p \in Q$ and $\ell^{\prime} \leqslant \ell$ all paths $q \in D_{Q}\left(p, \ell^{\prime}\right)$ such that $\operatorname{cong}_{p}\left(q \mid Q^{\prime}\right)>2^{10} \ell^{2}$. Since we have no more marked paths, each $p$ has at most $(\log n)^{-1} \cdot(k / 32)^{\ell^{\prime}} \leqslant$ $1 /(8 \ell)^{2} \cdot(k / 32)^{\ell^{\prime}}$ descendants at distance $\ell^{\prime}$ in $R$. Applying Lemma 3.4 we can remove $R$ and obtain a solution $Q^{\prime} \subseteq Q \backslash R$, where each open path has $k /(64 \ell)$ children and the source remains in the solution. Moreover, the solution has no more paths with $\ell$-local congestion more than $2^{10} \ell^{2}$ and the global congestion is at most $2^{11} \log ^{3} n$. This finishes the proof of Lemma 3.7.

### 3.5 FROM SINGLE SOURCE TO MULTIPLE SOURCES

We devise an algorithm that solves the problem with multiple sources based on an algorithmic technique first introduced by Haxell in the context of hypergraph matchings and then applied to a range of other problems, including the restricted assignment version of the Santa Claus problem. Haxell's algorithmic technique can be thought of as a (highly non-trivial) generalization of the augmenting path algorithm for bipartite matching. Our algorithm makes only calls to the $\alpha$-approximation for a single source. The algorithm itself requires quasi-polynomial time as well (even if the $\alpha$-approximation ran in polynomial time). Our concrete variant relies on the following simple, but powerful subprocedure.
Lemma 3.21. Let $Q$ be a (single) degree- $k^{\prime}$ arborescence in a layered graph. Let $R \subseteq Q$ be a set of paths where at most $\left(k^{\prime} / 4\right)^{i}$ many end in each layer $i$. Then there is a degree- $k^{\prime} / 4$ arborescence $Q^{\prime} \subseteq Q \backslash R$. Furthermore, $Q^{\prime}$ can be computed in polynomial time.

Proof. We start pruning the arborescence from bottom to top and we maintain at every layer $i$ that we remove at most $\left(k^{\prime} / 2\right)^{i}$ paths in it. At the last layer $i=h$ we remove simply the paths that are in $R$. These are by definition at most $\left(k^{\prime} / 4\right)^{i} \leqslant\left(k^{\prime} / 2\right)^{i}$ many. Then assume we already pruned layers $i, i+1, \ldots, h$ and that we removed at most $\left(k^{\prime} / 2\right)^{i}$ many paths in layer $i$. Now in layer $i-1$ we again remove the at most $\left(k^{\prime} / 4\right)^{i-1}$ many paths in $R$, but also those open paths where more than $3 / 4 \cdot k^{\prime}$ many children were removed in layer $i$. The number of open paths removed this way is at most $\left(k^{\prime} / 2\right)^{i} /\left(3 / 4 \cdot k^{\prime}\right) \leqslant 2 / 3 \cdot\left(k^{\prime} / 2\right)^{i-1}<\left(k^{\prime} / 2\right)^{i-1}$. Clearly, this procedure maintains that every remaining open path has at least $k^{\prime} / 4$ children and the source never gets removed.

Throughout the section we assume that $k$ is the optimum, which is obtained through a binary search framework. We now assume that we have a degree$k / 256 \alpha$ solution that already covers all but one source $s_{0} \in S$. We will
augment this solution to one that covers all sources. This is without loss of generality, since we can apply the procedure iteratively $|S|$ times, adding one source at a time.
blocking trees and addable trees. Let $Q$ be our current solution (which does not cover $s_{0}$ ). On an intuitive level, blocking trees are $k / 256 \alpha-$ degree arborescences in our current solution $Q$, which we would like to remove from the solution. In order to remove them, we need to add other arborescences for their sources instead. The addable trees are $k / 32 \alpha$-degree arborescences not in our solution. Their sources are sources of blocking trees in $Q$ and we would like to add to them solution. The addable trees in turn may be blocked by blocking trees, which means they overlap on some non-source vertex with a tree in $Q$, preventing us from adding them to the solution. We note that addable trees are (by a factor of 8 ) better arborescences than what we ultimately need. As explained above, addable and blocking trees naturally form an alternating structure, which is usually referred to as layers. In order to avoid conflicts with our other notion of layers, we call these rings. The addable and blocking trees in the $i$ th ring will be denoted by $A_{i}$ and $B_{i}$.
the algorithm. Initially we have no rings. We run our $\alpha$-approximation to find a $k / \alpha$-degree arborescence for $s_{0}$ (without taking into account our current solution $Q$ ). We reduce this to a $k / 32 \alpha$-degree arborescence and store it as singleton set in $A_{1}$. Then we store in $B_{1}$ all arborescences in $Q$ that intersect on any vertex with it. If the total intersection on each layer $j$ is at most $(k / 128 \alpha)^{j}$, then we can find through Lemma 3.21 a $k / 128 \alpha$-degree arborescence for $s_{0}$ which is disjoint from $Q$. We reduce it to a $k / 256 \alpha-$ degree arborescence and add it to the solution and terminate. Otherwise we continue the algorithm, but now we indeed have one ring of addable and blocking trees.

Assume now the algorithm in its current state has rings $A_{1}, B_{1}, \ldots, A_{i}, B_{i}$. Then we intialize $A_{i+1}=\emptyset$ and add addable trees to it in the following Greedy manner. From our (layered) graph we produce a reduced instance by removing any vertices appearing in $A_{1}, B_{1}, \ldots, A_{i}, B_{i}, A_{i+1}$ (except for the sources). Then we iterate over any source $s$ that is either $s_{0}$ or the source of a blocking tree in $B_{1} \cup \cdots \cup B_{i}$. We call our $\alpha$-approximation for $s$ on the reduced instance. If it outputs a $k^{\prime}$-degree solution with $k^{\prime} \geqslant k / 32 \alpha$, we add it to $A_{i+1}$ (after reducing to degree exactly $k / 32 \alpha$ ). We repeat this until we can no longer add any addable trees to $A_{i+1}$. Then we construct $B_{i+1}$ by taking any arborescence in $Q$ that intersects with any addable tree in $A_{i+1}$. As before, we check if any addable tree in $A_{i+1}$ has a total intersection of at most $(k / 128 \alpha)^{j}$ on each layer $j$ with solution $Q$. If so, we collapse: we can compute through Lemma 3.21 a $k / 128 \alpha$-degree arborescence for the corresponding source $s$, which is disjoint from $Q$. We reduce it to a $k / 256 \alpha-$ degree arborescence and add it to $Q$. Now $s$ has two arborescences in $Q$. The other arborescence must have appeared as a blocking tree in an earlier ring $B_{i^{\prime}}$. We remove this blocking tree from $Q$ and delete all rings after $i^{\prime}$. Then we revisit the addable trees in $A_{i^{\prime}}$. Since we removed a blocking tree in $B_{i^{\prime}}$, one of them may now have a small intersection with all trees in $Q$ (as above). If so, we collapse this as well. We continue so until no more collapse
is possible, leaving us with a new solution $Q$ and a prefix of the previous rings. We then continue the algorithm. Once an addable tree for $s_{0}$ is added to $Q$ we terminate.
analysis. Our analysis relies on two lemmas.
Lemma 3.22. At any point of time in the local search algorithm we have for every ring that $\left|B_{i}\right| \geqslant\left|A_{i}\right|$.

Proof. Recall that the addable trees in $A_{i}$ are vertex disjoint by construction. At the same time, none of them can be collapsed. This means for each of them there must be a layer $j$ such that the addable tree intersects with $Q$ (that is, with $B_{i}$ ) on $(k / 128)^{j}$ many vertices. Let $R_{j}$ be the set of all intersecting vertices with $Q$ in layer $j$ (across all addable trees in $A_{i}$ ). Then

$$
\sum_{j \geqslant 1} \frac{\left|R_{j}\right|}{(k / 128)^{j}}>\left|A_{i}\right| .
$$

On the other hand, each blocking tree in $B_{i}$ has at most $(k / 256)^{j}$ many vertices in each layer $j$. Thus,

$$
\sum_{j \geqslant 1} \frac{\left|R_{j}\right|}{(k / 128)^{j}} \leqslant \sum_{j \geqslant 1}\left|B_{i}\right| \cdot \frac{1}{2^{j}} \leqslant\left|B_{i}\right| .
$$

Lemma 3.23. After the Greedy procedure to construct $A_{i+1}$, we have that $\left|A_{i+1}\right| \geqslant \sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right|$.

Proof. By Lemma 3.22 we have that

$$
\sum_{i^{\prime}=1}^{i}\left|A_{i^{\prime}}\right|+\left|B_{i^{\prime}}\right| \leqslant 2 \sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right| .
$$

Now assume toward constradiction that $\left|A_{i+1}\right|<\sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right|$ and no more addable trees can be added. Recall that the reduced instance we consider removes all vertices appearing in $A_{1}, B_{1}, \ldots, A_{i}, B_{i}, A_{i+1}$. By the bounds above we know that this removes a total of at most

$$
\begin{equation*}
(k / 32)^{j} \cdot 3 \sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right| \tag{3.17}
\end{equation*}
$$

many vertices from each layer $j$. Now consider the optimal degree- $k$ solution restricted to the sources in $\bigcup_{i^{\prime}=1}^{i} B_{i^{\prime}}$. Each of the arborescences in this solution must overlap with the removed vertices on some layer $j$ on at least $(k / 4)^{j}$ many vertices. Otherwise, by Lemma 3.22 there would be a degree- $k / 4$ arborescence rooted in one of the sources of a blocking tree which is disjoint from the removed vertices. Consequently, our $\alpha$-approximation would have found a degree- $k / 4 \alpha$ arborescence that would have been added as an addable
tree. Let $R_{j}$ be the intersection of the aforementioned optimal solution with the removed vertices. Then

$$
\sum_{j \geqslant 1} \frac{\left|R_{j}\right|}{(k / 4)^{j}} \geqslant \sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right| .
$$

However, from (3.17) it follows that

$$
\sum_{j \geqslant 1} \frac{\left|R_{j}\right|}{(k / 4)^{j}} \leqslant\left(3 \sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right|\right) \cdot \sum_{j \geqslant 1}\left(\frac{1}{8}\right)^{j}<\sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right|
$$

a contradiction.
From the lemmas above it follows that the local search never gets stuck: at any time it can either collapse or add a new (large) ring. It remains to show that it terminates in quasi-polynomially many steps. First we observe that the algorithm never has more than $O(\log n)$ many rings. By the previous two lemmas we have for every $i$ that

$$
\sum_{i^{\prime}=1}^{i+1}\left|B_{i^{\prime}}\right|=\left|B_{i+1}\right|+\sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right| \geqslant\left|A_{i+1}\right|+\sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right| \geqslant 2 \sum_{i^{\prime}=1}^{i}\left|B_{i^{\prime}}\right| .
$$

Consequently, the number of blocking trees grows exponentially and the number of rings can be only logarithmic. Now consider the potential

$$
\left(\left|B_{1}\right|,\left|B_{2}\right|, \ldots,\left|B_{i}\right|, \infty\right),
$$

where $i$ is the number of rings. This potential decreases lexicographically with every collapse or newly added ring. Number of possible potentials is bounded by $n^{O(\log n)}$, which implies that the local search terminates in quasi-polynomial time.

### 3.6 APX-HARDNESS OF MAXMIN ARBORESCENCES

We finish this chapter by showing that the MaxMin arborescence problem is APX-hard, even in a layered graph with 2 layers and a single source (hence a stronger statement than Theorem 2.2). Our proof is based on a reduction from the max- $k$-cover problem, in which one is given a universe $\mathcal{U}$ of $m$ elements, a family of subsets $S_{1}, S_{2}, \ldots, S_{n}$ of subsets of $\mathcal{U}$ and one has to cover a maximum number of elements of $\mathcal{U}$ using at most $k$ sets. In a seminal result by Feige [38], it is showed that it is NP-hard to approximate max- $k$-cover within factor better than $\frac{e}{e-1}$. More specifically, it is showed that on instances where all sets have the same cardinality $m / k$, it is NP-hard to decide whether there exist $k$ disjoint sets that cover the whole universe $\mathcal{U}$ (YES instance) or if any $k$ sets cover at most $(1-1 / e) \cdot m$ elements (NO instance).

Theorem 3.24. For any $\varepsilon>0$ there is no $\left(\sqrt{\frac{e}{e-1}}-\varepsilon\right)$-approximation algorithm to the single source MaxMin Arborescences problem on 2-layered graphs, unless $P=N P$.

Proof. From a max- $k$-cover instance $\left(\mathcal{U}, S_{1}, S_{2}, \ldots, S_{n}\right)$ we make a layered instance of max-min arborescence as follows. In layer $L_{0}$ there is a single source $s$. In layer $L_{1}$, there are $m / k^{2}$ vertices $\left(v_{i, j}\right)_{1 \leqslant j \leqslant m / k^{2}}$ for each set $S_{i}$. In layer $L_{2}$, there are $m / k^{2}$ sinks $\left(s_{a, j}\right)_{1 \leqslant j \leqslant m / k^{2}}$ for each element $a \in \mathcal{U}$.

Then we connect the source $s$ to all vertices in layer $L_{1}$. Each vertex $v_{i, j}$ in $L_{1}$ is connected to all sinks $s_{a, j}$ such that $a \in S_{i}$. Let us denote by OPT the optimum max-min degree of an arborescence on this instance.

Claim 3.25. If the max- $k$-cover instance is a YES instance, then $O P T \geqslant m / k$. Otherwise, OPT $\leqslant \sqrt{1-1 / e} \cdot m / k$.

Proof. If the max- $k$-cover instance is a YES instance then for the source we select as neighbors in $L_{1}$ all the vertices $v_{i, j}$ such that $S_{i}$ is selected in the $k$-cover solution. The source gets in this way $k \cdot m / k^{2}=m / k$ outgoing neighbors. Each vertex $v_{i, j}$ that was selected then selects all the sinks $s_{a, j}$ such that $a \in S_{i}$. Since we are in a YES instance, all the selected sets are disjoint and each one receives exactly $m / k$ sinks. This is a valid arborescence of value $m / k$.

In the NO case, consider any valid arborescence of value OPT. If there is a $j$ such that the source selects more than $(1+\beta) \cdot k$ vertices (for any $\beta \geqslant 0$ ) in the set $V_{j}=\left(v_{i, j}\right)_{1 \leqslant i \leqslant m}$ then this arborescence cannot have a maxmin degree more than

$$
\max _{\beta \geqslant 0} \min \left\{\frac{1-1 / e+\beta}{1+\beta}, \frac{1}{1+\beta}\right\} \cdot \frac{m}{k} \leqslant \frac{1}{1+1 / e} \cdot \frac{m}{k} \leqslant(\sqrt{1-1 / e}) \cdot \frac{m}{k}
$$

To see this, note that in a NO instance, any union of $k$ sets contains at most ( $1-1 / e) \cdot m$ elements, hence any union of $(1+\beta) \cdot k$ sets contains at most ( $1-1 / e$ ) $\cdot m+\beta m$ elements (recall that each set has cardinality $m / k$ ), to be shared between $(1+\beta) \cdot k$ vertices. This gives the first ratio. The second comes from the fact that any union of some sets contains at most $m$ elements. Optimizing over all $\beta \geqslant 0$ gives the upper-bound. Hence assume there is not such $j$.
If the source $s$ has outdegree more than $\sqrt{1-1 / e} \cdot m / k$ then it must be that there is a $j \in\left[1, m / k^{2}\right]$ such that $s$ has more than

$$
\frac{\sqrt{1-1 / e} \cdot m / k}{m / k^{2}}=\sqrt{1-1 / e} \cdot k
$$

out-neighbors in the set $V_{j}=\left(v_{i, j}\right)_{1 \leqslant i \leqslant m}$. By assumption, the source has also at most $k$ out-neighbors in this set hence the number of sinks connected to vertices in $V_{j}$ is at most $(1-1 / e) \cdot m$ (recall we are in a NO case). Hence there must be a vertex $v_{i, j}$ that gets less than

$$
\frac{(1-1 / e) \cdot m}{\sqrt{1-1 / e} \cdot k}=(\sqrt{1-1 / e}) \cdot \frac{m}{k}
$$

sinks as neighbors. In all cases, there must be a vertex in the arborescence that gets out-degree at most $(\sqrt{1-1 / e}) \cdot \frac{m}{k}$, which ends the proof of the claim.

By the previous claim, a $\left(\sqrt{\frac{e}{e-1}}-\varepsilon\right)$-approximation would be able to distinguish between the YES and NO instances of the max- $k$-cover problem, which is NP-hard.

This chapter contains the proof of Theorem 2.3, starting by an introduction to our techniques. This chapter is based on the publication "The Submodular Santa Claus in the Restricted Assignment Case", a joint work with Lars Rohwedder and Paritosh Garg that appeared at the International Colloquium on Automata, Languages and Programming (ICALP '21) [11]. Our way to the result is organised as follows. In Section 4.2 , we first reduce our problem to a hypergraph matching problem (see next paragraph for a formal definition). We then solve this problem using Lovasz Local Lemma (LLL) in Section 4.3. In [14] the authors also reduce to a hypergraph matching problem which they then solve using LLL, although both parts are substantially simpler. The higher generality of our utility functions is reflected in the more general hypergraph matching problem. Namely, our problem is precisely the weighted variant of the (unweighted) problem in [14]. We will elaborate later why the previous techniques do not easily extend to the weighted variant.

### 4.1 OVERVIEW OF PREVIOUS TECHNIQUES AND OUR NEW IDEAS

After the reduction in Section 4.2 we arrive at the following problem. There is a hypergraph $\mathcal{H}=(P \cup R, \mathcal{C})$ with hyperedges $\mathcal{C}$ over the vertices $P$ and $R$. We write $m=|P|$ and $n=|R|$. We will refer to hyperedges as configurations, the vertices in $P$ as players and $R$ as resources ${ }^{1}$. Moreover, a hypergraph is said to be regular if all vertices in $P$ and $R$ have the same degree, that is, they are contained in the same number of configurations. The hypergraph may contain multiple copies of the same configuration. Each configuration $C \in \mathcal{C}$ contains exactly one vertex in $P$, that is, $|C \cap P|=1$. Additionally, for each configuration $C \in \mathcal{C}$ the resources $j \in C$ have weights $w_{j, C} \geqslant 0$. We emphasize that the same resource $j$ can be given different weights in two different configurations, that is, we may have $w_{j, C} \neq w_{j, C^{\prime}}$ for two different configurations $C, C^{\prime}$.

We require to select for each player $i \in P$ one configuration $C$ that contains $i$. For each configuration $C$ that was selected we require to assign a subset of the resources in $C$ which has a total weight of at least $(1 / \alpha)$. $\sum_{j \in C} w_{j, C}$ to the player in $C$. A resource can only be assigned to one player. We call such a solution an $\alpha$-relaxed perfect matching. One seeks to minimize $\alpha$.

We show that every regular hypergraph has an $\alpha$-relaxed perfect matching for some $\alpha=O(\log \log (n))$ assuming that $w_{j, C} \leqslant(1 / \alpha) \cdot \sum_{j^{\prime} \in C} w_{j^{\prime}, C}$ for all $j, C$, that is, all weights are small compared to the total weight of the configuration. Moreover, we can find such a matching in randomized polynomial time. In the reduction we use this result to round a certain LP

1 We note that these do not have to be the same players and resources as in the Santa Claus problem we reduced from, but $n$ and $m$ do not increase.
relaxation and $\alpha$ essentially translates to the approximation rate. This result generalizes that of Bansal and Srividenko [14] on hypergraph matching in the following way. They proved the same result for unit weights and uniform hyperedges, that is, $w_{j, C}=1$ for all $j, C$ and all hyperedges have the same number of resources ${ }^{2}$. In the next paragraph we briefly go over the techniques to prove our result for the hypergraph matching problem.
our techniques. Already the extension from uniform to non-uniform hypergraphs (assuming unit weights) is highly non-trivial and captures the core difficulty of our result. Indeed, we show with a (perhaps surprising) reduction, that we can reduce our weighted hypergraph matching problem to the unweighted (but non-uniform) version by introducing some bounded dependencies between the choices of the different players. For sake of brevity we therefore focus in this section on the unweighted non-uniform variant, that is, we need to assign to each player a configuration $C$ and at least $|C| / \alpha$ resources in $C$. We show that for any regular hypergraph there exists such a matching for $\alpha=O(\log \log (n))$ assuming that all configurations contain at least $\alpha$ resources and we can find it in randomized polynomial time. Without the assumption of uniformity the problem becomes significantly more challenging. To see this, we lay out the techniques of Bansal and Srividenko that allowed them to solve the problem in the uniform case. We note that for $\alpha=O(\log (n))$ the statement is easy to prove: We select for each player $i$ one of the configurations containing $i$ uniformly at random. Then by standard concentration bounds each resource is contained in at most $O(\log (n))$ of the selected configurations with high probability. This implies that there is a fractional assignment of resources to configurations such that each of the selected configurations $C$ receives $\lfloor|C| / O(\log (n))\rfloor$ of the resources in $C$. By integrality of the bipartite matching polytope, there is also an integral assignment with this property.

To improve to $\alpha=O(\log \log (n))$ in the uniform case, Bansal and Srividenko proceed as follows. Let $k$ be the size of each configuration. First they reduce the degree of each player and resource to $O(\log (n))$ using the argument above, but taking $O(\log (n))$ configurations for each player. Then they sample uniformly at random $O(n \log (n) / k)$ resources and drop all others. This is sensible, because they manage to prove the (perhaps surprising) fact that an $\alpha$-relaxed perfect matching with respect to the smaller set of resources is still an $O(\alpha)$-relaxed perfect matching with respect to all resources with high probability (when assigning the dropped resources to the selected configurations appropriately). Indeed, the smaller instance is easier to solve: With high probability all configurations have size $O(\log (n))$ and this greatly reduces the dependencies between the bad events of the random experiment above (the event that a resource is contained in too many selected configurations). This allows them to apply Lovász Local Lemma (LLL) in order to show that with positive probability the experiment succeeds for $\alpha=O(\log \log (n))$.
It is not obvious how to extend this approach to non-uniform hypergraphs: Sampling a fixed fraction of the resources will either make the small configura-

[^2]tions empty-which makes it impossible to retain guarantees for the original instance-or it leaves the big configurations big-which fails to reduce the dependencies enough to apply LLL. Hence it requires new sophisticated ideas for non-uniform hypergraphs, which we describe next.

Suppose we are able to find a set $\mathcal{K} \subseteq \mathcal{C}$ of configurations (one for each player) such that for each $K \in \mathcal{K}$ the sum of intersections $\left|K \cap K^{\prime}\right|$ with smaller configurations $K^{\prime} \in \mathcal{K}$ is very small, say at most $|K| / 2$. Then it is easy to derive a 2 -relaxed perfect matching: We iterate over all $K \in \mathcal{K}$ from large to small and reassign all resources to $K$ (possibly stealing them from the configuration that previously had them). In this process every configuration gets stolen at most $|K| / 2$ of its resources, in particular, it keeps the other half. However, it is non-trivial to obtain a property like the one mentioned above. If we take a random configuration for each player, the dependencies of the intersections are too complex. To avoid this we invoke an advanced variant of the sampling approach where we construct not only one set of resources, but a hierarchy of resource sets $R_{0} \supseteq \cdots \supseteq R_{d}$ by repeatedly dropping a fraction of resources from the previous set. We then formulate bad events based on the intersections of a configuration $C$ with smaller configurations $C^{\prime}$, but we write it only considering a resource set $R_{k}$ of convenient granularity (chosen based on the size of $C^{\prime}$ ). In this way we formulate a number of bad events using various sets $R_{k}$. This succeeds in reducing the dependencies enough to apply LLL. Unfortunately, even with this new way of defining bad events, the guarantee that for each $K \in \mathcal{K}$ the sum of intersections $\left|K \cap K^{\prime}\right|$ with smaller configurations $K^{\prime} \in \mathcal{K}$ is at most $|K| / 2$ is still too much to ask. We can only prove some weaker property which makes it more difficult to reconstruct a good solution from it. The reconstruction still starts from the biggest configurations and iterates to finish by including the smallest configurations but it requires a delicate induction where at each step, both the resource set expands and some new small configurations that were not considered before come into play.

A REMARK ON LOCAL SEARCH TECHNIQUES. In our proof, we focus here on an extension of the LLL technique of Bansal and Srividenko. However, another technique proved itself very successful for the Santa Claus problem in the restricted assignment case with a linear utility function. This is a local search technique discovered by Asadpour, Feige, and Saberi [6] who used it to give a non-constructive proof that the integrality gap of the configuration LP of Bansal and Srividenko is at most 4. One may wonder if this technique could also be extended to the submodular case as we did with LLL. Unfortunately, this seems problematic as the local search arguments heavily rely on amortizing different volumes of configurations (i.e., the sum of their resources' weights or the number of resources in the unweighted case). Amortizing the volumes of configurations works well, if each configuration has the same volume, which is the case for the problem derived from linear valuation functions, but not the one derived from submodular functions. If the volumes differ then the amortization arguments break and we believe this is a fundamental problem for this approach.

### 4.2 REDUCTION TO HYPERGRAPH MATCHING PROBLEM

In this section we give a reduction of the restricted submodular Santa Claus problem to the hypergraph matching problem. As a starting point we solve the configuration LP, a linear programming relaxation of our problem. The LP is constructed using a parameter $T$ which denotes the value of its solution. The goal is to find the maximal $T$ such that the LP is feasible. In the LP we have a variable $x_{i, C}$ for every player $i \in P$ and every configuration $C \in \mathcal{C}(i, T)$. The configurations $\mathcal{C}(i, T)$ are defined as the sets of resources $C \subseteq \Gamma_{i}$ such that $f(C) \geqslant T$. We require every player $i \in P$ to have at least one configuration and every resource $j \in R$ to be contained in at most one configuration.

$$
\begin{gathered}
\sum_{C \in \mathcal{C}(i, T)} x_{i, C} \geqslant 1 \quad \text { for all } i \in P \\
\sum_{i \in P} \sum_{C \in \mathcal{C}(i, T): j \in C} x_{i, C} \leqslant 1 \quad \text { for all } j \in R \\
x_{i, C} \geqslant 0
\end{gathered} \quad \text { for all } i \in P, C \in \mathcal{C}(i, T) \text { }
$$

Since this linear program has exponentially many variables, we cannot directly solve it in polynomial time. We will give a polynomial time constant approximation for it via its dual. This is similar to the linear variant in [14], but requires some more work. In their case they can reduce the problem to one where the separation problem of the dual can be solved in polynomial time. In our case even the separation problem can only be approximated. Nevertheless, this is sufficient to approximate the linear program in polynomial time.

Theorem 4.1. The configuration LP of the Restricted Submodular Santa Claus problem can be approximated within a factor of $(1-1 / e) / 2$ in polynomial time.

We defer the proof of this theorem to Appendix B.1.1. Given a solution $x^{*}$ of the configuration LP we want to arrive at the hypergraph matching problem from the introduction such that an $\alpha$-relaxed perfect matching of that problem corresponds to an $O(\alpha)$-approximate solution of the Restricted Submodular Santa Claus problem. Let $T^{*}$ denote the value of the solution $x^{*}$. We will define a resource $j \in R$ as fat if $f(\{j\}) \geqslant T^{*} /(100 \alpha)$.

Resources that are not fat are called thin. We call a configuration $C \in$ $\mathcal{C}(i, T)$ thin, if it contains only thin resources and denote by $\mathcal{C}_{t}(i, T) \subseteq$ $\mathcal{C}(i, T)$ the set of thin configurations. Intuitively in order to obtain an $O(\alpha)-$ approximate solution, it suffices to give each player $i$ either one fat resource $j \in \Gamma_{i}$ or a thin configuration $C \in \mathcal{C}_{t}\left(i, T^{*} / O(\alpha)\right)$. For our next step towards the hypergraph problem we use a technique borrowed from Bansal and Srividenko [14]. This technique allows us to simplify the structure of the problem significantly using the solution of the configuration LP. Namely, one can find a partition of the players into clusters such that we only need to cover one player from each cluster with thin resources. All other players can then be covered by fat resources. Informally speaking, the following lemma
is proved by sampling configurations randomly according to a distribution derived in a non-trivial way from the configuration LP.

Lemma 4.2. Let $\ell \geqslant 12 \log (n)$. Given a solution of value $T^{*}$ for the configuration LP in randomized polynomial time we can find a partition of the players into clusters $K_{1} \cup \cdots \cup K_{k} \cup Q=P$ and multisets of configurations $\mathcal{C}_{h} \subseteq \bigcup_{i \in K_{h}} \mathcal{C}_{t}\left(i, T^{*} / 5\right), h=1, \ldots, k$, such that

1. $\left|\mathcal{C}_{h}\right|=\ell$ for all $h=1, \ldots, k$ and
2. Each small resource appears in at most $\ell$ configurations of $\bigcup_{h} \mathcal{C}_{h}$.
3. given any $i_{1} \in K_{1}, i_{2} \in K_{2}, \ldots, i_{k} \in K_{k}$ there is a matching of fat resources to players $P \backslash\left\{i_{1}, \ldots, i_{k}\right\}$ such that each of these players $i$ gets a unique fat resource $j \in \Gamma_{i}$.

The role of the players $Q$ in the lemma above is that each one of them gets a fat resource for certain. The proof follows closely that in [14]. For completeness we include it in Appendix B.1.2. We are now ready to define the hypergraph matching instance. The vertices of our hypergraph are the clusters $K_{1}, \ldots, K_{k}$ and the thin resources. Let $\mathcal{C}_{1}, \ldots, \mathcal{C}_{k}$ be the multisets of configurations as in Lemma 4.2. For each $K_{h}$ and $C \in \mathcal{C}_{h}$ there is a hyperedge containing $K_{h}$ and all resources in $C$. Let $\left\{j_{1}, \ldots, j_{m}\right\}=C$ ordered arbitrarily, but consistently. Then we define the weights as normalized marginal gains of resources if they are taken in this order, that is,

$$
\begin{aligned}
& w_{j_{i}, C}=\frac{5}{T^{*}} f\left(\left\{j_{i}\right\} \mid\left\{j_{1}, \ldots, j_{i-1}\right\}\right) \\
&=\frac{5}{T^{*}}\left(f\left(\left\{j_{1}, \ldots, j_{i-1}, j_{i}\right\}\right)-f\left(\left\{j_{1}, \ldots, j_{i-1}\right\}\right)\right)
\end{aligned}
$$

This implies that $\sum_{j \in C} w_{j, C} \geqslant 5 f(C) / T^{*} \geqslant 1$ for each $C \in \mathcal{C}_{h}, h=$ $1, \ldots, k$.

Lemma 4.3. Given an $\alpha$-relaxed perfect matching to the instance as described by the reduction, one can find in polynomial time an $O(\alpha)$-approximation to the instance of Restricted Submodular Santa Claus.

Proof. The $\alpha$-relaxed perfect matching implies that each cluster $K_{h}$ gets some small resources $C^{\prime}$ where $C^{\prime} \subseteq C$ for some $C \in \mathcal{C}_{h}$ and $\sum_{j \in C^{\prime}} w_{j, C} \geqslant 1 / \alpha$. By submodularity we have that $f\left(C^{\prime}\right) \geqslant T^{*} /(5 \alpha)$. Therefore we can satisfy one player in each cluster using thin resources and by Lemma 4.2 all others using fat resources.

The proof above is the most critical place in the paper where we make use of the submodularity of the valuation function $f$. We note that since all resources considered are thin resources we have, by submodularity of $f$, the assumption that

$$
w_{j, C} \leqslant \frac{5}{T^{*}} f(\{j\}) \leqslant \frac{5}{T^{*}} \frac{T^{*}}{100 \alpha} \leqslant \frac{5}{100 \alpha} \sum_{j \in C} w_{j, C}
$$

for all $j, C$ such that $j \in C$. This means that the weights are all small enough, as promised in introduction. From now on, we will assume that
$\sum_{j \in C} w_{j, C}=1$ for all configurations $C$. This is without loss of generality, since we can just rescale the weights inside each configuration. This does not hurt the property that all weights are small enough.

### 4.2.1 Reduction to unweighted hypergraph matching

Before proceeding to the solution of this hypergraph matching problem, we first give a reduction to an unweighted variant of the problem. We will then solve this unweighted variant in the next section. First, we note that we can assume that all the weights $w_{j, C}$ are powers of 2 by standard rounding arguments. This only loses a constant factor in the approximation rate. Second, we can assume that inside each configuration $C$, each resource has a weight that is at least a $1 /(2 n)$. Formally, we can assume that $\min _{j \in C} w_{j, C} \geqslant 1 /(2 n)$ for all $C \in \mathcal{C}$. If this is not the case for some $C \in \mathcal{C}$, simply delete from $C$ all the resources that have a weight less than $1 /(2 n)$. By doing this, the total weight of $C$ is only decreased by a factor $1 / 2$ since it looses in total at most a weight of $n \cdot(1 / 2 n)=1 / 2$. (Recall that we rescaled the weights so that $\left.\sum_{j \in C} w_{j, C}=1\right)$.

Hence after these two operations, an $\alpha$-relaxed perfect matching in the new hypergraph is still an $O(\alpha)$-relaxed perfect matching in the original hypergraph. From there we reduce to an unweighted variant of the matching problem. Note that each configuration contains resources of at most $\log (n)$ different possible weights (powers of 2 from $1 /(2 n)$ to $1 / \alpha$ ). We create the following new unweighted hypergraph $\mathcal{H}^{\prime}=\left(P^{\prime} \cup R, \mathcal{C}^{\prime}\right)$. The resource set $R$ remains unchanged. For each player $i \in P$, we create $\log (n)$ players, which later correspond each to a distinct weight. We will say that the players obtained from duplicating the original player form a group. For every configuration $C$ containing player $i$ in the hypergraph $\mathcal{H}$, we add a set $\mathcal{S}_{C}=\left\{C_{1}, \ldots, C_{s}, \ldots, C_{\log (n)}\right\}$ of configurations in $\mathcal{H}^{\prime} . C_{s}$ contains player $i_{s}$ and all resources that are given a weight $2^{-(s+1)}$ in $C$. In this new hypergraph, the resources are not weighted. Note that if the hypergraph $\mathcal{H}$ is regular then $\mathcal{H}^{\prime}$ is regular as well.

Additionally, for a group of player and a set of $\log (n)$ configurations (one for each player in the group), we say that this set of configurations is consistent if all the configurations selected are obtained from the same configuration in the original hypergraph $\mathcal{H}$ (i.e. the selected configurations all belong to $\mathcal{S}_{C}$ for some $C$ in $\mathcal{H}$ ).

Formally, we focus of the following problem. Given the regular hypergraph $\mathcal{H}^{\prime}$, we want to select, for each group of $\log (n)$ players, a consistent set of configurations $C_{1}, \ldots, C_{s}, \ldots, C_{\log (n)}$ and assign to each player $i_{s}$ a subset of the resources in the corresponding configuration $C_{s}$ so that $i_{s}$ is assigned at least $\left\lfloor\left|C_{s}\right| / \alpha\right\rfloor$ resources. No resource can be assigned to more than one player. We refer to this assignment as a consistent $\alpha$-relaxed perfect matching. Note that in the case where $\left|C_{s}\right|$ is small (e.g. of constant size) we are not required to assign any resource to player $i_{s}$.

We finish this section by formally proving that this reduction is sufficient for our purposes. More precisely, we show that we can easily transform a
consistent $\alpha$-relaxed matching in $\mathcal{H}^{\prime}$ into a good matching in the original hypergraph $\mathcal{H}$.

Lemma 4.4. A consistent $\alpha$-relaxed matching in $\mathcal{H}^{\prime}$ induces a $O(\alpha)$-relaxed matching in $\mathcal{H}$.

Proof. Let us consider a group of $\log (n)$ players $i_{1}, \ldots, i_{s}, \ldots, i_{\log (n)}$ in $\mathcal{H}^{\prime}$ corresponding to a player $i$ in $\mathcal{H}$. These players are assigned a consistent set of configurations $C_{1}, \ldots, C_{s}, \ldots, C_{\log (n)}$ that correspond to a partition of a configuration in $\mathcal{H}$. Moreover, each player $i_{s}$ is assigned $\left\lfloor\left|C_{s}\right| / \alpha\right\rfloor$ resources from $C_{s}$. We have two cases. If $\left|C_{s}\right| \geqslant \alpha$ then we have that $i_{s}$ is assigned at least

$$
\left\lfloor\left|C_{s}\right| / \alpha\right\rfloor \geqslant\left|C_{s}\right| /(2 \alpha)
$$

resources from $C_{s}$. On the other hand, if $\left\lfloor\left|C_{s}\right| / \alpha\right\rfloor=0$ then the player $i_{s}$ might not be assigned anything. However, we claim that that the configurations $C_{s}$ of cardinality less than $\alpha$ can represent at most a $1 / 5$ fraction of the total weight of the configuration $C$ in the original weighted hypergraph. To see this note that the total weight they represent is upper bounded by

$$
\alpha\left(\sum_{k=\log (100 \alpha / 5)}^{\infty} \frac{1}{2^{k}}\right)=\alpha\left(\frac{5}{100 \alpha} \sum_{k=0}^{\infty} \frac{1}{2^{k}}\right) \leqslant \frac{10}{100}=\frac{1}{10} \sum_{j \in C} w_{j, C}
$$

Hence, the consistent $\alpha$-relaxed matching in $\mathcal{H}^{\prime}$ induces in a straightforward way a matching in $\mathcal{H}$ where every player gets at least a fraction $1 /(2 \alpha) \cdot(1-1 / 10) \geqslant 1 /(3 \alpha)$ of the total weight of the appropriate configuration. This means that the consistent $\alpha$-relaxed perfect matching in $\mathcal{H}^{\prime}$ is indeed a (3 3 )-relaxed perfect matching in $\mathcal{H}$.

### 4.3 MATCHINGS IN REGULAR HYPERGRAPHS

In this section we solve the hypergraph matching problem we arrived to in the previous section. For convenience, we give a self contained definition of the problem before formulating and proving our result.

InPUT: We are given $\mathcal{H}=(P \cup R, \mathcal{C})$ a hypergraph with hyperedges $\mathcal{C}$ over the vertices $P$ (players) and $R$ (resources) with $m=|P|$ and $n=|R|$. As in previous sections, we will refer to hyperedges as configurations. Each configuration $C \in \mathcal{C}$ contains exactly one vertex in $P$, that is, $|C \cap P|=1$. The set of players is partitioned into groups of size at most $\log (n)$, we will use $A$ to denote a group. These groups are disjoint and contain all players. Finally there exists an integer $\ell$ such that for each group $A$ there are $\ell$ consistent sets of configurations. A consistent set of configurations for a group $A$ is a set of $|A|$ configurations such that all players in the group appear in exactly one of these configurations. We will denote by $\mathcal{S}_{A}$ such a set and for a player $i \in A$, we will denote by $\mathcal{S}_{A}^{(i)}$ the unique configuration in $\mathcal{S}_{A}$ containing $i$. Finally, no resource appears in more than $\ell$ configurations. We say that the hypergraph is regular (although some resources may appear in less than $\ell$ configurations).
output: We wish to select a matching that covers all players in $P$. More precisely, for each group $A$ we want to select a consistent set of configurations (denoted by $\left\{\mathcal{S}_{A}^{(i)}\right\}_{i \in A}$ ). Then for each player $i \in A$, we wish to assign a subset of the resources in $\mathcal{S}_{A}^{(i)}$ to the player $i$ such that:

1. No resource is assigned to more than one player in total.
2. For any group $A$ and any player $i \in A$, player $i$ is assigned at least $\left\lfloor\left|\mathcal{S}_{A}^{(i)}\right| / \alpha\right\rfloor$ resources from $\mathcal{S}_{A}^{(i)}$.

We call this a consistent $\alpha$-relaxed perfect matching. Our goal in this section will be to prove the following theorem.

Theorem 4.5. Let $\mathcal{H}=(P \cup R, \mathcal{C})$ be a regular (non-uniform) hypergraph where the set of players is partitioned into groups of size at most $\log (n)$. Then we can, in randomized polynomial time, compute a consistent $\alpha$-relaxed perfect matching for $\alpha=O(\log \log (n))$.

We note that Theorem 4.5 together with the reduction from the previous section will prove our main result (Theorem 2.3) stated in introduction.

### 4.3.1 Overview and notations

To prove Theorem 4.5, we introduce the following notations. Let $\ell \in \mathbb{N}$ be the regularity parameter as described in the problem input (i.e. each group has $\ell$ consistent sets and each resource appears in no more than $\ell$ configurations). As we proved in Lemma 4.2 we can assume with standard sampling arguments that $\ell=300.000 \log ^{3}(n)$ at a constant loss. If this is not the case because we might want to solve the hypergraph matching problem by itself (i.e. not obtained by the reduction in Section 4.2), the proof of Lemma 4.2 can be repeated in a very similar way here.
For a configuration $C$, its size will be defined as $|C \cap R|$ (i.e. its cardinality over the resource set). For each player $i$, we denote by $\mathcal{C}_{i}$ the set of configurations that contain $i$. We now group the configurations in $\mathcal{C}_{i}$ by size: We denote by $\mathcal{C}_{i}^{(0)}$ the configurations of size in $\left[0, \ell^{4}\right)$ and for $k \geqslant 1$ we write $\mathcal{C}_{i}^{(k)}$ for the configurations of size in $\left[\ell^{k+3}, \ell^{k+4}\right)$. Moreover, define $\mathcal{C}^{(k)}=\bigcup_{i} \mathcal{C}_{i}^{(k)}$ and $\mathcal{C}^{(\geqslant k)}=\bigcup_{h \geqslant k} \mathcal{C}^{(h)}$. Let $d$ be the smallest number such that $\mathcal{C}^{(\geqslant d)}$ is empty. Note that $d \leqslant \log (n) / \log (\ell)$. Now consider the following random process.

Random Experiment 4.6. We construct a nested sequence of resource sets $R=R_{0} \supseteq R_{1} \supseteq \ldots \supseteq R_{d}$ as follows. Each $R_{k}$ is obtained from $R_{k-1}$ by deleting every resource in $R_{k-1}$ independently with probability $(\ell-1) / \ell$.

In expectation only a $1 / \ell$ fraction of resources in $R_{k-1}$ survives in $R_{k}$. Also notice that for $C \in \mathcal{C}^{(k)}$ we have that $\mathbb{E}\left[\left|R_{k} \cap C\right|\right]=\operatorname{poly}(\ell)$.

The proof of Theorem 4.5 is organized as follows. In Section 4.3.2, we give some properties of the resource sets constructed by Random Experiment 4.6 that hold with high probability. Then in Section 4.3.3, we show that we can find a single consistent set of configurations for each group of players such that for each configuration selected, its intersection with smaller selected configurations is bounded if we restrict the resource set to an appropriate
$R_{k}$. Restricting the resource set is important to bound the dependencies of bad events in order to apply Lovasz Local Lemma. Finally in Section 4.3.4, we demonstrate how these configurations allow us to reconstruct a consistent $\alpha$-relaxed perfect matching for an appropriate assignment of resources to configurations.

### 4.3.2 Properties of resource sets

In this subsection, we give a precise statement of the key properties that we need from Random Experiment 4.6. The first two lemmas have a straightforward proof. The last one is a generalization of an argument used by Bansal and Srividenko [14]. Since the proof is more technical and tedious, we also defer it to Appendix B. 2 along with the proof of the first two statements.

We start with the first property which bounds the size of the configurations when restricted to some $R_{k}$. This property is useful to reduce the dependencies while applying LLL later.

Lemma 4.7. Consider Random Experiment 4.6 with $\ell \geqslant 300.000 \log ^{3}(n)$. For any $k \geqslant 0$ and any $C \in \mathcal{C}^{(\geqslant k)}$ we have

$$
\frac{1}{2} \ell^{-k}|C| \leqslant\left|R_{k} \cap C\right| \leqslant \frac{3}{2} \ell^{-k}|C|
$$

with probability at least $1-1 / n^{10}$.
The next property expresses that for any configuration the sum of intersections with configurations of a particular size does not deviate much from its expectation. In particular, for any configuration $C$, the sum of it's intersections with other configurations is at most $|C| \ell$ as each resource is in at most $\ell$ configurations. By the lemma stated below, we recover this up to a multiplicative constant factor when we consider the appropriately weighted sum of the intersection of $C$ with other configurations $C^{\prime}$ of smaller sizes where each configuration $C^{\prime} \in \mathcal{C}^{(k)}$ is restricted to the resource set $R_{k}$.
Lemma 4.8. Consider Random Experiment 4.6 with $\ell \geqslant 300.000 \log ^{3}(n)$. For any $k \geqslant 0$ and any $C \in \mathcal{C}^{(\geqslant k)}$ we have

$$
\sum_{C^{\prime} \in \mathcal{C}^{(k)}}\left|C^{\prime} \cap C \cap R_{k}\right| \leqslant \frac{10}{\ell^{k}}\left(|C|+\sum_{C^{\prime} \in \mathcal{C}^{(k)}}\left|C^{\prime} \cap C\right|\right)
$$

with probability at least $1-1 / n^{10}$.
We now define the notion of good solutions which is helpful in stating our last property. Let $\mathcal{F}$ be a set of configurations, $\alpha: \mathcal{F} \rightarrow \mathbb{N}, \gamma \in \mathbb{N}$, and $R^{\prime} \subseteq R$. We say that an assignment of $R^{\prime}$ to $\mathcal{F}$ is $(\alpha, \gamma)$-good if every configuration $C \in \mathcal{F}$ receives at least $\alpha(C)$ resources of $C \cap R^{\prime}$ and if no resource in $R^{\prime}$ is assigned more than $\gamma$ times in total.

Below we obtain that given a $(\alpha, \gamma)$-good solution with respect to resource set $R_{k+1}$, one can construct an almost $(\ell \cdot \alpha, \gamma)$-good solution with respect to the bigger resource set $R_{k}$. Informally, starting from a good solution with respect to the final resource set and iteratively applying this lemma would give us a good solution with respect to our complete set of resources.

Lemma 4.9. Consider Random Experiment 4.6 with $\ell \geqslant 300.000 \log ^{3}(n)$. Fix $k \geqslant 0$. Conditioned on the event that the bounds in Lemma 4.7 hold for $k$, then with probability at least $1-1 / n^{10}$ the following holds for all $\mathcal{F} \subseteq \mathcal{C}^{(\geqslant k+1)}$, $\alpha: \mathcal{F} \rightarrow \mathbb{N}$, and $\gamma \in \mathbb{N}$ such that $\ell^{3} / 1000 \leqslant \alpha(C) \leqslant n$ for all $C \in \mathcal{F}$ and $\gamma \in\{1, \ldots, \ell\}:$ If there is $a(\alpha, \gamma)$-good assignment of $R_{k+1}$ to $\mathcal{F}$, then there is a $\left(\alpha^{\prime}, \gamma\right)$-good assignment of $R_{k}$ to $\mathcal{F}$ where

$$
\alpha^{\prime}(C) \geqslant \ell\left(1-\frac{1}{\log (n)}\right) \alpha(C)
$$

for all $C \in \mathcal{F}$. Moreover, this assignment can be found in polynomial time.
Given the lemmata above, by a simple union bound one gets that all the properties of resource sets hold.

### 4.3.3 Selection of configurations

In this subsection, we give a random process that selects one consistent set of configurations for each group of players such that the intersection of the selected configurations with smaller configurations is bounded when considered on appropriate sets $R_{k}$. We will denote $\mathcal{S}_{A}$ the selected consistent set for group $A$ and for ease of notation we will denote $K_{i}=\mathcal{S}_{A}^{(i)}$ the selected configuration for player $i \in A$. For any integer $k$, we write $\mathcal{K}_{i}^{(k)}=\left\{K_{i}\right\}$ if $K_{i} \in \mathcal{C}_{i}^{(k)}$ and $\mathcal{K}_{i}^{(k)}=\emptyset$ otherwise. As for the configuration set, we will also denote $\mathcal{K}^{(k)}=\bigcup_{i} \mathcal{K}_{i}^{(k)}$ and $\mathcal{K}=\bigcup_{k} \mathcal{K}^{(k)}$. The following lemma describes what are the properties we want to have while selecting the configurations. For better clarity we also recall what the properties of the sets $R_{0}, \ldots, R_{d}$ that we need are. These hold with high probability by the lemmata of the previous section.

Lemma 4.10. Let $R=R_{0} \supseteq \ldots \supseteq R_{d}$ be sets of fewer and fewer resources. Assume that for each $k$ and $C \in \mathcal{C}_{i}^{(k)}$ we have

$$
1 / 2 \cdot \ell^{k-h} \leqslant\left|C \cap R_{h}\right| \leqslant 3 / 2 \cdot \ell^{-h}|C|<3 / 2 \cdot \ell^{k-h+4}
$$

for all $h=0, \ldots, k$. Then there exists a selection of one consistent set $\mathcal{S}_{A}$ for each group $A$ such for all $k=0, \ldots, d, C \in \mathcal{C}^{(k)}$ and $j=0, \ldots, k$ then we have

$$
\begin{aligned}
& \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h}\left|K \cap C \cap R_{h}\right| \leqslant \frac{1}{\ell} \sum_{j \leqslant h \leqslant k} \sum_{C^{\prime} \in \mathcal{C}^{(h)}} \ell^{h}\left|C^{\prime} \cap C \cap R_{h}\right| \\
&+1000 \frac{d+\ell}{\ell} \log (\ell)|C|
\end{aligned}
$$

Moreover, this selection of consistent sets can be found in polynomial time.
Before we prove this lemma, we give an intuition of the statement. Consider the sets $R_{1}, \ldots, R_{d}$ constructed as in Random Experiment 4.6. Then for $C^{\prime} \in \mathcal{C}^{(h)}$ we have $\mathbb{E}\left[\ell^{h}\left|C^{\prime} \cap C \cap R_{h}\right|\right]=\left|C^{\prime} \cap C\right|$. Hence

$$
\sum_{h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}}|K \cap C|=\mathbb{E}\left[\sum_{h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h}\left|K \cap C \cap R_{h}\right|\right]
$$

Similarly for the right-hand side we have

$$
\begin{aligned}
& \mathbb{E}\left[\frac{1}{\ell} \sum_{j \leqslant h \leqslant k} \sum_{C^{\prime} \in \mathcal{C}^{(h)}} \ell^{h}\left|C^{\prime} \cap C \cap R_{h}\right|+O\left(\frac{d+\ell}{\ell} \log (\ell)|C|\right)\right] \\
& =\frac{1}{\ell} \underbrace{\sum_{j \leqslant h \leqslant k} \sum_{C^{\prime} \in \mathcal{C}^{(h)}}\left|C^{\prime} \cap C\right|}_{\leqslant \ell|C|}+O\left(\frac{d+\ell}{\ell} \log (\ell)|C|\right) \\
& \\
& =O\left(\frac{d+\ell}{\ell} \log (\ell)|C|\right)
\end{aligned}
$$

Hence, the lemma says that each resource in $C$ is roughly covered $O((d+\ell) / \ell \cdot \log (\ell))$ times by smaller configurations.

We now proceed to the proof of Lemma 4.10.
Proof. We perform the following random experiment and show with LLL that there is a positive probability of success.

Random Experiment 4.11. For each group $A$, select one consistent set $\mathcal{S}_{A}$ uniformly at random. Then for each player $i \in A$ set $K_{i}=\mathcal{S}_{A}^{(i)}$.

Given this experiment we can define the following random variables. For all $h=0, \ldots, d$ and $i \in P$ we define

$$
X_{i, C}^{(h)}=\sum_{K \in \mathcal{K}_{i}^{(h)}}\left|K \cap C \cap R_{h}\right| \leqslant \min \left\{3 / 2 \cdot \ell^{4},\left|C \cap R_{h}\right|\right\}
$$

Let $X_{C}^{(h)}=\sum_{i=1}^{m} X_{i, C}^{(h)}$. Then

$$
\mathbb{E}\left[X_{C}^{(h)}\right] \leqslant \frac{1}{\ell} \sum_{C^{\prime} \in \mathcal{C}^{(h)}}\left|C^{\prime} \cap C \cap R_{h}\right| \leqslant\left|C \cap R_{h}\right|
$$

We are now ready to define the bad events on which we will apply the Lovasz Local Lemma. As we will show later, if none of them occur, Lemma 4.10 will hold. For each $k, C \in \mathcal{C}^{(k)}$, and $h \leqslant k$ let $B_{C}^{(h)}$ be the event that

$$
X_{C}^{(h)} \geqslant \begin{cases}\mathbb{E}\left[X_{C}^{(h)}\right]+63\left|C \cap R_{h}\right| \log (\ell) & \text { if } k-5 \leqslant h \leqslant k \\ \mathbb{E}\left[X_{C}^{(h)}\right]+135\left|C \cap R_{h}\right| \log (\ell) \cdot \ell^{-1} & \text { if } h \leqslant k-6\end{cases}
$$

The intuitive reason as to why we define these two different bad events can be summarized as follows. In the case $h \leqslant k-6$, we are counting how many times $C$ is intersected by configurations that are much smaller than $C$. Hence the size of this intersection can be written as a sum of independent random variables of value at most $O\left(\ell^{4}\right)$ which is much smaller than the total size of the configuration $\left|C \cap R_{h}\right|$. Since the random variables are in a much smaller range, Chernoff bounds give much better concentration guarantees and we can afford a very small deviation from the expectation. In the other case, we do not have this property hence we need a bigger deviation to maintain a sufficiently low probability of failure. However, this does not hurt the statement of Lemma 4.10 since we sum this bigger deviation only
a constant number of times. One key idea to be able to apply Lovasz Local Lemma here is also to consider intersection of $C$ with smaller configurations but restricted to a set $R_{h}$ of convenient granularity. One can notice that $\left|C^{\prime} \cap R_{h}\right|=\operatorname{poly}(\ell)$ if $C^{\prime} \in \mathcal{C}^{(h)}$ (by the assumption made in Lemma 4.10). This allows to reduce significantly the dependencies between bad events which is crucial to make any use of LLL here.

With this in mind, we claim that the probability of each bad event happening is small.

Claim 4.12. For each $k, C \in \mathcal{C}^{(k)}$, and $h \leqslant k$ we have

$$
\mathbb{P}\left[B_{C}^{(h)}\right] \leqslant \exp \left(-2 \frac{\left|C \cap R_{h}\right|}{\ell^{9}}-18 \log (\ell)\right)
$$

Proof. Consider first the case that $h \geqslant k-5$. By a Chernoff bound (see full version for the precise formulation) with

$$
\delta=63 \frac{\left|C \cap R_{h}\right| \log (\ell)}{\mathbb{E}\left[X_{C}^{(h)}\right]} \geqslant 1
$$

we get

$$
\begin{aligned}
\mathbb{P}\left[B_{C}^{(h)}\right] \leqslant \exp \left(-\frac{\delta \mathbb{E}\left[X_{C}^{(h)}\right]}{3\left|C \cap R_{h}\right|}\right) & \leqslant \exp (-21 \log (\ell))) \\
& \leqslant \exp (-2 \underbrace{\frac{\left|C \cap R_{h}\right|}{\ell^{9}}}_{\leqslant 3 / 2}-18 \log (\ell))
\end{aligned}
$$

Now consider $h \leqslant k-6$. We apply again a Chernoff bound with

$$
\delta=135 \frac{\left|C \cap R_{h}\right| \log (\ell)}{\ell \mathbb{E}\left[X_{C}^{(h)}\right]} \geqslant \frac{1}{\ell}
$$

This implies

$$
\begin{aligned}
\mathbb{P}\left[B_{C}^{(h)}\right] & \leqslant \exp \left(-\frac{\min \left\{\delta, \delta^{2}\right\} \mathbb{E}\left[X_{C}^{(h)}\right]}{3 \cdot 3 / 2 \cdot \ell^{4}}\right) \\
& \leqslant \exp \left(-30 \frac{\left|C \cap R_{h}\right| \log (\ell)}{\ell^{6}}\right) \\
& \leqslant \exp \left(-2 \frac{\left|C \cap R_{h}\right|}{\ell^{9}}-18 \log (\ell)\right)
\end{aligned}
$$

We are now ready to instanciate the Lovász Local Lemma and use it in our setting. Let $k \in\{0, \ldots, d\}, C \in \mathcal{C}^{(k)}$ and $h \leqslant k$. For event $B_{C}^{(h)}$ we set

$$
x\left(B_{C}^{(h)}\right)=\exp \left(-\left|C \cap R_{h}\right| / \ell^{9}-18 \log (\ell)\right) .
$$

We now analyze the dependencies of $B_{C}^{(h)}$. The event depends only on random variables $\mathcal{S}_{A}$ for groups $A$ that contain at least one player $i$ that has
a configuration in $\mathcal{C}_{i}^{(h)}$ which overlaps with $C \cap R_{h}$. The number of such configurations (in particular, of such groups) is at most $\ell\left|C \cap R_{h}\right|$ since the hypergraph is regular.

In each of these groups, we count at $\operatorname{most} \log (n)$ players, each having $\ell$ configurations hence in total at most $\ell \cdot \log (n)$ configurations.

Each configuration $C^{\prime} \in \mathcal{C}^{\left(h^{\prime}\right)}$ can only influence those events $B_{C^{\prime \prime}}^{\left(h^{\prime}\right)}$ where $C^{\prime} \cap C^{\prime \prime} \cap R_{h^{\prime}} \neq \emptyset$. Since $\left|C^{\prime} \cap R_{h^{\prime}}\right| \leqslant 3 / 2 \cdot \ell^{4}$ and since each resource appears in at most $\ell$ configurations, we see that each configuration can influence at most $3 / 2 \cdot \ell^{5}$ events.

Putting everything together, we see that the bad event $B_{C}^{(h)}$ is independent of all but at most

$$
\left(\ell\left|C \cap R_{h}\right|\right) \cdot(\ell \cdot \log (n)) \cdot\left(3 / 2 \cdot \ell^{5}\right)=3 / 2 \cdot \ell^{7} \cdot \log (n)\left|C \cap R_{h}\right| \leqslant\left|C \cap R_{h}\right| \ell^{8}
$$

other bad events.
We can now verify the condition for LLL (see Lemma 2.5) by calculating

$$
\begin{aligned}
& x\left(B_{C}^{(h)}\right) \prod_{\left(B_{C}^{(h)}, B_{C^{\prime}}^{\left(h^{\prime}\right)}\right) \in E}\left(1-x\left(B_{C^{\prime}}^{\left(h^{\prime}\right)}\right)\right) \\
& \geqslant \exp \left(-\left|C \cap R_{h}\right| / \ell^{9}-18 \log (\ell)\right) \cdot\left(1-\ell^{-18}\right)^{\left|C \cap R_{h}\right| \ell^{8}} \\
& \geqslant \exp \left(-\left|C \cap R_{h}\right| / \ell^{9}-18 \log (\ell)\right) \cdot \exp \left(-\left|C \cap R_{h}\right| / \ell^{9}\right) \\
& \geqslant \exp \left(-2\left|C \cap R_{h}\right| / \ell^{9}-18 \log (\ell)\right) \geqslant \mathbb{P}\left[B_{C}^{(h)}\right] .
\end{aligned}
$$

By LLL we have that with positive probability none of the bad events happen. Let $k \in\{0, \ldots, d\}$ and $C \in \mathcal{C}^{(k)}$. Then for $k-5 \leqslant h \leqslant k$ we have

$$
\ell^{h} X_{C}^{(h)} \leqslant \ell^{h} \mathbb{E}\left[X_{C}^{(h)}\right]+63 \ell^{h}\left|C \cap R_{h}\right| \log (\ell) \leqslant \ell^{h} \mathbb{E}\left[X_{C}^{(h)}\right]+95|C| \log (\ell)
$$

Moreover, for $h \leqslant k-6$ it holds that
$\ell^{h} X_{C}^{(h)} \leqslant \ell^{h} \mathbb{E}\left[X_{C}^{(h)}\right]+135 \ell^{h-1}\left|C \cap R_{h}\right| \log (\ell) \leqslant \ell^{h} \mathbb{E}\left[X_{C}^{(h)}\right]+203|C| \log (\ell) \cdot \ell^{-1}$.
We conclude that, for any $0 \leqslant j \leqslant k$,

$$
\begin{aligned}
\sum_{j \leqslant h \leqslant k} & \sum_{K \in \mathcal{K}^{(h)}} \ell^{h}\left|K \cap C \cap R_{h}\right| \\
& \leqslant \sum_{j \leqslant h \leqslant k} \ell^{h} \mathbb{E}\left[X_{C}^{(h)}\right]+1000 \frac{(k-j+1)+\ell}{\ell}|C| \log (\ell) \\
& \leqslant \frac{1}{\ell} \sum_{j \leqslant h \leqslant k} \ell^{h} \sum_{C^{\prime} \in \mathcal{C}^{(h)}}\left|C^{\prime} \cap C \cap R_{h}\right|+1000 \frac{d+\ell}{\ell}|C| \log (\ell)
\end{aligned}
$$

This proves Lemma 4.10.

Remark 4.13. Since there are at most $\operatorname{poly}(n, m, \ell)$ bad events and each bad event $B$ has $\frac{x(B)}{1-x(B)} \leqslant 1 / 2$ (because $x(B) \leqslant \ell^{-18}$ ), the constructive variant of LLL by Moser and Tardos [70] can be applied to find a selection of configurations such that no bad events occur in randomized polynomial time.

In this subsection, we show how all the previously established properties allow us to find, in polynomial time, a good assignment of resources to the configurations $\mathcal{K}$ chosen as in the previous subsection. We will denote as in the previous subsection $\mathcal{K}_{i}^{(k)}=\left\{K_{i}\right\}$ if $K_{i} \in \mathcal{C}_{i}^{(k)}$ and $\mathcal{K}_{i}^{(k)}=\emptyset$ otherwise. We also define $\mathcal{K}^{(k)}=\bigcup_{i} \mathcal{K}_{i}^{(k)}$ and $\mathcal{K}^{(\geqslant k)}=\bigcup_{h \geqslant k} \mathcal{K}^{(k)}$. Finally we define the parameter

$$
\gamma=100.000 \frac{d+\ell}{\ell} \log (\ell)
$$

which will define how many times each resource can be assigned to configurations in an intermediate solution. Note that $d \leqslant \log (n) / \log (\ell)$. By our choice of $\ell=300.000 \log ^{3}(n)$, we have that $\gamma \leqslant 310.000 \log \log (n)$. Lemma 4.10 implies the following bound.
Claim 4.14. For any $k \geqslant 0$, any $0 \leqslant j \leqslant k$, and any $C \in \mathcal{K}^{(k)}$

$$
\sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h}\left|K \cap C \cap R_{h}\right| \leqslant 2000 \frac{d+\ell}{\ell} \log (\ell)|C|
$$

Proof. By Lemma 4.10 we have that

$$
\begin{aligned}
\sum_{j \leqslant h \leqslant k} & \sum_{K \in \mathcal{K}^{(h)}} \ell^{h}\left|K \cap C \cap R_{h}\right| \\
& \leqslant \frac{1}{\ell} \sum_{j \leqslant h \leqslant k} \sum_{C^{\prime} \in \mathcal{C}^{(h)}} \ell^{h}\left|C^{\prime} \cap C \cap R_{h}\right|+1000 \frac{d+\ell}{\ell} \log (\ell)|C|
\end{aligned}
$$

Furthermore, by Lemma 4.8 , we get

$$
\sum_{C^{\prime} \in \mathcal{C}^{(h)}} \ell^{h}\left|C^{\prime} \cap C \cap R_{h}\right| \leqslant \ell^{h} \frac{10}{\ell^{h}}\left(|C|+\sum_{C^{\prime} \in \mathcal{C}^{(h)}}\left|C^{\prime} \cap C\right|\right)
$$

Finally note that each resource appears in at most $\ell$ configurations, hence

$$
\sum_{j \leqslant h \leqslant k} \sum_{\left.C^{\prime} \in \mathcal{C}^{( }\right)}\left|C^{\prime} \cap C\right| \leqslant \ell|C| .
$$

Putting everything together we conclude

$$
\begin{aligned}
\sum_{j \leqslant h \leqslant k} & \sum_{K \in \mathcal{K}^{(h)}} \ell^{h}\left|K \cap C \cap R_{h}\right| \\
& \leqslant \frac{1}{\ell} \sum_{j \leqslant h \leqslant k} \sum_{C^{\prime} \in \mathcal{C}^{(h)}} \ell^{h}\left|C^{\prime} \cap C \cap R_{h}\right|+1000 \frac{d+\ell}{\ell} \log (\ell)|C| \\
& \leqslant \frac{1}{\ell} \sum_{j \leqslant h \leqslant k} 10\left(|C|+\sum_{C^{\prime} \in \mathcal{C}^{(h)}}\left|C^{\prime} \cap C\right|\right)+1000 \frac{d+\ell}{\ell} \log (\ell)|C| \\
& \leqslant \frac{k-j}{\ell} 10|C|+10|C|+1000 \frac{d+\ell}{\ell} \log (\ell)|C| \\
& \leqslant 20|C|+1000 \frac{d+\ell}{\ell} \log (\ell)|C| \\
& \leqslant 2000 \frac{d+\ell}{\ell} \log (\ell)|C| .
\end{aligned}
$$

We can now proceed to the main technical part of this section which is the following lemma proved by induction.

Lemma 4.15. For any $j \geqslant 0$, there exists an assignment of resources of $R_{j}$ to configurations in $\mathcal{K}(\geqslant j)$ such that no resource is taken more than $\gamma$ times and each configuration $C \in \mathcal{K}^{(k)}(k \geqslant j)$ receives at least

$$
\left(1-\frac{1}{\log (n)}\right)^{2(k-j)} \ell^{k-j}\left|C \cap R_{k}\right|-\frac{3}{\gamma} \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-j}\left|K \cap C \cap R_{h}\right|
$$

resources from $R_{k}$.
Before going through the proof, we give here the intuition of why this is what we want to prove. Note that the term $\ell^{k-j}\left|C \cap R_{k}\right|$ is roughly equal to $\ell^{-j}|C|$ by the properties of the resource sets (precisely Lemma 4.7). The second term

$$
\sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}(h)} \ell^{h-j}\left|K \cap C \cap R_{h}\right|
$$

can be shown to be

$$
O\left(\ell^{-j} \frac{d+\ell}{\ell} \log (\ell)|C|\right)=O\left(\ell^{-j} \log \log (n)|C|\right)
$$

by Claim 4.14. Hence by choosing $\gamma$ to be $\Theta(\log \log (n))$ we get that the bound in Lemma 4.15 will be $\Theta\left(\ell^{-j}|C|\right)$. At the end of the induction, we have $j=0$ which indeed implies that we have an assignment in which configurations receive

$$
\Theta\left(\ell^{-0}|C|\right)=\Theta(|C|)
$$

resources and such that each resource is assigned to at most $O(\log \log (n))$ configurations. With this in mind, we give the formal proof of Lemma 4.15 .

Proof. We start from the biggest configurations and then iteratively reconstruct a good solution for smaller and smaller configurations. Recall $d$ is the smallest integer such that $\mathcal{K}^{(\geqslant d)}$ is empty. Our base case for these configurations in $\mathcal{K}(\geqslant d)$ is vacuously satisfied.

Now assume that we have a solution at level $j$, i.e. an assignment of resources to configurations in $\mathcal{K}^{(\geqslant j)}$ such that no resource is taken more than $\gamma$ times and each configuration $C \in \mathcal{K}^{(k)}$ such that $k \geqslant j$ receives at least

$$
\left(1-\frac{1}{\log (n)}\right)^{2(k-j)} \ell^{k-j}\left|C \cap R_{k}\right|-\frac{3}{\gamma} \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-j}\left|K \cap C \cap R_{h}\right|
$$

resources from $R_{j}$. We show that this implies a solution at level $j-1$ in the following way. First by Lemma 4.9 , this implies an assignment of resources of $R_{j-1}$ to configurations in $\mathcal{K}(\geqslant j)$ such that each $C \in \mathcal{K}^{(k)}$ receives at least

$$
\begin{aligned}
& \left(1-\frac{1}{\log (n)}\right) \ell\left[\left(\ell^{k-j}\left(1-\frac{1}{\log (n)}\right)^{2(k-j)}\left|C \cap R_{k}\right|\right.\right. \\
& \left.\quad-\frac{3}{\gamma} \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-j}\left|K \cap C \cap R_{h}\right|\right] \\
& =\left(1-\frac{1}{\log (n)}\right)^{2(k-(j-1))-1} \ell^{k-(j-1)}\left|C \cap R_{k}\right| \\
& \quad-\frac{3}{\gamma}\left(1-\frac{1}{\log (n)}\right) \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-(j-1)}\left|K \cap C \cap R_{h}\right| \\
& \geqslant\left(1-\frac{1}{\log (n)}\right)^{2(k-(j-1))-1} \ell^{k-(j-1)}\left|C \cap R_{k}\right| \\
& \quad-\frac{3}{\gamma} \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-(j-1)}\left|K \cap C \cap R_{h}\right|
\end{aligned}
$$

resources and no resource of $R_{j-1}$ is taken more than $\gamma$ times. Note that we can apply Lemma 4.9 since we have by Claim 4.14 and Lemma 4.7

$$
\begin{aligned}
& \left(1-\frac{1}{\log (n)}\right)^{2(k-j)} \ell^{k-j}\left|C \cap R_{k}\right| \\
& \quad-\frac{3}{\gamma} \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-j}\left|K \cap C \cap R_{h}\right| \\
& \geqslant \frac{\ell^{k-j}}{e^{2}}\left|C \cap R_{k}\right|-\frac{3}{\gamma} 2000 \ell^{-j} \frac{d+\ell}{\ell} \log (\ell)|C| \\
& \geqslant \ell^{-j}|C|\left(\frac{1}{2 e^{2}}-\frac{6000}{\gamma} \frac{d+\ell}{\ell} \log (\ell)\right) \\
& \geqslant \frac{\ell^{-j}|C|}{3 e^{2}}>\frac{\ell^{3}}{1000} .
\end{aligned}
$$

Now consider configurations in $\mathcal{K}^{(j-1)}$ and proceed for them as follows. Give to each $C \in \mathcal{K}^{(j-1)}$ all the resources in $C \cap R_{j-1}$ except all the resources that appear in more than $\gamma$ configurations in $\mathcal{K}^{(j-1)}$. Since each deleted resource is counted at least $\gamma$ times in the sum $\sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1}\right|$, we have that each configuration $C$ in $\mathcal{K}^{(j-1)}$ receives at least

$$
\left|C \cap R_{j-1}\right|-\frac{1}{\gamma} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1}\right|
$$

resources and no resource is taken more than $\gamma$ times by configurations in $\mathcal{K}^{(j-1)}$. Notice that now every resource is taken no more than $\gamma$ times by configurations in $\mathcal{K}^{(\geqslant j)}$ and no more than $\gamma$ times by configurations in $\mathcal{K}^{(j-1)}$ which in total can sum up to $2 \gamma$ times.
Therefore to finish the proof consider an resource $i \in R_{j-1}$. This resource is taken $b_{i}$ times by configurations in $\mathcal{K}(\geqslant j)$ and $a_{i}$ times by configurations
in $\mathcal{K}^{(j-1)}$. If $a_{i}+b_{i} \leqslant \gamma$, nothing needs to be done. Otherwise, denote by $O$ the set of problematic resources (i.e. resources $i$ such that $a_{i}+b_{i}>\gamma$ ). For every $i \in O$, select uniformly at random $a_{i}+b_{i}-\gamma$ configurations in $\mathcal{K}(\geqslant j)$ that currently contain resource $i$ and delete the resource from these configurations. When this happens, each configuration in $C \in \mathcal{K}(\geqslant j)$ that contains $i$ has a probability of $\left(a_{i}+b_{i}-\gamma\right) / b_{i}$ to be selected to loose this resource. Hence the expected number of resources that $C$ looses with such a process is

$$
\mu=\sum_{i \in O \cap C} \frac{a_{i}+b_{i}-\gamma}{b_{i}}
$$

It is not difficult to prove the following claim.
Claim 4.16. For any $C \in \mathcal{K}(\geqslant j)$,

$$
\frac{1}{\gamma^{2}} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right| \leqslant \mu \leqslant \frac{2}{\gamma} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|
$$

Proof. Note that we can write
$\mu=\sum_{i \in O \cap C} \frac{a_{i}+b_{i}-\gamma}{b_{i}} \leqslant \max _{i \in O \cap C}\left\{\frac{a_{i}+b_{i}-\gamma}{a_{i} b_{i}}\right\} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|$.
The reason for this is that each resource $i$ accounts for an expected loss of $\left(a_{i}+b_{i}-\gamma\right) / b_{i}$ while it is counted $a_{i}$ times in the sum

$$
\sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|
$$

Similarly,
$\mu=\sum_{i \in O \cap C} \frac{a_{i}+b_{i}-\gamma}{b_{i}} \geqslant \min _{i \in O \cap C}\left\{\frac{a_{i}+b_{i}-\gamma}{a_{i} b_{i}}\right\} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|$.
Note that by assumption we have that $a_{i}+b_{i}>\gamma$. This implies that either $a_{i}$ or $b_{i}$ is greater than $\gamma / 2$. Assume w.l.o.g. that $a_{i} \geqslant \gamma / 2$. Since by assumption $a_{i} \leqslant \gamma$ we have that

$$
\frac{a_{i}+b_{i}-\gamma}{a_{i} b_{i}} \leqslant \frac{b_{i}}{a_{i} b_{i}}=\frac{1}{a_{i}} \leqslant \frac{2}{\gamma}
$$

In the same manner, since $a_{i}+b_{i}>\gamma$ and that $a_{i}, b_{i} \leqslant \gamma$, we can write

$$
\frac{a_{i}+b_{i}-\gamma}{a_{i} b_{i}} \geqslant \frac{1}{a_{i} b_{i}} \geqslant \frac{1}{\gamma^{2}}
$$

We therefore get the following bounds

$$
\frac{1}{\gamma^{2}} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right| \leqslant \mu \leqslant \frac{2}{\gamma} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|
$$

which is what we wanted to prove.

Assume then that $\mu \leqslant \frac{\left|C \cap R_{k}\right|}{10^{12} \log ^{3}(n)}$. Note that $C$ cannot loose more than $\sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|$ resources in any case. Therefore, by assumption on $\mu$, and since

$$
\mu \geqslant \frac{1}{\gamma^{2}} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right|,
$$

we have that

$$
\begin{aligned}
\sum_{K \in \mathcal{K}^{(j-1)}} \mid K \cap C & \left.\cap R_{j-1} \cap O\left|\leqslant \frac{\gamma^{2}}{10^{12} \log ^{3}(n)}\right| C \cap R_{k} \right\rvert\, \\
& \leqslant \frac{10^{11} \log ^{2} \log (n)}{10^{12} \log ^{3}(n)}\left|C \cap R_{k}\right| \leqslant \frac{1}{\log (n)}\left|C \cap R_{k}\right|
\end{aligned}
$$

Therefore $C$ looses at most $\left|C \cap R_{k}\right| / \log (n)$ resources. Otherwise we have that

$$
\mu>\frac{\left|C \cap R_{k}\right|}{10^{12} \log ^{2}(n)} \geqslant \frac{\ell^{3}}{10^{12} \log ^{3}(n)} \geqslant 200 \log (n)
$$

by Lemma 4.7. Hence noting $X$ the number of deleted resources in $C$ we have that

$$
\mathbb{P}\left(X \geqslant \frac{3}{2} \mu\right) \leqslant \exp \left(-\frac{\mu}{12}\right) \leqslant \frac{1}{n^{10}}
$$

With high probability no configuration looses more than

$$
\frac{3}{2} \mu \leqslant \frac{3}{\gamma} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1} \cap O\right| \leqslant \frac{3}{\gamma} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1}\right|
$$

resources. Hence each configuration $C \in \mathcal{K}(\geqslant j)$ ends with at least

$$
\begin{aligned}
& \left(1-\frac{1}{\log (n)}\right)^{2(k-(j-1))-1} \ell^{k-(j-1)}\left|C \cap R_{k}\right| \\
& \quad-\frac{3}{\gamma} \sum_{j \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-(j-1)}\left|K \cap C \cap R_{h}\right| \\
& \quad-\frac{1}{\log (n)}\left(1-\frac{1}{\log (n)}\right)^{2(k-(j-1))-1} \ell^{k-(j-1)}\left|C \cap R_{k}\right| \\
& \quad-\frac{3}{\gamma} \sum_{K \in \mathcal{K}^{(j-1)}}\left|K \cap C \cap R_{j-1}\right| \\
& \geqslant\left(1-\frac{1}{\log (n)}\right)^{2(k-(j-1))} \ell^{k-(j-1)}\left|C \cap R_{k}\right| \\
& \quad-\frac{3}{\gamma} \sum_{j-1 \leqslant h \leqslant k} \sum_{K \in \mathcal{K}^{(h)}} \ell^{h-(j-1)}\left|K \cap C \cap R_{h}\right|
\end{aligned}
$$

resources which concludes the proof of Lemma 4.15.
Given Lemma 4.15 and the intuition below it, it is straightforward to prove the following corollary which will complete the proof of Theorem 4.5 .

Corollary 4.17. There exists an assignment of resources $R$ to $\mathcal{K}$ such that each configuration $C \in \mathcal{K}$ receives at least $\lfloor|C| /(100 \gamma)\rfloor$ resources. Moreover, this assignment can be found in polynomial time.

Proof. Lemma 4.15 with $k=0$ and Claim 4.14 together imply that we can assign at least

$$
\frac{|C|}{2 e^{2}}-\frac{6000}{100.000}|C| \geqslant \frac{|C|}{100}
$$

resources to every $C \in \mathcal{K}$ such that no resource in $R$ is assigned more than $\gamma$ times. In particular, we can fractionally assign at least $|C| /(100 \gamma)$ resources to each $C \in \mathcal{K}$ such that no resource is assigned more than once. By integrality of the bipartite matching polytope, the corollary follows.

Part II
NETWORK DESIGN PROBLEMS

## INTRODUCTION

In this part, we focus on the design and analysis of simple algorithms for the Matching Augmentation problem (MAP) and the Steiner Forest problem, two fundamental problems in network design. Designing cheap networks that are robust to edge failures is a basic and important problem in the field of approximation algorithms. The area containing these problems is often referred to as survivable network design. Generally, one has to compute the cheapest network that satisfies some connectivity requirements in-between some prespecified set of vertices. Classic examples are for instance the Minimum Spanning Tree problem in which one has to augment the connectivity of a graph from 0 to 1 or related questions such as the Steiner Tree or the Steiner Forest problem. Another type of network design is to build 2-edge connected spanning subgraph (2-ECSS) or multisubgraph (2-ECSM), where one has to augment the connectivity of a graph from 0 to 2 . Unfortunately, most of the problems in this area are NP-hard (or even APX-hard), and what one can hope for is generally to compute an approximate solution in polynomial time. In this part, we will focus on two interesting connectivity problems: the Steiner Forest problem and the Matching Augmentation problem (MAP).

Powerful and versatile techniques such as primal-dual [47, 85] or iterative rounding $[58,66]$ guarantee an approximation within factor 2 for many of these problems (including the Steiner Forest problem and the MAP) but improving on this bound for any connectivity problem is often quite challenging. In addition, the mentioned techniques are already fairly involved while some extremely simple heuristics are still poorly understood.

### 5.1 THE STEINER FOREST PROBLEM

One of the most classic problems in network design is arguably the Steiner Tree problem. Given a weighted graph $G=\left(V, E, w: E \mapsto \mathbb{R}_{+}\right)$and a set $\mathcal{T} \subseteq V$ of terminals, one has to compute the cheapest tree that connects all the terminals. A straightforward generalization of this problem is the socalled Steiner Forest problem in which one is given a set $\mathcal{P}$ of pairs of vertices that are required to be connected. One has to buy the cheapest forest that connects all the pairs of terminals. These two problems are also interesting in the study of online algorithms, which deal with uncertain future. Indeed, the assumption that we know all the instance upfront might not always apply in real life, and it makes sense to relax this assumption. In the online version of the Steiner Tree (or Forest) problem, the terminals in $\mathcal{T}$ (or $\mathcal{P}$ ) are revealed one by one (in a possibly adversarial order), and the algorithm has to connect all previously arrived terminals (or pairs) before seeing the next one. The challenge lies in the fact that the algorithm is not allowed to remove edges that were bought before. By opposition, we refer to the standard model of computation where all the information is known upfront as the offline case.

In a seminal paper, Imase and Waxman [56] introduced the online version of the Steiner Tree problem and provided tight bounds in this scenario by proving that (1) the natural greedy algorithm which simply connects the latest arrived terminal to the closest previously arrived terminal is $O(\log (|\mathcal{T}|))$ competitive ${ }^{1}$ and (2) no algorithm can be better than $\Omega(\log (|\mathcal{T}|))$-competitive. Shortly after, Westbrook and Yan [84] introduced the online Steiner Forest problem. As it is a more general problem, the negative result of Imase and Waxman shows that no algorithm can be better than $\Omega(\log (|\mathcal{P}|))$-competitive. From now on, with a slight abuse of notation, we will use $k$ to denote either the number of terminals $|\mathcal{T}|$ in the online Steiner Tree problem or the number of pairs $|\mathcal{P}|$ in the online Steiner Forest problem. A natural generalization of the greedy algorithm of Imase and Waxman for online Steiner Tree to the case of online Steiner Forest can be described informally as follows:

Upon the arrival of a new pair $\left\{s_{i}, t_{i}\right\}$, connect $s_{i}$ and $t_{i}$ with the shortest path in the current metric, contract the metric along the chosen path and wait for the next pair.

We mention that there are some subtleties about how the metric is contracted exactly, but for the sake of clarity, we will postpone these details to later in the introduction. The reader might think for now that greedy contracts the edges that it selects (i.e. for any edge $e$ selected by greedy, its weight $w(e)$ is set to 0 ). For now, an algorithm will be considered as "greedy" if it always buys the shortest path in the current metric, and nothing else. Compared to the algorithm we introduced in Chapter 1 , the only difference is that here the algorithm cannot choose the order in which the pairs are revealed.
Westbrook and Yan [84] showed that a wide class of greedy algorithms are $O(\sqrt{k} \log (k))$-competitive. This bound was quickly improved by Awerbuch, Azar, and Bartal [8] who showed with an elegant dual fitting argument that greedy algorithms are in fact $O\left(\log ^{2}(k)\right)$-competitive and conjectured that the right bound should be $O(\log (k))$. Since then, their conjecture has remained open. It is quite remarkable that the state-of-the-art for this problem is the same also in the offline case, where greedy can choose the order in which it connects the pairs.

This remark also applies to the lower bound side. The best lower-bound is $\Omega(\log (k))$ both in online and offline settings. In fact, all lower bounds for greedy that appeared so far in the literature (see [3, 8, 23, 56]) are instances where the underlying optimum forest is a single tree. Surprisingly, even in this case, nothing better than the $O\left(\log ^{2}(k)\right)$ upper bound is known. We note that the $O(\log (k))$-competitive analysis of Imase and Waxman for online Steiner Tree does not extend to the online Steiner Forest problem, even if we additionally assume that the optimum is a single tree. Indeed, the solution constructed by greedy may not be a single connected component even if the

[^3]underlying optimum solution is. This problem is highlighted by the lower bounds that appear in [8,23], which showcase the limitations of the current analysis techniques.
further related work. Apart from the two results [8, 84] mentioned above, many papers are related to the problem of understanding the greedy algorithm for Steiner Forest. The competitive ratio of greedy was first mentioned again in a list of open problems by Fiat and Woeginger [40]. Around the same time, Berman and Coulston [16] designed a more complex (nongreedy) algorithm which they showed to be $O(\log (k))$-competitive in the online setting. However, their algorithm is not greedy because it can buy additional edges that could be helpful in the future but are useless right now. Later, Chen, Roughgarden, and Valiant [23] applied the result of Awerbuch et al. regarding the greedy algorithm to design network protocols for good equilibrium. Interestingly, they mention that the non-greedy algorithm of Berman and Coulston was not possible to apply in their setting. More recently, the performance of greedy algorithms for online Steiner Forest was further raised as an "important open problem" in [36] and also cited in [75].

Finally, we note that there are several other situations where the Steiner Forest problem is significantly harder to the understand than the Steiner Tree problem. In the offline case, giving a better-than-2 approximation for Steiner Forest is a major open problem in approximation algorithms, while such a result already exists for the Steiner Tree problem [19, 77, 86]. Similarly, it was known for a very long time that greedy (i.e. compute a minimum spanning tree of the metric completion on terminals) gives a constant factor approximation to the optimum Steiner Tree in the offline case, but it has only been recently proved by Gupta and Kumar [50] that a simple "gluttonous" algorithm also yields a constant factor approximation in the case of offline Steiner Forest (we emphasize that the gluttonous algorithm cannot apply to the online setting, in constrast with the simpler greedy algorithm).

### 5.1.1 Our results

As announced, there are some subtleties about how the metric is contracted when running greedy. Hence we will first define three variants of greedy that contract the metric in slightly different ways. However, we emphasize that the best current upper bound for all these three variants is the $O\left(\log ^{2}(k)\right)$ upper bound of Awerbuch et al. Furthermore, these three variants behave exactly the same on the problematic examples of [8, 23]. In particular, all the discussion so far applies to any of these three variants and our main technical theorem will apply to all three variants of greedy, but we will be able to obtain more specialized results for some variants. In particular, for one of the variants we show that greedy guarantees a $O(\log (k) \cdot \log \log (k))$ approximation in the offline case. Before getting into the precise definition, it is worthwhile to mention that Gupta and Kumar [50] also discussed some subtleties about contracting the metric and also defined several algorithms based on that. Hence it is not the first time that altering the contraction procedure has been considered.

DEFINITION OF THE GREEDY ALGORITHM. After greedy connected the latest arrived pair $p=\{s, t\}$, it is clear that the distance between $s$ and $t$ can be set to 0 in the current metric. This is the only property of the metric contraction that the proof of Awerbuch et al. uses to obtain a good upper bound. As long as we obtain a new graph $G^{\prime}$ such that $d_{G^{\prime}}(s, t)=0$, the $O\left(\log ^{2}(k)\right)$ upper bound applies. In our paper, each greedy algorithm will formally maintain a graph $G^{(\tau)}=\left(V, E^{(\tau)}\right)$ that accurately describes the current metric. In any case, greedy will always take the shortest path in the current metric $G^{(\tau)}$ to connect a newly arrived pair. After connecting the $\tau$-th pair, the set of edges will be defined as $E^{(\tau)}=E \cup S^{(\tau)}$ where $S^{(\tau)}$ is a set of edges of weight 0 over vertex set $V$ that we will call shortcuts. It will be clear from definition that we will have the natural condition that $\emptyset=S^{(0)} \subseteq S^{(1)} \subseteq \ldots \subseteq S^{(k)}$. The cost incurred by greedy will always be the sum of lengths of all the shortest paths taken for connecting the pairs. We now proceed to define the three contraction rules formally.

Rule 1: When greedy connects a pair $\{s, t\}$ through a path $s=$ $v_{0}, v_{1}, v_{2}, \ldots, v_{\ell}, t=v_{\ell+1}$, add the following shortcuts:

- For all $0 \leqslant i \leqslant \ell$ add the edge $\left\{v_{i}, v_{i+1}\right\}$ of weight 0 .

Rule 1 is what Awerbuch et al. intended in their original paper. It can be seen as simply contracting all the edges on the path taken.

Rule 2: When greedy connects a pair $\{s, t\}$ add an edge $\{s, t\}$ with weight 0 .

Rule 2 is actually how the metric is contracted in [50] for their main algorithm. Rule 2 might seem much weaker than Rule 1 as fewer shortcuts are added. One can see that for any $u, v \in V$, the distance between $u$ and $v$ can only be smaller when using Rule 1 over Rule 2 . Hence the cost of greedy equipped with Rule 2 is always an upper bound on the cost of greedy equipped with Rule 1. However, the proof of [8] already applies in this case. Hence the greedy algorithm that uses Rule 2 is already $O\left(\log ^{2}(k)\right)$-competitive.

Rule 3: When greedy connects a pair $\{s, t\}$ through a path $s, v_{1}, v_{2}, \ldots, v_{\ell}, t$, let $s=v_{0}^{\prime}, v_{1}^{\prime}, \ldots, v_{\ell^{\prime}}^{\prime}, t=v_{\ell^{\prime}+1}^{\prime}$ be the subsequence of $s, v_{1}, v_{2}, \ldots, v_{\ell}, t$ in which we keep only the vertices that appeared in a previous pair (i.e. previously arrived terminals). Then, add the following shortcuts:

- For all $0 \leqslant i \leqslant \ell^{\prime}$ add the edge $\left\{v_{i}^{\prime}, v_{i+1}^{\prime}\right\}$ of weight 0 .

Intuitively, Rule 3 is in-between rules 1 and 2. It is also reminiscent of the contraction rule of the second algorithm in [50]. Again, using this rule, it is clear that we obtain shorter paths than with Rule 2. The $O\left(\log ^{2}(k)\right)$ upper bound also applies when using Rule 3. The formal definition of Greedy $i_{i}$ follows naturally for any $i \in\{1,2,3\}$.

```
Algorithm 1 Greedy \(_{i}\)
    Upon arrival of pair \(\{s, t\}\), buy the shortest path in the current metric.
    Update the metric from \(G^{(\tau)}\) to \(G^{(\tau+1)}\) using Rule \(i\) and wait for the next
    pair.
```

As a shorthand, we will denote by A the greedy algorithm at hand. If we do not specify which contraction rule we use, it will implicitly mean that the statement that follows holds for any of our three rules.

OUR RESULTS. First, we introduce an intuitive measure of the efficiency of Greedy $_{i}$. In general, it might be that the cost incurred by the pair $p=\{s, t\}$ is much smaller than $d_{G}(s, t)=d_{G^{(0)}}(s, t)$. Indeed, because of additional shortcuts, it might be that the ratio $d_{G}(s, t) / d_{G^{(\tau)}}(s, t)$ is unbounded when $d_{G^{(\tau)}}(s, t)=0$. We will define this ratio as the contraction of pair $p$ and denote it $\alpha(p)$. We have that

$$
1 \leqslant \alpha(p) \leqslant \infty
$$

for all pairs $p$, in any instance $\mathcal{I}$. Intuitively, a very high contraction means that Greedy ${ }_{i}$ did a good job at reusing edges bought before. Following this remark, one can note that all known lower bounds in [3, 8, 23, 56] have a contraction of exactly 1 for all pairs in the instance (in the case of Steiner Tree instances in $[3,56]$, one can always choose one of the two endpoints of each pair so that this is the case). This seems to confirm the intuitive reasoning that a hard instance should have most of the pairs with low contraction.

After this remark, we use the shorthand $\mathrm{A}=\mathrm{Greedy}_{i}$ to denote the greedy algorithm at hand (using any of our three contraction rules). We denote by $\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \mathcal{P}_{<\alpha}\right)$ the cost incurred by A because of pairs of contraction strictly less than $\alpha$ (i.e. pairs $p$ with $\alpha(p)<\alpha$ ) when running instance $\mathcal{I}$. Note that we do not count the cost incurred by A because of pairs with contraction higher than $\alpha$. Furthermore, we will denote $\operatorname{cost}_{\mathrm{A}}(\mathcal{I})$ the total cost incurred by A. Finally, denote by $w(\mathrm{OPT}(\mathcal{I}))$ the cost of the offline optimum. Our main technical result is the following. We remark that it applies regardless of the ordering of the pairs in $\mathcal{P}$.

Theorem 5.1 (Main technical theorem). Fix a sequence $\left(\alpha_{k}\right)_{k \geqslant 0}$ with $\alpha_{k} \geqslant 1$ for all $k$. Let A be a greedy algorithm running on the instance $\mathcal{I}$ containing $k$ pairs. Then,

$$
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \mathcal{P}_{<\alpha_{k}}\right)=O(\log (k)) \cdot\left(\log \left(\alpha_{k}\right)+\log \log (k)\right) \cdot w(O P T(\mathcal{I}))
$$

As mentioned, this theorem applies to greedy with any of our three contraction rules. This already implies the theorem of Awerbuch et al. as it is straightforward to see that pairs with contraction at least $k$ can make greedy pay at most $O(1) \cdot w(\mathrm{OPT}(\mathcal{I}))$. To see this, simply denote by $\mathcal{P}^{\prime}$ these pairs with contraction more than $k$ and $c$ the greedy cost of the most expensive pair in $\mathcal{P}^{\prime}$. Then greedy pays for those pairs in $\mathcal{P}^{\prime}$ a total cost of at most $k \cdot c$ while $\operatorname{OPT}(\mathcal{I})$ must pay at least $k \cdot c$ to connect the most expensive pair in $\mathcal{P}^{\prime}$. With this observation and plugging in $\alpha_{k}=k$ in our bound, we obtain an upper bound of $O\left(\log ^{2}(k)\right)$ on the competitive ratio.

A consequence of our result is that if one wants to have an $\Omega\left(\log ^{1+\varepsilon}(k)\right)$ lower bound for any fixed $\varepsilon>0$, it must be that the lower bound on the cost incurred by greedy comes from pairs with contraction at least $2^{\Omega\left(\log ^{\varepsilon}(k)\right)}$. This already changes the perspective on how to obtain a stronger lower bound (if it exists) and shows that all previous lower bounds (that have contraction $\alpha_{k}=1$ for all pairs) cannot give anything stronger than $\Omega(\log (k) \cdot \log \log (k))$. The fact that the cost should come from pairs with high contraction seems quite counter-intuitive, and we believe this is strong evidence that the old conjecture of Awerbuch et al. should be true, even in adversarial order. Unfortunately, it is not clear to us how to formalize such an intuition. However, this result still has a number of additional consequences that are interesting. In the following, $w\left(T^{\star}(\mathcal{I})\right)$ denotes the optimum tree solution to instance $\mathcal{I}$ (i.e. we restrict the solution to be a single connected component).

Theorem 5.2. Let A be the greedy algorithm using Rule 3. Then we have that

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I})=O(\log (k) \cdot \log \log (k)) \cdot w\left(T^{\star}(\mathcal{I})\right)
$$

even in adversarial order.
As mentioned in the introduction, even in the case of a single tree spanning the whole graph, nothing better than the general $O\left(\log ^{2}(k)\right)$ was known, and all lower bounds in the literature are instances where the optimum is a single component. We also obtain the following result, which was highlighted in the introduction of this thesis (Chapter 1).

Theorem 5.3 (Main theorem). Let A be the greedy algorithm using contraction rule 3. Denote by $c_{i}$ the cost paid by A when connecting pair $p_{i}=\left\{s_{i}, t_{i}\right\}$ and define $d_{i}=d_{G}\left(s_{i}, t_{i}\right)$ (note that we take the distance in the original graph $G$ ). If either of the two sequences $\left(c_{1}, c_{2}, \ldots, c_{k}\right)$ or $\left(d_{1}, d_{2}, \ldots, d_{k}\right)$ is non-increasing then

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I})=O(\log (k) \cdot \log \log (k)) \cdot w(O P T(\mathcal{I}))
$$

This last result can be derived by using our main theorem, combined with a potential function argument from [50]. This result proves that greedy is an $O(\log (k) \cdot \log \log (k))$-approximation in the offline setting (for instance, we can choose to order the pairs so that the sequence $\left(d_{1}, d_{2}, \ldots, d_{k}\right)$ is non-increasing).
These results are proven in Chapter 6. This chapter is organized as follows. Section 6.1, we present the main obstacles in previous approaches and the key idea underlying our new result. In Section 6.2 we present the proof of our main technical result, Theorem 5.1. In Section 6.3, we present our proof of Theorem 5.2 and Theorem 5.3.

### 5.2 THE MATCHING AUGMENTATION PROBLEM

As mentioned earlier, an important network design problem is the 2-edge connected spanning subgraph (2-ECSS) where has to select the cheapest 2 -edge connected subgraph of a given graph (note that taking two copies of the same edge is not allowed in this problem. This problem is closely related to the famous Traveling Salesman Problem (TSP). As for the Steiner Forest
problem, iterative rounding gives a factor 2 approximation for this problem [58], and it is known that the problem is APX-hard (see for instance [34]). In the case of 2-ECSS, a $118 / 89+\varepsilon<1.321$-approximation is known [44] if the underlying graph $G$ is unweighted (this improves on a the previous bound of $4 / 3[55,79]$ ). However, a similar result for the weighted case has remained elusive, and the best approximation algorithm only guarantee a factor 2 approximation. A prominent special case of the weighted 2-ECSS problem is the so-called Forest Augmentation Problem (FAP). In such instances of 2-ECSS all edge weights are either 0 or 1 (we will refer to edges of cost 0 as light edges and edges of cost 1 as heavy edges). The name stems from the fact that one can assume that the light edges form a forest $F$, and the goal is to find the smallest set of heavy edges $E^{\prime}$ such that $F \cup E^{\prime}$ is 2-edge connected. Independently of our work, a very recent breakthrough [48] gives an 1.9973-approximation algorithm for FAP, breaching for the first time the barrier of 2 in that setting.

A famous special case of FAP is the Tree Augmentation Problem (TAP) which has been extensively studied for decades. In this problem, the forest $F$ is a single spanning tree, and one has to find the smallest set of edges to make the tree 2 -edge connected. For this problem, several better-than-2 approximations were designed in a long line of research [1, 20, 28-30, 32 , $37,41,43,49,61,63,64,72,74,81,82]$. One can see TAP as an extreme case of FAP where the forest is a single component. Another interesting special case is the Matching Augmentation Problem (MAP), in which the forest of light edges forms a matching $M$ and one has to find the smallest set of heavy edges $E^{\prime}$ such that $M \cup E^{\prime}$ is 2 -edge connected. It can be seen as the other extreme case in which the forest forms as many components as possible. We also remark that MAP generalizes the unweighted 2-ECSS problem, which can be viewed as an instance of MAP with an empty matching. For MAP, only recently several better-than- 2 approximation $[26,27,45,48]$ were given. These works culminate in a $13 / 8 \approx 1.625$-approximation for MAP.

For many of these network design problems, there is a simple linear programming relaxation called the cut $L P$. In the case of FAP, for a given graph $G=(V, E)$, forest $F \subseteq E$ the cut LP is written as follows, with a variable $x_{e}$ to decide to take each edge $e$ or not. Recall that $\delta(S)$ denotes the edges with exactly one endpoint in $S$.

$$
\begin{aligned}
L P(G, F): & \min \sum_{e \in E \backslash F} x_{e} \\
& \sum_{e \in \delta(S)} x_{e} \geqslant 2, \quad \text { for all } S, \emptyset \subsetneq S \subsetneq V \\
& 0 \leqslant x_{e} \leqslant 1, \quad \forall e \in E .
\end{aligned}
$$

The integrality gap of this linear program is an interesting question by itself. Recently, in the case of TAP (i.e. $F$ is a spanning tree), Nutov [74] showed that the integrality gap is at most $2-2 / 15 \approx 1.87$. Cheriyan et al. [30] showed that the integrality gap is at least $3 / 2$ in the case of TAP while in the case of MAP, the best upper bound on the integrality gap is 2 , and the best lower bound is $9 / 8[2,79]$. We note that the recent works on MAP [26,
$27,45,48$ ] do not compare against the cut LP, and therefore do not show an integrality gap better than 2 for MAP. Moreover, these algorithms and their analysis are quite involved.

### 5.2.1 Our results

In Chapter 7 of this thesis, we give a simple algorithm that guarantees an approximation ratio $2-c$ (for some absolute constant $c>0$ ) with respect to the best fractional solution of the cut LP. In contrast to the other better-than-2 approximation algorithms, our algorithm is very simple. We note that some of our techniques are reminiscent of the algorithm of Mömke and Svensson [69] for the travelling salesman problem (see also [71, 73] for follow-up works). The algorithm is the following.

## The LP-based algorithm:

1. Compute an optimal extreme point solution $x^{*}$ to $L P(G, M)$.
2. Let $E^{\prime}=\left\{e \in E, x_{e}^{*}>0\right\}$ be the support of $x^{*}$, and run a DFS on the support graph $G^{\prime}=\left(V, E^{\prime}\right)$ which always give priority first to an available light edge and second to the available heavy edge $e$ maximizing $x_{e}^{*}$.
3. Compute an optimum augmentation $A$ to the TAP problem with respect to the DFS tree $T$ computed in the previous step and return $H=T \cup A$.

We note that the LP-based algorithm indeed runs in polynomial time. Step 2 computes a DFS in which some edges are explored in priority (if possible). Step 3 can also be done in polynomial time because the tree $T$ is a DFS tree. This implies that all non-tree edges are back-edges (i.e. one endpoint is an ancestor of the other). In the language of TAP, these edges are often referred to as "uplinks", and it is well-known that TAP instances in which the edges are only "uplinks" are solvable in polynomial time [32, 42]. Finally, the solution given by the algorithm is feasible since Step 2 increases connectivity from 0 to 1 and Step 3 from 1 to 2 . One can check that no edge is taken twice in the process since $A$ and $T$ are disjoint.
In this part of the thesis, our main result shows that this simple algorithm guarantees an approximation within factor strictly better than 2 with respect to the cut LP relaxation.

Theorem 5.4. The LP-based algorithm returns a feasible solution to any MAP instance of cost at most $2-c$ times the cost of the fractional solution $x^{*}$, for some absolute constant $c>0$.

For the sake of exposition, we did not try to optimize the constant $c$ but we believe that improving the ratio of $13 / 8$ in [45] (that holds with respect to the optimum integral solution) would require new techniques in the analysis. Since Nutov [74] proved the integrality gap of the cut LP to be strictly better than 2 for TAP, the cut LP seems a promising relaxation for the general FAP. Additionally, we prove the following simple theorem.

Theorem 5.5. The integrality gap of the cut LP for MAP is at least $4 / 3$.
Proof. Consider the example given in Figure 5.1, which is a simple adaptation of a classic example for the related TSP problem. One can check that the
—_ heavy edge
--- - light edge


Figure 5.1: An integrality gap example.
fractional solution that gives $1 / 2$ fractional value to all heavy edges and value 1 to all light edges is feasible for a total cost of $6 / 2=3$. However, any integral solution costs at least 4 .

In this chapter, we give the proof of Theorem 5.1 in Section 6.2, Theorem 5.2, and Theorem 5.3 in Section 6.3. We first give an overview of the techniques in Section 6.1 before providing the proofs. All the results presented in this chapter are based on the paper "An Improved Analysis of Greedy for Online Steiner Forest", a joint work with Marina Drygala and Andreas Maggiori. It appeared at the ACM-SIAM Symposium on Discrete Algorithms (SODA '22) [9].

### 6.1 THE IDEA BEHIND THE PROOF

In this section, we give the necessary technical details to illustrate our techniques but keeping it rather informal for clarity.

PREVIOUS TECHNIQUES. Before going into our techniques, we will briefly mention the main ingredients of previous proofs. For ease of notation, we will denote by $\mathcal{T}$ the set of terminals that appear in at least one pair of $\mathcal{P}$. For a terminal $u \in \mathcal{T}$, denote by $\bar{u}$ its mate which is the terminal that $u$ should be connected to. Note that we can assume without loss of generality that all terminals have exactly one mate, by duplicating vertices if this is not the case. For ease of presentation, we will assume until the end of the section that for each pair $\{u, \bar{u}\}$ greedy pays exactly the distance of $u$ and $\bar{u}$ in the initial graph, i.e. no previously arrived pair helps greedy in paying less for the newly arrived terminal pair $\{u, \bar{u}\}$. Put otherwise, we assume that the contraction of all the pairs is equal to 1 .

By standard arguments, we can assume that $w(\operatorname{OPT}(\mathcal{I}))=k$ and that for each pair of terminal greedy pays a cost that belongs to the set $\left\{k / 2^{i}\right\}_{1 \leqslant i \leqslant \log (k)}$, with only the loss of a constant factor. Indeed, we can rescale all the edges in the graph $G$ by the factor $k / w(\mathrm{OPT}(\mathcal{I}))$ and for the second assumption, note that by standard geometric grouping, we can assume that greedy pays a cost that belongs to the set $\left\{k / 2^{i}\right\}_{1 \leqslant i \leqslant \infty}$. It is then straightforward to see that the total cost incurred by greedy for pairs cheaper than $k / 2^{\log (k)}=1$ is at most $O(w(\mathrm{OPT}(\mathcal{I}))$. Based on this observation, we will partition $\mathcal{P}$ into disjoint sets $\mathcal{P}^{(i)}$ where each set contains pairs of terminals for which greedy paid $k / 2^{i}$. These sets will be called cost classes. Moreover, in order to introduce the first observation let $B(v, r)=\left\{u \in V \mid d_{G}(v, u)<r\right\}$ be an open ball with center the terminal $u$ and radius $r$. A classic observation is the following.
Observation 6.1.1. Let $\mathcal{B}=\left\{B_{1}, B_{2}, \ldots, B_{\ell}\right\}$ be a collection of balls around terminals, such that (1) all the balls are pairwise disjoint and (2) each ball $B_{i}$ is centered as some terminal $u \in \mathcal{T}$ and its radius $r_{i}$ satisfies $r_{i}<d_{G}(u, \bar{u})$. Then $\sum_{i} r_{i} \leqslant w(O P T(\mathcal{I}))$.

The proof of this observation is straightforward, and we will come back to it later in Section 6.2. These balls can be viewed as a solution to the dual
of the natural linear programming relaxation of the Steiner Forest problem; hence we will refer to these balls as dual balls. We continue by restating (informally) a key lemma in the analysis of [8].

Lemma 6.1 ([8]). For any cost class $\mathcal{P}^{(i)}$ (associated with cost $c_{i}=k / 2^{i}$ ), it is possible to place disjoint dual balls in $G$ such that these balls are centered around terminals that belong to $\mathcal{P}^{(i)}$ and they all have a radius of

$$
\frac{k}{2^{i} \cdot \log \left(\left|\mathcal{P}^{(i)}\right|\right)}
$$

Moreover, at the cost of losing a constant factor we can assume that every pair $p=\{s, t\}$ in $\mathcal{P}^{(i)}$ has at least one dual ball centered around $s$ or $t$.

By taking Lemma 6.1 together with Observation 6.1.1, we obtain that greedy pays at most $O\left(\log \left(\left|\mathcal{P}^{(i)}\right|\right)\right)=O(\log (k))$ times the dual solution for each cost class. This proves that greedy is $O(\log (k))$-competitive for each cost class $\mathcal{P}_{i}$. Hence we see that the previous proof technique has mainly two ingredients:
(1) Partition the set $\mathcal{P}$ into disjoint cost classes $\mathcal{P}^{(i)}$ such that $\mathcal{P}=\bigcup_{i=1}^{\log (k)} \mathcal{P}^{(i)}$ and for each pair in $\mathcal{P}^{(i)}$ greedy paid $k / 2^{i}$.
(2) Prove that greedy is $O(\log (k))$-competitive for each cost class separately by building a dual solution.

By (1) and (2) we get that greedy is $O(\log (k) \cdot \log (k))=O\left(\log ^{2}(k)\right)-$ competitive. Interestingly, in the case of Steiner Tree, it is possible to improve the second step by showing that greedy is $O(1)$-competitive for each cost class, hence the $O(\log (k))$ competitive ratio in general (see [3]). Unfortunately, this is impossible in the case of Steiner Forest. Even if all the pairs have contraction 1 , it might be that greedy is already $\Omega(\log (k))$-competitive for a single cost class (see [8, 23] for an example). We also note that the Berman-Coulston algorithm [16] is designed so that the algorithm is $O(1)$ competitive for each cost class, so the analysis of this more complex algorithm cannot apply to greedy.

OUR NEW APPROACH. The first step of our new proof relies in a different partitioning of the set $\mathcal{P}$. Indeed we will partition $\mathcal{P}$ into $N=\Theta(\log \log (k))$ classes $\tilde{\mathcal{P}}^{(j)}$ such that each class is defined as follows.

$$
\tilde{\mathcal{P}}^{(j)}=\bigcup_{i \equiv j \bmod (N)} \mathcal{P}^{(i)}
$$

Note that we have $\Theta(\log \log (k))$ groups with this partitioning, each containing $\Theta\left(\frac{\log (k)}{\log \log (k)}\right)$ cost classes. Inside each group, the cost classes have the nice property that they are well separated, that is, the multiplicative gap between two consecutive costs is polylog $(k)$. We will make good use of this property to disentangle the interactions between pairs that have different costs. Using these techniques we prove that the competitive ratio of greedy for each set $\tilde{\mathcal{P}}^{(j)}$ is $O(\log (k))$ ending up with a competitive ratio of


Figure 6.1: An example with two cost classes before charging or clustering.


Figure 6.2: An example with two cost classes after charging or clustering. In the top left corner, smaller pairs are much more expensive than the pairs that created the big dual ball while we have the opposite situation in the bottom right corner.
$O(\log (k) \cdot \log \log (k))$ overall. The main technical challenge lies in proving such a result.

If we use Lemma 6.1 to place dual balls around pairs in each cost class $\mathcal{P}^{(i)} \subseteq \tilde{\mathcal{P}}^{(j)}$ (hence creating several collections $\mathcal{B}^{(i)}$ of balls) it might be that two dual balls $B, B^{\prime}$ that belong to different sets $\mathcal{B}^{(i)}$ and $\mathcal{B}^{\left(i^{\prime}\right)}$ overlap. This is the critical issue in the previous proofs, and we will proceed differently. For simplicity, assume we have two cost classes in our set $\tilde{\mathcal{P}}^{(j)}$. Let $c_{i}$ be the cost of the larger class and $c_{i^{\prime}}$ the cost of the smaller class (hence $i^{\prime}>i$ ). We place the dual balls only for the biggest cost class (using Lemma 6.1). Intuitively, the worst case in the analysis will be when all pairs of the smaller cost lie inside the dual balls from the bigger cost class (as depicted in Figure 6.1). If this happens, it will be impossible to place dual balls for the small cost class without intersecting the bigger balls already placed. To overcome this issue, we consider a ball $B$ from the big cost class, and we look at the number of pairs from the small cost class that lie inside this ball. Let the number of such pairs be $k^{\prime}$. We have two cases.
(1) $k^{\prime} \cdot c_{i^{\prime}} \geqslant \operatorname{polylog}(k) \cdot c_{i}$ which is the easy case. In this case, instead of charging the cost of the big pairs to the dual ball $B$ we can instead charge this cost to the smaller pairs inside $B$. By Lemma 6.1, the cost that was initially charged to the ball $B$ was $O\left(c_{i}\right)$. Hence if we evenly distribute this cost among all the small pairs inside $B$, each pair will get a cost of roughly.

$$
O\left(\frac{c_{i}}{k^{\prime}}\right) \leqslant O\left(\frac{c_{i^{\prime}} \cdot c_{i}}{\operatorname{poly} \log (k) \cdot c_{i}}\right)=O\left(\frac{c_{i^{\prime}}}{\operatorname{poly} \log (k)}\right) .
$$

Since the cost was transferred to smaller pairs, we can also safely delete the big dual ball $B$, hence making this space available to place the smaller balls. Note that smaller pairs can be charged at most once in this way because the balls in the dual solution $\mathcal{B}^{(i)}$ for big pairs are pairwise disjoint. This case is depicted in the top left corner of Figure 6.2.
(2) $k^{\prime} \cdot c_{i^{\prime}}<\operatorname{polylog}(k) \cdot c_{i}$. This is the most challenging case and is depicted in the bottom right corner of Figure 6.2. We cannot proceed as in the previous case as we cannot guarantee that small pairs do not get charged too much. Here lies the crux of our proof. First, by re-scaling slightly the ball $B$, we can assume that almost all the small pairs in $B$ are far from the border of $B$ (a pair $\{s, t\}$ is far from the border if one of $s$ or $t$ is at a distance much bigger than $c_{i^{\prime}}$ from the border of $B$ ). For simplicity we assume that all the small pairs are far from the border of $B$. From here we construct an instance $\mathcal{I}^{\prime}$ as follows. Consider the graph $G^{\prime}$ that is induced by vertices inside $B$. The instance $\mathcal{I}^{\prime}$ will be composed of the set $\mathcal{P}^{\prime}$ containing all the pairs of small cost that are inside $B$, and the metric will be the graph $G^{\prime}$. Here the assumption that the contraction is 1 implies that both endpoints of each pair in $\mathcal{P}^{\prime}$ should be inside $B$ (one endpoint well inside the ball and one outside would cost too much). It also implies that greedy behaves exactly the same for the pairs in $\mathcal{P}^{\prime}$ in instance $\mathcal{I}^{\prime}$ as it was behaving for these
pairs $\mathcal{P}^{\prime}$ in instance $\mathcal{I}$, that is, for each pair $p \in \mathcal{P}^{\prime}$ greedy buys exactly the same path to connect $p$ whether instance $\mathcal{I}$ or $\mathcal{I}^{\prime}$ is running. Recall that we assumed

$$
k^{\prime} \cdot c_{i^{\prime}}<\operatorname{polylog}(k) \cdot c_{i}
$$

hence we have

$$
k^{\prime}<\frac{\operatorname{poly} \log (k) \cdot c_{i}}{c_{i^{\prime}}}
$$

We know by previous results that greedy is $O(\log (k))$-competitive on a single cost class; hence the competitive ratio of greedy on instance $\mathcal{I}^{\prime}$ will be bounded by

$$
\begin{aligned}
O\left(\log \left(k^{\prime}\right)\right)=O(\log ( & \left.\left.\frac{\operatorname{polylog}(k) \cdot c_{i}}{c_{i^{\prime}}}\right)\right) \\
& =O(\log \log (k))+O\left(\log \left(\frac{c_{i}}{c_{i^{\prime}}}\right)\right)
\end{aligned}
$$

Now the crucial question: What is the value of $w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right)$ ? As we defined the graph $G^{\prime}$ now, it is not clear. But because we assumed all the small pairs are far from the border of $B$, we can allow ourselves to modify the metric on the border of $B$ without changing the behavior of greedy. If we consider $V^{\prime}$ the set of vertices that lie exactly on the border of $B$ we will add an edge of length 0 between any pair of vertices in $V^{\prime}$. This does not change the behavior of greedy on instance $\mathcal{I}^{\prime}$ because these extra edges are already too far from the pairs in $\mathcal{I}^{\prime}$ to be used (see Figure 6.2, bottom right corner). The interesting fact is now that

$$
w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right) \leqslant w(\mathrm{OPT}(\mathcal{I}) \cap B)
$$

where we denote by $w(\operatorname{OPT}(\mathcal{I}) \cap B)$ the cost of edges bought by $\operatorname{OPT}(\mathcal{I})$ inside the ball $B$.

These observations suggest a Top-Down approach where we first try to place dual balls around big pairs. Then proceed by the case distinction described above. Then we move to the next cost class but ignoring all the pairs that got into case (2). We repeat this until we reach the bottom of the cost hierarchy. We end up with dual balls that have different radii but are all pairwise disjoint (because we ignored the pairs that were in case (2) of any iteration). During this process, each pair got into case (1) at most $\log (k)$ times, hence the total additional charge is $O\left(\frac{\log (k)}{\operatorname{poly} \log (k)}\right)=O(1)$. It remains to handle all the pairs that were ignored. The idea is now that these pairs can be partitioned into disjoint instances with not too many pairs (recall that we have an upper bound on $k^{\prime}$ in case (2)) and such that the optimum solution is at most what $\operatorname{OPT}(\mathcal{I})$ pays inside the ball that created this instance. For instance in the case of two consecutive cost classes $\mathcal{P}^{(i)}, \mathcal{P}^{\left(i^{\prime}\right)}$ (hence $\left.\frac{c_{i}}{c_{i^{\prime}}}=\operatorname{poly} \log (k)\right)$, the total cost of ignored pairs would be equal to:

$$
\begin{aligned}
\sum_{B \in \mathcal{B}^{(i)}}(O(\log \log (k))+ & \left.O\left(\log \left(\frac{c_{i}}{c_{i^{\prime}}}\right)\right)\right) \cdot w(\mathrm{OPT}(\mathcal{I}) \cap B) \\
& =\sum_{B \in \mathcal{B}^{(i)}} O(\log \log (k)) \cdot w(\mathrm{OPT}(\mathcal{I}) \cap B)
\end{aligned}
$$

But because the pairs in $\mathcal{B}^{(i)}$ are pairwise disjoint we have

$$
\sum_{B \in \mathcal{B}^{(i)}} w(\mathrm{OPT}(\mathcal{I}) \cap B) \leqslant w(\mathrm{OPT}(\mathcal{I}))
$$

Hence in total the ignored pairs cost at most $O(\log \log (k)) \cdot w(\operatorname{OPT}(\mathcal{I}))$ to greedy. Because we have $\Theta\left(\frac{\log (k)}{\log \log (k)}\right)$ cost classes inside a set $\tilde{\mathcal{P}}^{(j)}$, it feels that a $\log (k)$ competitive ratio for this set is now possible. Of course we took two consecutive cost classes so that $\frac{c_{i}}{c_{i^{\prime}}}=\operatorname{poly} \log (k)$ but this is intuitively the worst case in the analysis. All this is formally handled via a delicate induction that is done in Section 6.2.

### 6.2 PROOF OF THEOREM 5.1

This section is devoted to the proof of Theorem 5.1. Recall that this theorem applies to the three variants of greedy as defined in the introduction. Hence in this section, A will denote Greedy ${ }_{i}$ for any $i \in\{1,2,3\}$. This section is organized as follows. In Subsection 6.2.1, we introduce some basic definitions that will be needed. In Subsection 6.2.2, we detail some results implied by previous work as well as some pre-processing of the instance needed for the rest of the proof. Namely, we recall the concept of dual fitting used by [8]. In addition, we pre-process the instance so that the different costs greedy pays upon the arrival of different pairs is well-structured (i.e. there is a geometric grouping and a big gap in-between two consecutive cost classes). In Subsection 6.2.3, we give an overview of the main body of the proof, and finally, in Subsections 6.2.4 and 6.2.5, we finish the proof.

### 6.2.1 Problem definition and notation

We will consider a slightly more general problem than Online Steiner Forest. Formally, we are given a weighted graph $G=(V, E, w)$ with weight function $w: E \mapsto \mathbb{R}_{\geqslant 0}$. Along with graph $G$ we are given an ordered sequence of pairs of vertices $\mathcal{P}=\left\{p_{1}, p_{2}, \ldots, p_{k}\right\}$ revealed one by one, and an ordered sequence of sets $\mathcal{S}=\left\{E_{1}, E_{2}, \ldots, E_{k}\right\}$ of additional weighted edges. These edges will be made available to the online algorithm A over time as follows. Before revealing the first pair $p_{1}$, the set of edges in $E_{1}$ is added in the graph $G$ to form the graph $G_{1}=\left(V, E \cup E_{1}\right)$. These edges in $E_{1}$ will remain available to the greedy algorithm A until the end. Then A buys some path and contracts the metric according to its contraction rule as defined in the Introduction to obtain graph $G_{1}^{\prime}$. Next, before revealing $p_{2}$, we add the edge set $E_{2}$ into the graph $G_{1}^{\prime}$ to obtain the graph $G_{2}^{\prime}$ (hence A updates the metric accordingly). Then A sees the pair $p_{2}$ and so on. In general, if $G^{(\tau)}$ denotes the current metric available to greedy after reading pairs $p_{1}, p_{2}, \ldots, p_{\tau}$, we first add the edges of $E_{\tau+1}$ to the graph $G^{(\tau)}$ and after this A connects the pair $p_{\tau+1}$ via the shortest path contracting the metric according to the chosen contraction rule. We call this variant online Steiner Forest in decreasing metrics. This generalizes the classic Online Steiner Forest which is the special case where $E_{i}=\emptyset$ for all $i$.

The goal is to compare the cost incurred by A on the instance $\mathcal{I}=$ $(G, \mathcal{P}, \mathcal{S})$ to the cost of the optimum Steiner forest in the graph $G$ with pairs $\mathcal{P}$. We insist that the offline optimum is not allowed to use edges from $\mathcal{S}$ while the algorithm A can use these edges in $\mathcal{S}$ after they are revealed to it. We will denote the optimum cost by $w(\operatorname{OPT}(\mathcal{I}))$. The size of an instance $\mathcal{I}$ is the number of pairs in $\mathcal{P}$. It will be denoted $k$ in the following (hence $k=|\mathcal{P}|)$. For each pair $p=\{s, t\}$, we will naturally call the endpoints of the pair $p$ the two vertices $s, t$.

For any subset $S \subseteq \mathcal{P}$, we will denote by $\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, S)$ the cost incurred by algorithm A on instance $\mathcal{I}$ because of pairs in $S$. By a slight abuse of notation, for a single pair $p$, we will denote by $\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)=\operatorname{cost}_{\mathrm{A}}(\mathcal{I},\{p\})$ the cost that A pays upon arrival of $p$.

The contraction of a pair with respect to instance $\mathcal{I}$ and algorithm A will be the ratio of the shortest path distance in-between the two endpoints of the pair in $G$ and the actual cost paid by A for this pair when running instance $\mathcal{I}$. Note that the shortest path is taken in the original graph $G$, without help of edges in $\mathcal{S}$. Formally, if we denote by $\alpha(p)$ the contraction of pair $p=\{s, t\}$, we have

$$
\alpha(p)=\frac{d_{G}(s, t)}{\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)}
$$

with the convention that $\alpha(p)=\infty$ if $\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)=0$. Given a fixed $\alpha \geqslant 1$ and an instance $\mathcal{I}=(G, \mathcal{P}, \mathcal{S})$, we denote by $\mathcal{P}_{<\alpha}$ the set of pairs of $\mathcal{P}$ that have contraction less than $\alpha$ when running A (i.e. the pairs $p$ with $\alpha(p)<\alpha$ ).

For simplicity we will assume that every edge is of weight exactly $\eta$ for some arbitrarily small $\eta>0$. If the graph $G$ does not satisfy this, we subdivide all the edges into chains of smaller edges. Of course, this increases the number of edges and vertices in the graph, but since our competitive ratio is only a function of the number of pairs of terminals, this subdivision will not hurt our analysis. It is also clear that subdividing edges changes neither the optimum solution nor the behavior of the greedy algorithm. It also does not change any of the parameters we just defined above. This assumption will be used for simplicity when constructing balls in the graph; we will assume that no edge in the graph $G$ has an endpoint inside and the other endpoint outside the balls (i.e. edges do not "jump over" the border of any ball).

### 6.2.2 Preliminary results and preprocessing of the instance

We describe here some key concepts that will be useful in the rest of the proof. We first introduce the following definition that gives much more structure to the instance $\mathcal{I}$.

Definition 6.2 (Canonical instance). For any $\alpha, \delta \geqslant 1$, any instance $\mathcal{I}=$ $(G, \mathcal{P}, \mathcal{S})$ of online Steiner Forest in decreasing metrics is said to be $(\alpha, \delta)$ canonical with respect to a greedy algorithm A if the following holds:
(a) There exist some real number $m>0$ such that for any pair $p \in \mathcal{P}$, there exists an integer $1 \leqslant j \leqslant \log (k) / \delta$ such that

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)=m / 2^{j \cdot(\delta+10)}
$$

(i.e. we have some geometric grouping of costs and two consecutive cost classes are separated by a multiplicative factor of at least $2^{\delta+10}$ ). We will say that cost classes are well separated, and define $c_{j}=m / 2^{j \cdot(\delta+10)}$.
(b) All pairs in $\mathcal{P}$ have contraction at most $\alpha$ when running A on instance $\mathcal{I}$. We say that all pairs have low contraction.
(c) For any $i \geqslant 1$, the set of additional edges $E_{i}$ contains exactly one edge with the same endpoints of the pair $p_{i}$ and whose weight is exactly $\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, p_{i}\right)$ (i.e. we can assume A connected the pair $p_{i}$ by simply using the single edge in $E_{i}$ ).

This definition suggests that we partition the set of pairs $\mathcal{P}$ into cost classes $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(j)}, \ldots, \mathcal{P}^{(\log (k) /(\delta+10))}$ where $\mathcal{P}^{(j)}$ is the subset of pairs in $\mathcal{P}$ that cost exactly $m / 2^{(\delta+10) \cdot j}$. Note that there are at most $M \leqslant \log (k) /(\delta+$ 10) distinct cost classes. Given this definition, we first claim the following lemma. Intuitively, the lemma states that worst-case instances can be reduced to $(\alpha, \delta)$-canonical instances (for some big $\delta$ ) at a multiplicative loss of $O(\log \log (k)+\log (\alpha))$.

Lemma 6.3. For any instance $\mathcal{I}$ of size $k$, any greedy algorithm $A$ and any $\alpha \geqslant 1$, there exists an $(\alpha, \delta)$-canonical instance $\mathcal{I}^{\prime}$ of size $k^{\prime} \leqslant k$ such that

$$
\begin{gathered}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}\right) \geqslant \frac{\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \mathcal{P}_{<\alpha}\right)}{O(\log \log (k)+\log (\alpha))} \\
\delta \geqslant 100 \cdot(\log (\alpha)+\log \log (k)), \text { and } \\
w\left(O P T\left(\mathcal{I}^{\prime}\right)\right) \leqslant w(O P T(\mathcal{I}))
\end{gathered}
$$

Proof. By standard geometric grouping arguments we can assume that there are at most $\log (k)$ cost classes $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(\log (k))}$ such that the greedy algorithm A pays a cost of $w(\mathrm{OPT}(\mathcal{I})) / 2^{i-1}$ for all pairs in $\mathcal{P}^{(i)}$. This first transformation already appeared in [75] and loses a constant factor. Then we consider these cost classes but we keep only the pairs of contraction at most $\alpha$. Fix $\delta=100 \cdot(\log \log (k)+\log (\alpha))$ We partition the pairs as follows:

$$
\tilde{\mathcal{P}}^{(j)}=\bigcup_{i \equiv j \bmod (\delta+10)} \mathcal{P}^{(i)}
$$

for all $0 \leqslant j<(\delta+10)$. Since there are $(\delta+10)$ groups, one of them represents at least a fraction $1 /(\delta+10)$ of the total cost. Keep only this group and transform the instance by adding additional edges to $\mathcal{S}$ as follows. Assume that we kept the group $\tilde{\mathcal{P}}^{(j)}$ then index the pairs in $\tilde{\mathcal{P}}^{(j)}$ by order of arrival i.e. $\tilde{\mathcal{P}}^{(j)}=\left\{p_{1}, \ldots, p_{i}, \ldots, p_{k^{\prime}}\right\}$. For each pair $p_{i} \in \tilde{\mathcal{P}}^{(j)}$, we define the set of additional edges $E^{(i)}$ as a single edge whose endpoints are exactly the endpoints of the pair $p_{i}$ and whose length is exactly what A paid for this pair in the original instance $\mathcal{I}$. This formally describes the instance $\mathcal{I}^{\prime}=\left(G, \tilde{\mathcal{P}}^{(j)}, \mathcal{S}\right)$. We claim that for any pair selected, the greedy algorithm A pays exactly the same cost for this pair regardless of which instance $\mathcal{I}$ or $\mathcal{I}^{\prime}$ is running. We can prove this simple fact by induction of the number of pairs already arrived in $\mathcal{I}^{\prime}$. If no pair has arrived this is clear. Now consider the
next pair $p_{i}$ to arrive. Note that when running instance $\mathcal{I}^{\prime}$, a path of length exactly $\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, p_{i}\right)$ is available to connect $p_{i}$ since we added an edge in $E^{(i)}$ of exactly this length connecting the endpoints of $p_{i}$. We claim that there cannot be a shorter path. Indeed, by induction we assumed that A paid the same in instance $\mathcal{I}$ and $\mathcal{I}^{\prime}$ for previously arrived pairs hence it must be that the greedy algorithm A used the additional edges in $\mathcal{S}$ to connect previously arrived pairs. Because A is $\mathrm{Greedy}_{i}$ for some $i \in\{1,2,3\}$, it must be that the shortcuts added by A on instance $\mathcal{I}^{\prime}$ so far are exactly edges of length o with endpoints at the endpoints of pairs arrived before $p_{i}$. Note that when running greedy on instance $\mathcal{I}$, the endpoints of previously arrived pairs must be at distance $o$ when the new pair $p_{i}$ arrives. Hence all the shortcuts available to A when receiving the pair $p_{i}$ in instance $\mathcal{I}^{\prime}$ are also available when receiving the pair $p_{i}$ in instance $\mathcal{I}$. In particular, the shortest path taken by A for pair $p_{i}$ in instance $\mathcal{I}^{\prime}$ can only be longer than the path taken for pair $p_{i}$ in instance $\mathcal{I}$.

One can see that in total we lose a multiplicative factor of at most $O(\delta)$ during the reduction. Finally, it is also clear that $w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right) \leqslant w(\mathrm{OPT}(\mathcal{I}))$ and $k^{\prime} \leqslant k$ since the graph $G$ has not changed and we keep in $\mathcal{I}^{\prime}$ only a subset of the pairs in $\mathcal{I}$. This ends the proof of the lemma.

The rest of the section will be devoted to the proof of the following theorem.
Theorem 6.4. Let A be any greedy algorithm that uses one of our 3 contraction rules. Let $\mathcal{I}$ be an $(\alpha, \delta)$-canonical instance (of size $k$ ) of online Steiner Forest in decreasing metrics. Assume $\delta \geqslant 100 \cdot(\log (\alpha)+\log \log (k))$. Then,

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I}) \leqslant O(\log (k)) \cdot w(O P T(\mathcal{I}))
$$

Note that Theorem 6.4 together with Lemma 6.3 imply Theorem 5.1. To see this, consider any instance $\mathcal{I}$. By losing a multiplicative factor of $O(\log \log (k)+\log (\alpha))$ and only considering the pairs in $\mathcal{P}_{<\alpha}$, we transform the instance $\mathcal{I}$ into an $(\alpha, \delta)$-canonical instance $\mathcal{I}^{\prime}$ using Lemma 6.3. Then we apply Theorem 6.4 on instance $\mathcal{I}^{\prime}$ and the total competitive ratio for pairs in $\mathcal{P}_{<\alpha}$ will be $O(\log (k)) \cdot O(\log \log (k)+\log (\alpha))$ which is exactly what we wanted to prove.
dual fitting. A key technical ingredient in the proof of Theorem 6.4 will be dual fitting, which was also used in $[3,8]$ and is a common technique in competitive analysis. In the case of Steiner Forest, a natural way to do dual fitting without explicitly writing a linear program is to consider a set of balls in the graph $G=(V, E)$. For some vertex $v \in V$ and some radius $r>0$, the ball $B(v, r)$ is the open ball of center $v$ and radius $r$, i.e.

$$
B(v, r)=\left\{u \in V \mid d_{G}(v, u)<r\right\} .
$$

Denote by $\mathcal{T}$ the set of terminals which are the vertices that appear in at least one pair. For a terminal $u \in \mathcal{T}$, denote by $\bar{u}$ its mate which is the terminal that $u$ should be connected to. Note that we can assume without loss of generality that all terminals have a only one mate, by duplicating vertices if this is not the case. Assume we have a collection $\mathcal{B}=\left\{B_{1}, B_{2}, \ldots, B_{\ell}\right\}$ of balls such that:

- All of the balls in $\mathcal{B}$ are pairwise disjoint, and
- Each ball $B_{i}$ is centered at some terminal $u \in \mathcal{T}$ and its radius $r_{i}$ satisfies $r<d_{G}(u, \bar{u})$.

Then if we define $y=\sum_{i} r_{i}$ the sum of radii of these balls it must be that

$$
w(\operatorname{OPT}(\mathcal{I})) \geqslant y
$$

A reason for this is that any feasible solution to the Steiner Forest instance must connect $u$ to $\bar{u}$. If we look at any ball $B_{i} \in \mathcal{B}$, then at its center lies at a terminal $u$, and we know that $\bar{u}$ is not in $B_{i}$. Therefore, to connect $u$ to $\bar{u}$, a feasible solution needs to buy at least a path from the center to the border of $B_{i}$, which will have length at least $r_{i}$. Since all balls are pairwise disjoint, we know that these paths will be disjoint, and we can sum the lower bounds on each ball.

An alternative view of this is that the dual balls can be seen as a feasible solution (because the balls are pairwise disjoint) to the dual of the natural LP relaxation of the Steiner Forest problem. Then by weak duality, we know that any feasible solution has cost at least the cost of the dual. Hence we will also refer to a collection of balls as above as a dual solution. A dual solution is feasible if all the corresponding balls are pairwise disjoint and are all centered around the endpoints of some pairs. Using previous works, we obtain the following lemma, which is a formal and slightly more general version of Lemma 6.1.

Lemma 6.5. Let $\mathcal{I}$ be an instance of online Steiner Forest in decreasing metrics. Consider a cost class $\tilde{\mathcal{P}} \subseteq \mathcal{P}$ of pairs. Let $\mathcal{P}^{\prime} \subseteq \tilde{\mathcal{P}}$ be an arbitrary subset of $\tilde{\mathcal{P}}$. Let $c$ be the constant such that

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)=c
$$

for all $p \in \tilde{\mathcal{P}}$ and A a greedy algorithm. Fix an arbitrary radius

$$
r \leqslant \frac{c}{8 \log (|\tilde{\mathcal{P}}|)}
$$

Then for any such radius $r$ it is possible to construct a feasible dual solution $\mathcal{B}$ such that:
(a) All the balls $B \in \mathcal{B}$ have a radius equal to $r$,
(b) $\left|\mathcal{P}^{\prime}\right| \leqslant 5 \cdot|\mathcal{B}|$,
(c) each pair $p \in \mathcal{P}^{\prime}$ has at most one ball $B \in \mathcal{B}$ (denoted $B(p)$ ) centered around one of its endpoints, and
(d) all balls in $\mathcal{B}$ are centered around endpoints of pairs in $\mathcal{P}^{\prime}$.

Proof. We start the proof by stating Moore's bound. We recall that the girth of a graph is the length of the shortest cycle in that graph.

Theorem 6.6 (Moore's bound, see $[4,18])$. Every graph with at least $2 n^{1+\frac{1}{p}}$ edges has girth at most $2 p$.


Figure 6.3: The girth argument.

We are interested in a subset of pairs $\mathcal{P}^{\prime}$ that all belong to the same cost class $\mathcal{P}^{\prime}$. That is, all the pairs in $\mathcal{P}^{\prime}$ cost the same value $c$. We will place dual balls as desired in the statement of Lemma 6.5. To do this we will try to place a dual ball around the endpoint of a pair $p$ as it arrives. If this is not possible without intersecting previously placed balls, we then skip the pair $p$. It will be clear by construction that the dual solution will be feasible. As we construct the feasible dual solution $\mathcal{B}$, we will maintain an auxiliary graph $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ which will be unweighted. The quantity of interest will be the number of edges in $E^{\prime}$. The vertex set $V^{\prime}$ will correspond to a subset of terminals that appear in the pairs in $\mathcal{P}^{\prime}$.

We proceed as follows. When a pair $p=\{s, t\}$ belonging to $\mathcal{P}^{\prime}$ arrives, we try to place a ball of radius

$$
r=\frac{c}{8 \log \left(\left|\mathcal{P}^{\prime}\right|\right)}
$$

around either one of $s$ or $t$. If one of these two balls can be placed without intersecting balls previously placed in $\mathcal{B}$ we place it. We also add its center (either $s$ or $t$ ) to the vertex set $V^{\prime}$ of the auxiliary graph $G^{\prime}$. Otherwise, if none of these balls can be placed without intersecting balls already placed in $\mathcal{B}$ we do not add any ball. However, we identify one ball $B \in \mathcal{B}$ that intersects with the ball of radius $r$ around $s$ and another ball $B^{\prime} \in \mathcal{B}$ for $t$. Let $s^{\prime}$ and $t^{\prime}$ be their centers. By construction we have that $s^{\prime}, t^{\prime} \in V^{\prime}$. We can add the edge $\left\{s^{\prime}, t^{\prime}\right\}$ in the auxiliary graph. We have the following claim.

Claim 6.7. The girth of $G^{\prime}$ is at least $2 \log \left(\left|\mathcal{P}^{\prime}\right|\right)$.
Proof of Claim 6.7. Suppose that the claim does not hold, and consider a cycle of length $\ell<2 \log \left(\left|\mathcal{P}^{\prime}\right|\right)$ in $G^{\prime}$ as in Figure 6.3. Consider the last edge that was added to the cycle. Suppose it was added because of terminal pair $\{s, t\}$ (in red on Figure 6.3). Each other edge in the cycle must have been created by a previous pair $\left\{s_{m}, t_{m}\right\}$ for $1 \leqslant m<\ell$. By triangle inequality we have, for all $m<\ell$ that

$$
d_{G}\left(t_{m}, s_{m+1}\right) \leqslant d_{G}\left(t_{m}, u\right)+d_{G}\left(u, s_{m+1}\right) \leqslant 4 r
$$

where $u$ is the center of the ball that was close to both $t_{m}$ and $s_{m+1}$. Similarly, $d_{G}\left(t, s_{1}\right) \leqslant 4 r$ and $d_{G}\left(t_{m-1}, s\right) \leqslant 4 r$. Since all pairs $\left\{s_{m}, t_{m}\right\}_{m \leqslant \ell-1}$ arrived before $\{s, t\}$, it must be that the shortest path in the current metric $G^{\prime}$ (when the pair $\{s, t\}$ arrives) satisfies

$$
d_{G^{\prime}}\left(s_{m}, t_{m}\right)=0
$$

(this holds for all three contraction rules). Hence, by going through the terminals $t, s_{1}, t_{1}, s_{2}, t_{2}, \ldots, s_{\ell-1}, t_{\ell-1}, s$, greedy could have paid at most

$$
(4 r) \cdot \ell<(8 r) \cdot \log \left(\left|\mathcal{P}^{\prime}\right|\right) \leqslant c
$$

which is a contradiction on the fact that greedy should always take the shortest path available. Hence the girth in $G^{\prime}$ is at least $2 \log \left(\left|\mathcal{P}^{\prime}\right|\right)$. We note that this whole argument also holds if we are in the case of Steiner Forest in decreasing metrics where additional edges are revealed over time. Indeed, these additional edges cannot make a shortest path longer.

Applying Moore's bound to the graph $G^{\prime}$ (which has at most $\left|\mathcal{P}^{\prime}\right|$ vertices by construction) gives that $\left|E^{\prime}\right|<4\left|V^{\prime}\right| \leqslant 4\left|\mathcal{P}^{\prime}\right|$. Since the cardinality of $\mathcal{P}^{\prime}$ equals the number of vertices $\left|V^{\prime}\right|$ in $G^{\prime}$ plus the number of edges $\left|E^{\prime}\right|$ in $G^{\prime}$, and $\left|V^{\prime}\right|=|\mathcal{B}|$ property (b) of Lemma 6.5 follows. Properties (a), (c) and (d) follow by construction.

We are now ready to start the overview of our main proof.

### 6.2.3 Overview of the proof

Recall that we aim to prove Theorem 6.4. By assumption we have that all pairs have contraction at most $\alpha$ and that the cost classes are well separated. Recall that this means the multiplicative gap between two consecutive cost classes is at least $2^{\delta}$ for some $\delta \geqslant 100 \cdot(\log \log (k)+\log (\alpha))$. We use $M$ to denote the number of cost classes where class $j$ is denoted by $\mathcal{P}^{(j)}$ for $1 \leqslant j \leqslant M$.

The goal will be to construct a feasible dual solution $\mathcal{B}$ that has some special properties. This dual solution will be constructed by taking a subset of the dual balls in each of the $\log (k)$ dual solutions constructed with the technique of Awerbuch, Azar, and Bartal [8]. In the end we will charge a portion of the cost that A pays to the dual solution $\mathcal{B}$. The remaining portion of the cost A pays will be handled by an inductive argument. To this end we will have a charging scheme, charge : $\mathcal{P} \mapsto \mathbb{R}_{+}$that will redistribute $\operatorname{cost}_{A}(\mathcal{I})$ amongst the terminal pairs.

Precisely, the total cost that a pair $p$ carries will be

$$
\operatorname{charge}(p) \cdot \operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)
$$

Hence we see the charge as an additional multiplicative factor on the cost of a given pair (initially, the charge is set to 1 for all pairs).

Note that we might sometimes decrease or increase the charge of a pair or transfer the cost, but we will always make sure that when a charge of a pair $p$ is decreasing, the charge of some other pairs are increasing accordingly so that
no cost is lost. For any set $S$ of terminal pairs, we will let $\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, S$, charge) denote the total charged cost carried by pairs in $S$. Formally,

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, S, \text { charge })=\sum_{p \in S} \operatorname{charge}(p) \cdot \operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)
$$

In the proof the pairs in $\mathcal{P}$ will be classified into three types: surviving pairs, charged pairs, and dangerous pairs. The surviving pairs will contain pairs $p$ such that there is a ball $B(p) \in \mathcal{B}$ centered around one of the endpoints. Intuitively these pairs are good for us since we can charge their cost directly to the dual ball $B(p)$. The other pairs will be by default classified as non-surviving. Non-surviving pairs are further partitioned into two subsets, charged or dangerous pairs. Charged pairs are those pairs that have their charge set to o (i.e. charge $(p)=0$ ). Intuitively, they are also in an excellent situation for us since it means that we were able to transfer their cost entirely to other pairs. We do not need to count them in our total cost anymore. Finally, the dangerous pairs are those pairs have neither a charge equal to 0 nor a dual ball in $\mathcal{B}$ centered at one of the endpoints. These pairs will be handled carefully via an inductive argument since we cannot charge them to the dual solution nor to some other pair. To keep careful track of all these elements, we will store a triple ( $\mathcal{B}$, charge, $D$ ) where charge is the charge function as described above, $\mathcal{B}$ is a feasible dual solution and $D \subseteq \mathcal{P}$ a set of dangerous pairs. The family of dual balls will be a union of subsets of dual balls $\mathcal{B}^{(j)}$ for $1 \leqslant j \leqslant M$. Each of the balls in $\mathcal{B}^{(j)}$ will account for a subset of pairs in $\mathcal{P}^{(j)}$ and have some radius of roughly

$$
r_{j}=\frac{c_{j}}{8 \log \left(k_{j}\right)}
$$

where $c_{j}$ is the cost of pairs in $\mathcal{P}^{(j)}$ and $k_{j}=\left|\mathcal{P}^{(j)}\right|$. This choice of radius is coming from previous work, summarized in Lemma 6.5.

Note that in our procedure, it might be that some pairs are not yet classified into one of the three categories (surviving, charged, or dangerous). However, at the beginning of iteration $j$, all pairs in cost classes $j^{\prime}<j$ will be classified. The procedure contains two main steps.

STEP 1. In this step, we start taking into account interactions in-between cost classes. Informally, we do an iterative procedure from $j=1$ to $M$, where we try to build the dual solution from top to bottom. When we start iteration $j$ of this procedure we have a feasible dual solution composed of $\mathcal{B}^{(1)}, \ldots, \mathcal{B}^{(j-1)}$ (of radius $\Theta\left(r_{j}\right)$ with $r_{j}$ specified above) centered around pairs in $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(j-1)}$. All the pairs in $\mathcal{P}^{(j)}$ that are not yet classified are guaranteed to be far from the dual balls already in place. We then look at pairs in $\mathcal{P}^{(j)}$ that are not yet classified, and build a dual solution $\mathcal{B}^{(j)}$ around these pairs using Lemma 6.5. Because unclassified pairs are far from previously placed balls in $\mathcal{B}^{(1)}, \ldots, \mathcal{B}^{(j-1)}$, it is guaranteed that this new dual solution $\mathcal{B}^{(j)}$ will not overlap with the previous dual solution. We then proceed as follows. For any ball $B \in \mathcal{B}^{(j)}$ that we just added, we let $B_{\mathcal{P}}$ denote the set
of pairs of $\bigcup_{j^{\prime}>j} \mathcal{P}^{\left(j^{\prime}\right)}$ (note that we only consider pairs of smaller cost) such that one of its endpoints is at a distance at most

$$
r \cdot\left(1+\frac{1}{200 \cdot \log ^{2}(K)}\right)
$$

from the center of $B$ (denoted $p(B)$ ). Here we note that for a technical reason, $K$ is only an upper bound on the real number of pairs $k$ (i.e. $K \geqslant k$ ). We will let $\partial B_{\mathcal{P}}$ denote the set of pairs of $B_{\mathcal{P}}$ that have one endpoint at a distance of at least

$$
r \cdot\left(1-\frac{1}{200 \cdot \log ^{2}(K)}\right)
$$

from the center of $B$. These pairs are on the border of $B$ hence the choice of notation. With a similar analogy, we will denote the interior of $B_{\mathcal{P}}$ by $\dot{B}_{\mathcal{P}}$. This is the complement of $\partial B_{\mathcal{P}}$ in $B_{\mathcal{P}}$, i.e.

$$
\stackrel{\circ}{B}_{\mathcal{P}}=B_{\mathcal{P}} \backslash \partial B_{\mathcal{P}}
$$

Then for the current ball $B \in \mathcal{B}^{(j)}$ at hand, centered at an endpoint of $p$, we look at the total charged cost of the pairs in $\circ_{\mathcal{P}}$ and make a case distinction based on this value. If this cost is more than polylog $(K)$ times the charged cost of the pair $p$, then we can safely set the charge of $p$ to 0 , delete the dual ball $B$, and increase the charge of pairs in $B_{\mathcal{P}}$ to account for this lost cost. Note that the charges of pairs in $B_{\mathcal{P}}$ increase by at most a multiplicative $\left.\left(1+\frac{1}{\operatorname{polylog}(K)}\right)\right)$. This is pictured in the top left corner of Figure 6.2. Since the number of cost classes is at most $M<\log (K)$ such accumulation of charges is not a problem (note that the charge of each pair can increase at most once per cost class in this way, since we only charge pairs inside a dual ball and the dual balls in a single cost class are pairwise disjoint).

On the contrary, if the charged cost of pairs in $\stackrel{\circ}{B}_{\mathcal{P}}$ is less than polylog $(K)$ times the charged cost of the pair $p$, we first halve the radius of $B$ to get $B^{\prime}$. If the charged cost of the pairs inside $B_{\mathcal{P}}^{\prime}$ is at most a constant factor times the charged cost of $p$, we classify all the pairs in $B_{\mathcal{P}}^{\prime}$ as charged and charge their cost to the pair $p$. Note that the charge of $p$ only increases by a constant factor when doing this, and this happens at most once per pair $p$ (when we place the dual ball around $p$ ). In addition we update $B$ to be $B^{\prime}$, and we classify the pair $p$ as surviving. If on the other hand, the charge of the pairs inside $B_{\mathcal{P}}^{\prime}$ is larger than a specified constant factor times the charge of $B$, we scale the radius of $B^{\prime}$ up until we reach a point where most of the $\operatorname{cost}$ in $B_{\mathcal{P}}^{\prime}$ is carried by $\stackrel{\circ}{\mathcal{P}}_{\prime}$ and not $\partial B_{\mathcal{P}}^{\prime}$. Then we mark all the pairs in $B_{\mathcal{P}}^{\prime}$ as dangerous and add them to the set $D$. We update $B$ to be the ball $B^{\prime}$ and classify $p$ as surviving. This case is pictured in the bottom right corner of Figure 6.2. Note that if the ball $B$ is not deleted, then all the pair in $B_{\mathcal{P}}$ will be classified as either charged or dangerous. In particular, we will never try to place a dual ball around these pairs in the following iterations. We do this procedure for all the balls $B \in \mathcal{B}^{(j)}$ and then move to iteration $(j+1)$. This step is handled in Subsection 6.2.4.

STEP 2. After Step 1, we end with a feasible dual solution $\mathcal{B}$ consisting of the balls placed around surviving pairs, a charge function, and a set $D$ of dangerous pairs. Additionally, we guarantee that no pair is overcharged.

The pairs that are not dangerous are easily accounted for by the dual solution $\mathcal{B}$. Indeed, the surviving pairs still have a dual ball around an endpoint, and the charged pairs have their cost entirely redistributed to other pairs. The only problem might come from dangerous pairs. However, because of how we constructed the dual solution $\mathcal{B}$ and the set $D$, we will be able to cluster the dangerous pairs into disjoint sub-instances that are contained in dual balls corresponding to bigger cost classes. These instances are disjoint, and the crux of the argument is to show a statement of the form:

If the greedy algorithm were to run separately on each sub-instance, then the cost greedy would pay for these pairs would be the same cost that it was paying for these pairs in the bigger instance $\mathcal{I}$.

Hence we can argue that the total cost incurred for dangerous pairs is at most the sum of costs paid by greedy on each sub-instance separately. This helps because we only put pairs in $D$ in the case that their charged cost was bounded by polylog $(K)$ times the charged cost of the pair that created the ball $B$ that contains them. As a result, we have a strong upper bound on the number of pairs $k^{\prime}$ in each smaller sub-instance $\mathcal{I}^{\prime}$. To finish the proof, we need to bound the cost of the offline optimum for each sub-instance $\mathcal{I}^{\prime}$. We note that because all the pairs in $D$ are in the interior of $B$ (i.e. far from the border of $B$ ), we can modify the metric of the graph $G$ at the border of $B$. This will not change the behavior of greedy for the pairs in $D$ because the border is way too far from the interior of $B$ for greedy to be tempted to use the modified metric (recall that greedy always takes the shortest path). We will define a new graph $G^{\prime}$, which is the graph induced by vertices in $B$. We also say that all the vertices exactly on the border of $B$ are all at a distance 0 from each other (see bottom right corner of Figure 6.2). With this modification, it becomes clear that the offline optimum cost on instance $\mathcal{I}^{\prime}$ is at most the cost paid by $\operatorname{OPT}(\mathcal{I})$ inside $B$, which we will denote by $w(\operatorname{OPT}(\mathcal{I}) \cap B)$. Hence the offline optimum cost for each sub-instance is at most what the global optimum pays locally inside the ball that created the sub-instance. Since all balls in $\mathcal{B}$ are disjoint, these areas never overlap; hence the sum of all local optima is at most the global optimum of instance $\mathcal{I}$. Using this observation, we handle the cost incurred by pairs in $D$ via a delicate induction hypothesis on the number of cost classes in the instance. This induction is described formally in Subsection 6.2.5.

### 6.2.4 Building a balanced dual solution

We formalize here Step 1 of the previous subsection. We give a formal definition of all the properties that our triple $(\mathcal{B}$, charge, $D)$ should satisfy. Note that we are also given an upper bound $K$ on the real number of pairs $k$ (this is for technical reasons for handling the induction in the next subsection). Recall that for a ball $B \in \mathcal{B}^{(j)}$ we denote by $B_{\mathcal{P}}$ the set of pairs of $\bigcup_{j^{\prime}>j} \mathcal{P}^{\left(j^{\prime}\right)}$ such that one of its endpoint is at distance at most

$$
r \cdot\left(1+\frac{1}{200 \cdot \log ^{2}(K)}\right)
$$

from the center of $B$. We also have similar definitions for $\partial B_{\mathcal{P}}$ and $\stackrel{\circ}{B}_{\mathcal{P}}$. Now we can state the main definition of this subsection. Intuitively, conditions (a)
and (b) state that $\mathcal{B}$ is a feasible dual solution whose dual balls have radii large enough. Condition (c) states that the total charged cost of dangerous pairs inside the ball $B$ is never much more than polylog $(K)$ times the charged cost of the pair that created the ball $B$. Similarly, condition (d) states that the charged cost of dangerous pairs on the border of $B$ is not more than $1 / \log (K)$ times the charged cost of dangerous pairs strictly inside $B$. The last condition (e) states that no pair was charged too many times.

Definition 6.8 (Balanced dual solution). A balanced dual solution for an $(\alpha, \delta)$-canonical instance $\mathcal{I}$ with respect to algorithm A is a quadruple ( $\mathcal{B}$, charge, $D, K$ ) such that:
(a) All balls $B \in \mathcal{B}$ are pairwise disjoint and $D \subseteq \bigcup_{B \in \mathcal{B}} B_{\mathcal{P}}$. Moreover, $\mathcal{B}$ is partitioned into $M$ sub-collections of balls $\mathcal{B}^{(1)}, \mathcal{B}^{(2)}, \ldots, \mathcal{B}^{(M)}$ such that,
(b) For every $j \geqslant 1$, each ball in $\mathcal{B}^{(j)}$ has a radius $r^{\prime}$ that satisfies $r_{j} / 2 \leqslant$ $r^{\prime} \leqslant r_{j}=c_{j} / 8 \log \left(k_{j}\right)$, with $c_{j}$ the cost associated to cost class $\mathcal{P}^{(j)}$ and $k_{j}=\left|\mathcal{P}^{(j)}\right|$.
(c) For every ball $B \in \mathcal{B}^{(j)}$

$$
\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, \stackrel{\circ}{\mathcal{P}} \cap D, \text { charge }) \leqslant 10 \cdot \operatorname{charge}(p(B)) \cdot c_{j} \cdot \log ^{10}(K),
$$

(d) For any ball $B \in \mathcal{B}^{(j)}$

$$
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \partial B_{\mathcal{P}} \cap D, \text { charge }\right) \leqslant \frac{10 \cdot \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \circ_{\mathcal{P}} \cap D, \text { charge }\right)}{\log (K)},
$$

(e) For any $j>0$, any pair $p \in \mathcal{P}^{(j)}$,

$$
\operatorname{charge}(p) \leqslant\left\{\begin{array}{l}
55 \cdot e^{5} \text { ifp is a surviving pair, } \\
0 \text { ifp is a charged pair, }
\end{array}\right.
$$

charge $(p) \geqslant 1$ ifp is a surviving or dangerous pair, and finally if $p$ is a dangerous pair that belongs to $B_{\mathcal{P}}$ for some $B \in \mathcal{B}^{(j)}$ then

$$
\operatorname{charge}(p) \leqslant\left(1+\frac{5}{\log (K)}\right)^{j-1} .
$$

The main result of this subsection will be that it is always possible to find a balanced dual solution.

Lemma 6.9. Given an $(\alpha, \delta)$-canonical instance $\mathcal{I}$ with respect to A , a balanced dual solution $(\mathcal{B}$, charge, $D, K)$ always exists provided that $\delta \geqslant 100 \cdot(\log (\alpha)+$ $\log \log (K))$, the number of cost classes $M$ satisfies $M \leqslant \log (K)$, and the number of pairs $k$ satisfies $k \leqslant K$.

Proof. We build the solution quadruple ( $\mathcal{B}$, charge, $D, K$ ) with an iterative procedure from $j=1$ to $j=M$ (recall that $M$ is the number of cost classes in $\mathcal{I}$ ) that will maintain the following invariants at the beginning of any iteration $j$ :
(i) For any $j^{\prime}<j$, the dual balls in $\mathcal{B}^{\left(j^{\prime}\right)}$ are already fixed and all pairs in $\mathcal{P}^{\left(j^{\prime}\right)}$ are already classified as either surviving, charged, or dangerous. The dual balls of $\mathcal{B}^{\left(j^{\prime}\right)}$ satisfy conditions (b), (c), (d) of Definition 6.8.
(ii) For any $j^{\prime}<j$, the charge of pairs in $\mathcal{P}^{\left(j^{\prime}\right)}$ satisfy condition (e) of Definition 6.8.
(iii) For any $j^{\prime} \geqslant j$, the pairs in $\mathcal{P}^{\left(j^{\prime}\right)}$ can be classified as either dangerous or charged, or not be classified yet. No pair of $\mathcal{P}^{\left(j^{\prime}\right)}$ is classified as surviving yet. We also have $\mathcal{B}^{\left(j^{\prime}\right)}=\emptyset$.
(iv) For any $j^{\prime} \geqslant j$, all pairs in $\mathcal{P}^{\left(j^{\prime}\right)}$ classified as either dangerous or charged satisfy condition (e) in Definition 6.8. The pairs $p$ that are not yet classified satisfy

$$
1 \leqslant \operatorname{charge}(p) \leqslant\left(1+\frac{5}{\log (K)}\right)^{j-1}
$$

(v) For any $j^{\prime} \geqslant j$, all pairs in $\mathcal{P}^{\left(j^{\prime}\right)}$ not yet classified do not belong to any set $B_{\mathcal{P}}$ for some already existing ball (i.e. unclassified pairs are far from the balls already placed).

We start iteration $j$ by considering the pairs in $\mathcal{P}^{(j)}$ that are not yet classified. To avoid confusion, let us denote by $\mathcal{P}^{\prime(j)} \subseteq \mathcal{P}^{(j)}$ this set of unclassified pairs. Using Lemma 6.5, we get a dual solution $\mathcal{B}^{(j)}$ of balls all of radius

$$
r_{j}=\frac{c_{j}}{8 \log \left(k_{j}\right)}
$$

that are all centered around endpoints of pairs in $\mathcal{P}^{\prime(j)}$ and such that no pair in $\mathcal{P}^{\prime(j)}$ has more than one ball centered around an endpoint. All the pairs in $\mathcal{P}^{\prime(j)}$ that do not have a dual ball can already be classified as charged. We decrease the charge of these pairs to o increase the charge of the other pairs in $\mathcal{P}^{\prime(j)}$ to compensate. Since we have that $\left|\mathcal{P}^{\prime(j)}\right| \leqslant 5 \cdot\left|\mathcal{B}^{(j)}\right|$ and invariant (iv), we have that the charge of the remaining pairs in $\mathcal{P}^{\prime(j)}$ will be at most

$$
\begin{equation*}
5 \cdot\left(1+\frac{5}{\log (K)}\right)^{(j-1)} \leqslant 5 \cdot\left(1+\frac{5}{\log (K)}\right)^{M} \leqslant 5 \cdot e^{5} \tag{6.1}
\end{equation*}
$$

after this step (we use the assumption that $M \leqslant \log (K)$ and $\left.(1+x) \leqslant e^{x}\right)$. Then for each new ball $B \in \mathcal{B}^{(j)}$ that we created around the endpoint of a pair $p$ we consider the following process with two cases. We consider the pairs of $\stackrel{\circ}{B}_{\mathcal{P}}$ (i.e. pairs of smaller cost in the interior of $B$ ) and proceed to a case distinction on the value of their total charged cost

$$
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \stackrel{\circ}{B}_{\mathcal{P}}, \text { charge }\right)
$$

that we denote $\Sigma$.

CASE 1 (TOP LEFT CORNER in figure 6.2). If

$$
\begin{equation*}
\Sigma>10 \cdot \operatorname{charge}(p) \cdot c_{j} \cdot \log ^{10}(K) \tag{6.2}
\end{equation*}
$$

that we can rephrase intuitively as " $B$ does not satisfy condition (c)" then we simply erase the dual ball $B$ and charge all the cost charge $(p) \cdot c_{j}$ to the pairs in $\stackrel{\circ}{B}_{\mathcal{P}}$ proportionally to their weight. Note that by doing this, all the pairs in $\stackrel{\circ}{B}_{\mathcal{P}}$ see their charge increasing by a multiplicative factor of $\left(1+\frac{1}{\left(10 \log ^{10}(K)\right)}\right)<\left(1+\frac{5}{\log (K)}\right)$. The pair $p$ is now classified as charged. Note that the pairs in $\mathcal{P}^{\left(j^{\prime}>j\right)}$ have now a charge that is at most $\left(1+\frac{5}{\log (K)}\right)^{j}$ by applying invariant (iv). This will be the only place where their charge can increase during iteration $j$ and note that this happens only once per iteration since only the pairs inside a ball $B \in \mathcal{B}^{(j)}$ are charged and the balls in $\mathcal{B}^{(j)}$ are pairwise disjoint. Therefore invariant (iv) will hold at the beginning of iteration $j+1$ for these pairs.

CASE 2. In this case we assume that $\Sigma \leqslant 10 \cdot \operatorname{charge}(p(B)) \cdot c_{j} \cdot \log ^{10}(K)$. We consider the ball $B^{\prime}$ with the same center as $B$ but with only half the radius of $B$ (i.e. $B^{\prime}$ is a ball centered at $p(B)$ with radius $r=r_{j} / 2$ ). We look at the total charged cost of pairs in $B_{\mathcal{P}}^{\prime}$ and proceed to a case distinction on this value (note that we do not consider only the strict interior of $B$ but also its border). Denote by $\Sigma^{\prime}$ this new value, i.e.

$$
\Sigma^{\prime}=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{\prime}, \text { charge }\right)
$$

We proceed by a sub-case distinction of the value of $\Sigma^{\prime}$.
SUB-CASE $2(\mathrm{~A})$. If $\Sigma^{\prime} \leqslant 10 \cdot \operatorname{charge}(p(B)) \cdot c_{j}$ then we simply charge the cost of all the pairs in in $B_{\mathcal{P}}^{\prime}$ to the pair $p$. In this case the pairs in $B_{\mathcal{P}}^{\prime}$ become charged pairs. Note that the charge of $p$ in that case is multiplied by at most 11 . We also set the dual ball $B$ to $B^{\prime}$ (i.e. we scale down the radius of $B$ by a factor 2). By Equation 6.1, we get that the charge of $p$ will be at most

$$
\operatorname{charge}(p) \leqslant 11 \cdot 5 \cdot e^{5}=55 e^{5}
$$

This will be the final charge of this pair hence invariant (ii) will be satisfied.
SUB-CASE 2(в) (вотTOM RIGHT CORNER IN FIGURE 6.2). If $\Sigma^{\prime}>$ $10 \cdot \operatorname{charge}(p(B)) \cdot c_{j}$, then we will try to increase the radius $r$ of $B^{\prime}$ by an increment of

$$
\Delta r=\frac{r_{j}}{200 \cdot \log ^{2}(K)}
$$

until we find a radius which satisfies condition (d) for the ball $B^{\prime}$. We denote by $B^{(t)}$ the ball obtained after increasing the radius of $B^{\prime} t$ times by $\Delta r$. Hence the radius of $B^{(t)}$ is $r_{t}=r_{j} / 2+t \cdot(\Delta r)$. Each time we increase $r$, we check if

$$
\begin{equation*}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \partial B_{\mathcal{P}}^{(t)}, \text { charge }\right) \leqslant \frac{10 \cdot \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \stackrel{\circ}{B}_{\mathcal{P}}^{(t)}, \text { charge }\right)}{\log (K)} \tag{6.3}
\end{equation*}
$$

If this is the case we stop and add all the pairs in $B_{\mathcal{P}}^{(t)}$ to $D$. The pairs in $B_{\mathcal{P}}^{(t)}$ become in this case dangerous pairs. If not we continue increasing the radius $r$ until Equation 6.3 holds. We claim that this process will stop before reaching $r_{t}=r_{j}$. To see this we first prove that $B_{\mathcal{P}}^{(t)} \cap \partial B_{\mathcal{P}}^{(t+6)}=\emptyset$ for any $t \geqslant 0$. Indeed if a pair $p_{i} \in \mathcal{P}^{\left(j^{\prime}\right)}$ with $j^{\prime}>j$ belongs to $B_{\mathcal{P}}^{(t)} \cap \partial B_{\mathcal{P}}^{(t+6)}$ then it must be that one of the two endpoints of $p_{i}=\left\{s_{i}, t_{i}\right\}$ belongs to $B^{(t)}$ while the other belongs to the border $\partial B^{(t+6)}$ since $B^{(t)} \cap \partial B^{(t+6)}=\emptyset$. Hence the distance between the two endpoints of $p_{i}=\left\{s_{i}, t_{i}\right\}$ satisfies:

$$
\begin{aligned}
d_{G}\left(s_{i}, t_{i}\right) \geqslant\left(r_{t}+6 \Delta r\right) \cdot(1 & \left.-\frac{1}{200 \cdot \log ^{2}(K)}\right)-r_{t} \cdot\left(1+\frac{1}{200 \cdot \log ^{2}(K)}\right) \\
\geqslant & 5 \Delta r-\frac{2 r_{t}}{200 \cdot \log ^{2}(K)} \geqslant \frac{r_{j}}{200 \cdot \log ^{2}(K)}
\end{aligned}
$$

where the last inequality comes from the fact that $r_{t} \leqslant r_{j}$. However, we assumed that the instance $\mathcal{I}$ is $(\alpha, \delta)$-canonical with $\delta \geqslant 100 \cdot(\log (\alpha)+$ $\log \log (K))$ and this implies that all pairs have contraction at most $\alpha$ which means that the pair $p_{i}$ should have costed A at least

$$
\begin{array}{r}
\frac{r_{j}}{\alpha\left(200 \cdot \log ^{2}(K)\right)} \geqslant \frac{c_{j}}{\alpha\left(8 \cdot 200 \cdot \log ^{3}(K)\right)}>\frac{c_{j}}{2^{10} \cdot \alpha^{100} \log ^{100}(K)} \\
>\frac{c_{j}}{2^{\delta+10}}=c_{j+1}
\end{array}
$$

which is a contradiction. Thus $B_{\mathcal{P}}^{(t)} \cap \partial B_{\mathcal{P}}^{(t+6)}=\emptyset$. Since we assume the process does not stop, we have that

$$
\begin{aligned}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \partial B_{\mathcal{P}}^{(t+6)}, \text { charge }\right)> & \frac{10 \cdot \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \stackrel{\circ}{B}_{\mathcal{P}}^{(t+6)}, \text { charge }\right)}{\log (K)} \\
\geqslant & \frac{10 \cdot \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t)}, \text { charge }\right)}{\log (K)}
\end{aligned}
$$

Together with the fact that $B_{\mathcal{P}}^{(t)} \cap \partial B_{\mathcal{P}}^{(t+6)}=\emptyset$ we get

$$
\begin{aligned}
& \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t+6)}, \operatorname{charge}\right) \geqslant \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t)},{\operatorname{charge})+\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \partial B_{\mathcal{P}}^{(t+6)}, \text { charge }\right)}\right. \\
& \geqslant \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t)}, \text { charge }\right) \cdot\left(1+\frac{10}{\log (K)}\right)
\end{aligned}
$$

This implies that, for all $t \geqslant 0$,

$$
\begin{aligned}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t)}, \text { charge }\right) \geqslant \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(0)}, \text { charge }\right) \cdot & \left(1+\frac{10}{\log (K)}\right)^{\lfloor t / 6\rfloor} \\
& =\Sigma^{\prime} \cdot\left(1+\frac{10}{\log (K)}\right)^{\lfloor t / 6\rfloor}
\end{aligned}
$$

However, we have by assumption in our case that

$$
\frac{\Sigma}{\Sigma^{\prime}} \leqslant \log ^{10}(K)
$$

thus for $t=60 \log (K) \log \log (K)$, we can write simultaneously

$$
\begin{align*}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t)}, \text { charge }\right) & \geqslant \Sigma^{\prime} \cdot\left(1+\frac{10}{\log (K)}\right)^{10 \log (K) \log \log (K)}  \tag{6.4}\\
& >\Sigma^{\prime} \cdot \log ^{10}(K) \geqslant \Sigma \tag{6.5}
\end{align*}
$$

and

$$
\begin{equation*}
r_{t}=r_{j} / 2+t \cdot(\Delta r)=r_{j} \cdot\left(\frac{1}{2}+\frac{60 \log (K) \log \log (K)}{200 \cdot \log ^{2}(K)}\right)<r_{j} \tag{6.6}
\end{equation*}
$$

This is a contradiction since by Equation 6.6 we should have $B_{\mathcal{P}}^{(t)} \subseteq B_{\mathcal{P}}$ and by Equation 6.5 we have $\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}^{(t)}\right.$, charge $)>\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, B_{\mathcal{P}}\right.$, charge $)$. Thus the process must stop before reaching $t=60 \log (K) \log \log (K)$, hence before reaching $r_{t}=r_{j}$.

CORRECTNESS. We note that Sub-case $2(\mathrm{~b})$ is the only case in which we create dangerous pairs. By construction, properties (c) and (d) will hold for any ball. This is because if a ball $B$ remains around a pair $p$ at the end of the procedure, either all the pairs inside are charged to $p$, or they become dangerous pairs. In the case where the pairs inside become dangerous the procedure described in Sub-case 2(b) stops exactly when properties (c) and (d) are satisfied. Property (b) is also satisfied since the only place where a radius can be modified is in Sub-case 2(b). In this case, we show that after halving the radius, it cannot grow for too long before the number of pairs in the border is less concentrated than in the interior. This means that since the balls in $\mathcal{B}^{(j)}$ were disjoint when taken with radius $r_{j}$, they also have to be disjoint if the radius is slightly smaller. Condition (a) is also satisfied. To see this, note that the balls in $\mathcal{B}^{(j)}$ cannot intersect balls installed before since in Case 2, which is the only case where a ball survives, for any ball $B \in \mathcal{B}^{\left(j^{\prime}<j\right)}$, all the pairs in $B_{\mathcal{P}}$ become classified as either dangerous or charged (invariant (v)). In particular, we will never try to place a ball around these pairs since we only place dual balls around pairs that are not classified yet. Since $B_{\mathcal{P}}$ encompasses a region slightly bigger than the ball $B$, the smaller dual balls will not intersect $B$. We also have that $D \subseteq \bigcup_{B \in \mathcal{B}} B_{\mathcal{P}}$ since dangerous pairs are only created in Sub-case $2(\mathrm{~b})$ where we consider pairs inside a ball $B$. It remains to show that condition (e) holds. Note that for charged pairs, this is clear. For surviving pairs, we showed that our process maintains invariant (ii); hence after the end of the procedure, condition (e) is satisfied. For dangerous pairs, note that once a pair becomes dangerous, its charge will never increase anymore. Hence this shows that invariant (ii) also holds for these pairs, and in particular, at the end of the process, condition (e) holds.

### 6.2.5 Inductive proof using balanced dual solutions

Given the previous subsection, we are ready to state the main induction. Recall that in a balanced dual solution ( $\mathcal{B}$, charge, $D$ ), all the pairs except the dangerous pairs (the set $D$ ) are accounted for by the dual balls in $\mathcal{B}$. Our induction will precisely take advantage of this. In the following, for any ball
$B$ and instance $\mathcal{I}$, we denote by $w(\operatorname{OPT}(\mathcal{I}) \cap B)$ the total cost of edges that are contained in $B$ and bought by $\operatorname{OPT}(\mathcal{I})$ (note that we assume that edges in $G$ are arbitrarily small, so no edge is crossed by the border of $B$ ). Note that we also have the straightforward bound $r \leqslant w(\operatorname{OPT}(\mathcal{I}) \cap B)$ (with $r$ being the radius of $B$ ) since the optimum solution needs to connect at least the center of the ball $B$ to its border. For a collection of disjoint balls $\mathcal{B}$, we define

$$
w(\mathrm{OPT}(\mathcal{I}) \cap \mathcal{B})=\sum_{B \in \mathcal{B}} w(\mathrm{OPT}(\mathcal{I}) \cap B) .
$$

Lemma 6.10. Let $\mathcal{I}$ be an $(\alpha, \delta)$-canonical instance with respect to algorithm A. Assume that its size is $k \leqslant K$, and $\delta \geqslant 100 \cdot(\log (\alpha)+\log \log (K))$. Assume there are $M \leqslant \log (K)$ distinct cost classes when running A on instance $\mathcal{I}$. Let $\left(\mathcal{B}=\bigcup_{j=1}^{M} \mathcal{B}^{(j)}\right.$, charge, $\left.D, K\right)$ be a balanced dual solution. Then we have that

$$
\begin{aligned}
& \operatorname{cost}_{\mathrm{A}}(\mathcal{I}) \leqslant\left(880 e^{5}\right) \cdot\left(\sum_{j=1}^{M} \log \left(k_{j}\right) \cdot w\left(O P T(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right) \\
& \quad+e^{200+20 M / \log (K)} \cdot\left(\sum_{j=1}^{M} \delta \cdot(M-j) \cdot w\left(O P T(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right),
\end{aligned}
$$

where $k_{j}$ is the number of terminals in cost class $j$ for $j=1, \ldots, M$.
Intuitively, the first term of the right-hand side of the inequality corresponds to the pairs that are either surviving or charged and can be charged to the dual. The second term corresponds to the cost paid by A because of dangerous pairs in $D$. The proof of Lemma 6.10 will be done by induction on the number of cost classes $M$.

Proof. The base case of the induction is for $M=0$ in which case the statement is vacuously true (the instance is empty). Hence assume $M>0$ and that the statement is true for any $(\alpha, \delta)$-canonical instance with $M^{\prime}<M$ cost classes, and satisfying the conditions of the lemma.

Recall that in a balanced dual solution ( $\mathcal{B}=\bigcup_{j=1}^{M} \mathcal{B}^{(j)}$, charge, $D, K$ ) (that exists by Lemma 6.9) we have that charge $(p) \leqslant 55 e^{5}$ for all surviving pairs $p$, that $\mathcal{B}$ is a feasible dual solution, and that the radius of each ball $B \in \mathcal{B}^{(j)}$ is at least $1 /\left(16 \log \left(k_{j}\right)\right)$ times the cost of corresponding pair. Hence we have that the total cost of surviving or charged pairs (denote this set $\mathcal{P}^{\prime}$ ) is at most

$$
\begin{aligned}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \mathcal{P}^{\prime}\right) \leqslant & \left(16 \cdot 55 e^{5}\right) \cdot\left(\sum_{j=1}^{M} \log \left(k_{j}\right) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right) \\
& =\left(880 e^{5}\right) \cdot\left(\sum_{j=1}^{M} \log \left(k_{j}\right) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right) .
\end{aligned}
$$

All that remains to do is to upper bound the total cost incurred by A because of pairs in $D$. By property (a) of Definition 6.8, it suffices to consider each
ball $B \in \mathcal{B}$ and the dangerous pairs it contains. Hence fix a ball $B \in \mathcal{B}^{(j)}$ of radius $r$. We consider the instance $\mathcal{I}^{\prime}=\left(G^{\prime}, \mathcal{P}^{\prime}, \mathcal{S}^{\prime}\right)$ that is defined as follows (see bottom right corner of Figure 6.2).

- The graph $G^{\prime}$ is the graph induced by $G$ on the vertex set $B$, in which we add an edge of cost 0 between any pair of vertices at a distance exactly $r$ from the center of $B$. We denote by $E^{\prime}$ this set of additional edges of weight o (see Figure 6.2 in the bottom right corner).
- $\mathcal{P}^{\prime}$ is the set of pairs in $\stackrel{\circ}{B}_{\mathcal{P}} \cap D$, given in the same relative order to algorithm A. Note that these pairs must have both endpoints inside $B$ because all pairs have a low contraction, and the cost paid by these pairs is much smaller than the radius of the ball. Hence traveling from the interior of $B$ to the border is already way too far. We also emphasize that we ignore pairs in $\partial B_{\mathcal{P}} \cap D$.
- The set of additional edges $\mathcal{S}^{\prime}$ contains all the edges in $\mathcal{S}$ that we revealed just before reading a pair in $\mathcal{P}^{\prime}$. Moreover, the edges in $\mathcal{S}^{\prime}$ are revealed in the same way to $A$, that is, if an edge $e \in \mathcal{S}$ is revealed just before a pair $p \in \mathcal{P}^{\prime}$, then it is also revealed just before $p$ in instance $\mathcal{I}^{\prime}$.

We then make the following claim.
Claim 6.11. Given instance $\mathcal{I}^{\prime}$ we have that:
(a) For all $p \in \mathcal{P}^{\prime}, \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}, p\right)=\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)$. In particular, the cost incurred by A because of pairs in $\mathcal{P}^{\prime}$ is exactly the same as the cost that A would pay when running on instance $\mathcal{I}^{\prime}$.
(b) $w\left(O P T\left(\mathcal{I}^{\prime}\right)\right) \leqslant w(B \cap O P T(\mathcal{I}))$.

Proof. To see (b), let us buy the edge set $(\operatorname{OPT}(\mathcal{I}) \cap B) \cup E^{\prime}$. It is clear that this costs at most $w(B \cap \mathrm{OPT}(\mathcal{I}))$ since all edges in $E^{\prime}$ have cost o. We claim that this is a feasible solution for the instance $\mathcal{I}^{\prime}$. Consider any pair $p=\{s, t\} \in \mathcal{P}^{\prime}$. It must be that $\operatorname{OPT}(\mathcal{I})$ contains a path between $s$ and $t$. If this path does not leave the ball $B$ then it is contained in the edge set $(\operatorname{OPT}(\mathcal{I}) \cap B) \cup E^{\prime}$. Otherwise, denote this path by its sequence of vertices: $s, v_{1}, v_{2}, \ldots, v_{k}, t$. Denote by $v_{f}$ and $v_{\ell}$ the first and last vertices on this path that are outside $B$. Then it is clear that the edge set $(\mathrm{OPT}(\mathcal{I}) \cap B) \cup E^{\prime}$ contains the path $s, v_{1}, \ldots, v_{f-1}, v_{\ell+1}, \ldots, v_{k}, t$. In both case, $s$ and $t$ are connected and $(\operatorname{OPT}(\mathcal{I}) \cap B) \cup E^{\prime}$ is a feasible solution to instance $\mathcal{I}^{\prime}$.

To see (a), we prove by induction on the number of pairs $\mathcal{P}^{\prime}$ arrived so far that the claim holds. If no pair arrived yet, this is vacuously true. Otherwise, consider the arrival of pair $p$ and assume (a) holds for all pairs that arrived before. First note that the set of shortcuts added by A when connecting any pair $p^{\prime}=\{s, t\}$ in instance $\mathcal{I}$ is just one edge $e=\{s, t\}$ with weight o. This is because we assumed that just before $p^{\prime}$ arrived, an edge of $\mathcal{S}$ between $s$ and $t$ arrived with weight exactly $\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, p^{\prime}\right)$. Hence we can assume that A went through this edge to connect $p^{\prime}$. Here we also use the fact that when A goes through only one edge to connect a pair, the three contraction rules behave exactly the same.

Now by induction hypothesis, the set of shortcuts bought by A so far on instance $\mathcal{I}^{\prime}$ have exactly the same property. Hence this set of shortcuts is a subset of the shortcuts added by A on instance $\mathcal{I}$. We claim that when pair $p=\{s, t\}$ arrives, the shortest path is again through the corresponding edge of $\mathcal{S}^{\prime}$. Assume this is not the case. As a shorthand, define $c=\operatorname{cost}_{\mathrm{A}}(\mathcal{I}, p)$. Denote by $B$ the set of vertices in $B$ that are at distance at most

$$
r \cdot\left(1-\frac{1}{100 \cdot \log ^{2}(K)}\right)
$$

from the center of $B$. We first claim that all the shortcuts added so far by A in instance $\mathcal{I}^{\prime}$ have both endpoints in $B^{\circ}$ (i.e. far from the border of $B$ ). To see this, note that all the pairs in $\mathcal{P}^{\prime}$ are in $B_{\mathcal{P}}$ hence they have one endpoint at distance at most

$$
r \cdot\left(1-\frac{1}{200 \cdot \log ^{2}(K)}\right)
$$

from the center of $B$. If the other endpoint was not in $B$, the distance $d$ between the two endpoints would be at least

$$
d \geqslant \frac{r}{200 \cdot \log ^{2}(K)},
$$

and since pairs have contraction at most $\alpha$, we would have that the cost paid for this pair in $\mathcal{I}$ is at least

$$
\frac{d}{\alpha} \geqslant \frac{r}{200 \cdot\left(\alpha \log ^{2}(K)\right)}
$$

but we assume that cost classes would be separated by at least a multiplicative $2^{\delta+10}>2^{10} \cdot \alpha^{100} \cdot \log ^{100}(K)$ by assumption on $\delta$. Hence we would have a contradiction. Hence all pairs in $\mathcal{P}^{\prime}$ have both endpoints in $B$, but this also implies that all shortcuts added so far by A in instance $\mathcal{I}^{\prime}$ have both endpoints in $\stackrel{\circ}{B}$.

Now remark that because the instance is $(\alpha, \delta)$-canonical, we have that $d_{G}(s, t) \leqslant \alpha c$, and since the distance between a point in $\stackrel{B}{B}$ and the exact border of $B$ is much bigger than $\alpha c$, it cannot be that the shortest path uses some edges in $E^{\prime}$ (just to reach them is already too expensive because all previously bought shortcuts have endpoints in $\stackrel{B}{B}$ ). Hence the shortest path only uses edges that were available to A when connecting $p$ in instance $\mathcal{I}$. Therefore the shortest path available in $\mathcal{I}^{\prime}$ can only be longer than the shortest path for the same pair in $\mathcal{I}$. Since the corresponding edge of $\mathcal{S}^{\prime}$ of weight exactly $c$ is available, the shortest path is in fact exactly the same.

By Claim 6.11, we know that the cost incurred by A because of pairs in $B_{\mathcal{P}}$ is equal to the cost that A would pay on the instance $\mathcal{I}^{\prime}$. Denote by $k_{i}^{\prime}$ the number of pairs in $\mathcal{P}^{\prime(i)}$ in $\mathcal{I}^{\prime}$. Note that by Claim 6.11 we have $\mathcal{P}^{\prime(i)}=\mathcal{P}^{(i)} \cap \stackrel{\circ}{\mathcal{P}}_{\mathcal{P}} \cap D$. Because of properties (c) and (e) of Definition 6.8, it must be that the following inequalities hold

$$
\begin{align*}
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \stackrel{\circ}{B}_{\mathcal{P}}, \text { charge }\right) & \leqslant 10 \cdot \operatorname{charge}(p(B)) \cdot c_{j} \cdot \log ^{10}(K)  \tag{6.7}\\
& \leqslant\left(550 e^{5}\right) \cdot c_{j} \cdot \log ^{10}(K) \tag{6.8}
\end{align*}
$$

where $c_{j}$ is the cost of the pair that created the ball $B$, and

$$
\begin{equation*}
k_{i}^{\prime} \cdot c_{i} \leqslant \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, \stackrel{\circ}{B}_{\mathcal{P}}, \text { charge }\right) \tag{6.9}
\end{equation*}
$$

Putting Equations 6.8 and 6.9 together we obtain

$$
\begin{equation*}
k_{i}^{\prime} \leqslant\left(550 e^{5}\right) \cdot \log ^{10}(K) \cdot \frac{c_{j}}{c_{i}}=\left(550 e^{5}\right) \cdot \log ^{10}(K) \cdot 2^{(\delta+10) \cdot(i-j)} \leqslant 2^{3(\delta+10) \cdot(i-j)} \tag{6.10}
\end{equation*}
$$

for all $i>j$.
Note that in instance $\mathcal{I}^{\prime}$, we have at most $M^{\prime}=M-j<M$ cost classes. Finally, it is clear that $\mathcal{I}^{\prime}$ is still an $(\alpha, \delta)$-canonical instance, A behaves the same on the relevant pairs, and all conditions of Lemma 6.10 are satisfied (we keep the same upper bound $K$ on the number of pairs). Hence we can apply the induction hypothesis on the sub-instance $\mathcal{I}^{\prime}$ to obtain the following upper bound on the cost of dangerous pairs inside $B$.

$$
\begin{aligned}
& \operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}\right) \leqslant\left(880 e^{5}\right)\left(\sum_{i=1}^{M^{\prime}} \log \left(k_{i}^{\prime}\right) \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right) \\
& +e^{200+20 M^{\prime} / \log (K)} \cdot\left(\sum_{i=1}^{M^{\prime}} \delta \cdot\left(M^{\prime}-i\right) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right)
\end{aligned}
$$

with $k_{i}^{\prime} \leqslant 2^{3(\delta+10) \cdot i}$ by re-indexing cost classes from 1 to $M^{\prime}$ in Equation 6.10. The first term on the right-hand side is

$$
\begin{aligned}
& \left(880 e^{5}\right) \cdot\left(\sum_{i=1}^{M-j} \log \left(k_{i}^{\prime}\right) \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right) \\
& \leqslant\left(880 e^{5}\right) \cdot\left(\sum_{i=1}^{M-j} \log \left(2^{3(\delta+10) \cdot i}\right) \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right) \\
& \leqslant\left(2640 e^{5}\right) \cdot(\delta+10) \cdot \sum_{i=1}^{M-j} i \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right) \\
& \leqslant\left(\delta e^{200+20 M^{\prime} / \log (K)}\right) \cdot \sum_{i=1}^{M-j} i \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)
\end{aligned}
$$

The second term in the right-hand side is less than

$$
\left(e^{200+20 M^{\prime} / \log (K)}\right) \cdot\left(\sum_{i=1}^{M-j} \delta \cdot(M-j-i) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right)
$$

By summing both terms, we obtain a cost of at most

$$
\begin{aligned}
& \left(e^{200+20 M^{\prime} / \log (K)}\right) \cdot \sum_{i=1}^{M-j}(i \delta) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)+ \\
& \left(e^{200+20 M^{\prime} / \log (K)}\right) \cdot\left(\sum_{i=1}^{M-j} \delta \cdot(M-j-i) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right) \\
& \leqslant\left(\delta e^{200+20 M^{\prime} / \log (K)}\right) \cdot\left(\sum_{i=1}^{M-j}(M-j) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right) \cap \mathcal{B}^{\prime(i)}\right)\right)
\end{aligned}
$$

Now recall that in a balanced dual solution, all balls are pairwise disjoint hence we obtain the upper bound

$$
\begin{equation*}
\left(\delta e^{200+20 M^{\prime} / \log (K)}\right) \cdot(M-j) \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right) \tag{6.11}
\end{equation*}
$$

on the cost of pairs in $\stackrel{\circ}{B}_{\mathcal{P}} \cap D$ when summing on all $i$ in the previous inequality.

By Claim 6.11, we also have $w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right)\right) \leqslant w(B \cap \operatorname{OPT}(\mathcal{I}))$. Now we need to take into account the charges that were put on dangerous pairs in $\stackrel{\circ}{B}_{\mathcal{P}} \cap D$ by the balanced dual solution in the big instance $\mathcal{I}$. By property (d) of Definition 6.8, the total charged cost incurred by A for pairs in $\partial B_{\mathcal{P}} \cap D$ is a most $10 / \log (K)$ fraction of the cost of pairs in $\stackrel{\circ}{B}_{\mathcal{P}} \cap D$. Additionally, the charge of each pair in $p \in \stackrel{\circ}{B}_{\mathcal{P}} \cap D$ is at most (by property (e))

$$
\left(1+\frac{5}{\log (K)}\right)^{j-1} \leqslant e^{5(j-1) / \log (K)}
$$

if the ball that contains this pair is in $\mathcal{B}^{(j)}$. Hence we only need to increase the upper bound given by Equation 6.11 by a multiplicative term $e^{(5(j-1)+10) / \log (K)} \leqslant e^{20 j / \log (K)}$ to get a final upper bound of

$$
\begin{align*}
& \left(\delta e^{200+(20(M-j)+20 j) / \log (K)}\right) \cdot(M-j) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right)\right) \\
& \quad \leqslant\left(\delta e^{200+20 M / \log (K)}\right) \cdot(M-j) \cdot w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right)\right) \tag{6.12}
\end{align*}
$$

on the total charged cost incurred by A (in instance $\mathcal{I}$ ) because of pairs in $B_{\mathcal{P}} \cap D$. To finish the proof, we sum these upper bounds over all $B \in \mathcal{B}^{(j)}$ and all $j$ to obtain indeed the second term of the induction hypothesis.

Given Lemma 6.10, it is now straightforward to prove Theorem 6.4. Lemma 6.10 applied to the main canonical instance $\mathcal{I}$ states that A pays at most

$$
\begin{aligned}
& \left(880 e^{5}\right) \cdot\left(\sum_{j=1}^{M} \log \left(k_{j}\right) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right) \\
& +e^{200+20 M / \log (K)} \cdot\left(\sum_{j=1}^{M} \delta \cdot(M-j) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right) \\
& \leqslant O\left[\sum_{j=1}^{M} \log \left(k_{j}\right) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right. \\
& \left.+\sum_{j=1}^{M} \delta \cdot(M-j) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right)\right]
\end{aligned}
$$

For the first inequality, we use that $M \leqslant \log (K)$. We obtain that the first term is at most

$$
\log (k) \cdot \sum_{j=1}^{M} w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right) \leqslant \log (k) \cdot w(\mathrm{OPT}(\mathcal{I}))
$$

since all balls in $\mathcal{B}$ are pairwise disjoint. The second term, is at most

$$
\begin{aligned}
\sum_{j=1}^{M} \delta \cdot(M-j) \cdot w\left(\operatorname{OPT}(\mathcal{I}) \cap \mathcal{B}^{(j)}\right) \leqslant & (\delta M) \cdot w(\operatorname{OPT}(\mathcal{I})) \\
& \leqslant \log (k) \cdot w(\operatorname{OPT}(\mathcal{I}))
\end{aligned}
$$

since we have that $M \leqslant \log (k) / \delta$ (choosing the upper bound $K=k$ is a valid choice). Hence, if $\mathcal{I}$ is an $(\alpha, \delta)$-canonical instance, we indeed have that the cost paid by A on this instance is at most $O(\log (k)) \cdot w(\mathrm{OPT}(\mathcal{I}))$, which proves Theorem 6.4, and Theorem 5.1, combining it with Lemma 6.3.

### 6.3 PROOF OF THEOREM 5.2 AND THEOREM $5 \cdot 3$

The aim of this section is to establish Theorem 5.2 and Theorem 5.3. To this end we prove Lemma 6.12. We recall that in this section, all results apply only to Greedy $3_{3}$ so A will be a shorthand for Greedy 3 in this whole section, unlike in Section 6.2 where A would mean that we could use any of the contraction rules.

Lemma 6.12. Suppose we are given an instance $\mathcal{I}$ of Steiner Forest, with $k$ pairs of terminals. Then we can construct an instance $\mathcal{I}^{\prime}$ of Steiner Forest that satisfies the following, where $\mathrm{A}=$ Greedy $_{3}$ :
(a) $\operatorname{cost}_{\mathrm{A}}(\mathcal{I})=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}\right)$.
(b) The contraction of all pairs in instance $\mathcal{I}^{\prime}$ when running A is exactly equal to 1 .
(c) $\mathcal{I}^{\prime}$ has $k^{\prime}=O\left(k^{2}\right)$ terminal pairs.
(d) The set of terminals (vertices that appear in a pair) is the same for instance $\mathcal{I}$ and $\mathcal{I}^{\prime}$.

Proof. The idea is to subdivide each pair $p_{i}$ in instance $\mathcal{I}$ into $O(k)$ pairs that are made of pairs of previously arrived terminals that lie on the path $P$ that is used to connect $p_{i}$. It will be clear from construction that Property (d) holds. Formally, we will construct an instance $\mathcal{I}^{\prime}$ of Steiner Forest as follows. Beginning with the empty set of terminal pairs, for each pair $\left\{s_{i}, t_{i}\right\}$ that arrives in $\mathcal{I}$ we will construct a list of pairs $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ and add them to our instance $\mathcal{I}^{\prime}$. The order in which we add pairs to $\mathcal{I}^{\prime}$ determines the arrival order. We next describe how to construct $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$. Suppose that when $\left\{s_{i}, t_{i}\right\}$ appears in $\mathcal{I}$ the algorithm A connects the pair using the path $P=s_{i}, v_{1}, \ldots, v_{\ell}, t_{i}$. Let $s_{i}, v_{1}^{\prime}, \ldots, v_{\ell^{\prime}}^{\prime}, t_{i}$ be the sub-sequence of $P$ specified by the contraction rule 3. Recall that this sub-sequence contains only $s_{i}, t_{i}$ and previously arrived terminals. We add to $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ the subset of $\left\{\left\{s_{i}, v_{1}^{\prime}\right\},\left\{v_{1}^{\prime}, v_{2}^{\prime}\right\}, \ldots,\left\{v_{\ell^{\prime}-1}^{\prime}, v_{\ell^{\prime}}^{\prime}\right\},\left\{v_{\ell^{\prime}}^{\prime}, t_{i}\right\}\right\}$ containing all the pairs such that their distance in the contracted metric is not yet 0 (here we consider the contracted metric just before connecting the pair $p=\left\{s_{i}, t_{i}\right\}$ in instance $\mathcal{I})$. For any pair $\left\{s^{\prime}, t^{\prime}\right\}$ in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ we say that $\left\{s_{i}, t_{i}\right\}$ is the parent of $\left\{s^{\prime}, t^{\prime}\right\}$.

Now that we have constructed $\mathcal{I}^{\prime}$, we argue that it satisfies the properties of the lemma. To show that $\operatorname{cost}_{\mathrm{A}}(\mathcal{I})=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}\right)$ we will prove the following stronger statement. For any pair $p_{i}=\left\{s_{i}, t_{i}\right\}$,

1. A spends exactly the same cost for the pairs in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ in instance $\mathcal{I}^{\prime}$ that it was spending for the pair $p_{i}=\left\{s_{i}, t_{i}\right\}$ in instance $\mathcal{I}$, i.e.

$$
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, p_{i}\right)=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}, \mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)\right)
$$

2. The contracted metric $G^{(i)}$ after revealing all the pairs up to pair $p_{i}$ in instance $\mathcal{I}$ is exactly the same as the contracted metric $G^{\prime(\tau)}$ after revealing all the pairs $\bigcup_{i^{\prime} \leqslant i} \mathcal{P}\left(\left\{s_{i^{\prime}}, t_{i^{\prime}}\right\}\right)$ in instance $\mathcal{I}^{\prime}$.

We will prove this statement by induction on $i$. It is vacuously true for $i=0$ since no pair appeared so far in both instances. Assume this is true up to pair $i-1$, and let us prove the statement for step $i$. Let $v_{0}^{\prime}, v_{1}^{\prime}, \ldots, v_{\ell^{\prime}}^{\prime}, v_{\ell^{\prime}+1}^{\prime}$ be the sub-sequence of previously arrived terminals in the path $P$ that $A$ uses to connect $s_{i}$ and $t_{i}$ in instance $\mathcal{I}$, where $v_{0}^{\prime}=s_{i}$ and $v_{\ell^{\prime}+1}^{\prime}=t_{i}$. Suppose that for some $0 \leqslant j \leqslant \ell^{\prime}$ the pair $\left\{v_{j}^{\prime}, v_{j+1}^{\prime}\right\}$ is not in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$. Then by construction we know that $v_{j}^{\prime}$ and $v_{j+1}^{\prime}$ were already at distance 0 in the contracted metric before the arrival of $\left\{s_{i}, t_{i}\right\}$ during the execution of A on $\mathcal{I}$, i.e.

$$
d_{G^{(i-1)}}\left(v_{j}^{\prime}, v_{j+1}^{\prime}\right)=0
$$

As a result the portion of the path $P$ between $v_{j}^{\prime}$ and $v_{j+1}^{\prime}$ contributes 0 to the value $\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, p_{i}\right)$ hence no cost is lost. Now suppose that the pair $\left\{v_{j}^{\prime}, v_{j+1}^{\prime}\right\}$ is in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$, and that when presented with the pair $\left\{v_{j}^{\prime}, v_{j+1}^{\prime}\right\}$ in instance $\mathcal{I}^{\prime}$, A uses a path $P^{\prime}$ to connect the pair. Denote by $P_{\left[v_{j}^{\prime}, v_{j+1}^{\prime}\right]}$ the
restriction of the path $P$ to vertices that appear in-between vertices $v_{j}^{\prime}, v_{j+1}^{\prime}$ (included). We claim that $P^{\prime}=P_{\left[v_{j}^{\prime}, v_{j+1}^{\prime}\right]}$. To see this, let us consider the first time this is not the case, this would mean that the path $P^{\prime}$ costs less to A in instance $\mathcal{I}^{\prime}$ than what $P_{\left[v_{j}^{\prime}, v_{j+1}^{\prime}\right]}$ costed to A in instance $\mathcal{I}$. But since we assumed that the contracted metrics were identical up to pair $p_{i-1}$, A could have replaced the path $P_{\left[v_{j}^{\prime}, v_{j+1}^{\prime}\right]}$ by $P^{\prime}$ to pay less. This is a contradiction to the fact that A always takes the shortest path. Hence the total cost for the pair $p_{i}$ is preserved and we have that

$$
\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}, p_{i}\right)=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}, \mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)\right)
$$

To see why we have the second property, now note that we have that $P^{\prime}=$ $P_{\left[v_{j}^{\prime}, v_{j+1}^{\prime}\right]}$, in particular on the path $P^{\prime}$ there is no previously arrived terminal other than $v_{j}^{\prime}, v_{j+1}^{\prime}$. By contraction rule 3, it means that A simply adds an edge $\left\{v_{j}^{\prime}, v_{j+1}^{\prime}\right\}$ of weight 0 in the contracted metric. But by the definition of contraction rule 3 , this edge was also added by A when connecting the pair $p_{i}$ in instance $\mathcal{I}$. Reciprocally, it is clear that if a shortcut was added by A when connecting the pair $p_{i}$ in instance $\mathcal{I}$, it must be between two previously arrived terminals $v_{j^{\prime}}^{\prime}, v_{j^{\prime}+1}^{\prime}$ that are consecutive on the path $P$. If these two previously arrived terminals were already at distance 0 in the contracted metric, then the shortcut does not change the metric. Otherwise if $v_{j^{\prime}}^{\prime}$ and $v_{j^{\prime}+1}^{\prime}$ were not at distance 0 then we have that $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ also contains the pair $\left\{v_{j^{\prime}}^{\prime}, v_{j^{\prime}+1}^{\prime}\right\}$ hence this shortcut will also be added in the metric when running instance $\mathcal{I}^{\prime}$. This proves that the contracted metrics are indeed the same after step $i$. We also have Property (b) since, by construction, there is no previously arrived terminal on the path $P^{\prime}$. Hence it must be that A pays the cost of the path $P^{\prime}$ in the original metric (the only way to pay less than the length in the original metric is to use a shortcut that was added before, but such a shortcut must connect previously arrived terminals).

Lastly, we argue that the number of terminals $k^{\prime}$ in $\mathcal{I}^{\prime}$ is $O\left(k^{2}\right)$. Since there are $k$ pairs in instance $\mathcal{I}$, we know that there are at most $2 k$ terminals in total. Thus there are at most $2 k$ terminals along any path A takes to connect a terminal pair in $\mathcal{I}$ (going twice through the same terminal cannot happen since A takes the shortest path). Summing this bound over all terminal pairs in $\mathcal{I}$, we get that $\mathcal{I}^{\prime}$ has at most $2 k^{2}$ terminals.

With Lemma 6.12, we are able to prove Theorem 5.2.
Theorem 5.2. Construct $\mathcal{I}^{\prime}$ from $\mathcal{I}$ as in Lemma 6.12. Then we can write

$$
\begin{aligned}
\operatorname{cost}_{\mathrm{A}}(\mathcal{I})=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}\right) & \leqslant O\left(\log \left(k^{\prime}\right) \cdot \log \log \left(k^{\prime}\right)\right) \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right) \\
& \leqslant O\left(\log \left(k^{\prime}\right) \cdot \log \log \left(k^{\prime}\right)\right) \cdot w\left(T^{\star}(\mathcal{I})\right) \\
& \leqslant O(\log (k) \cdot \log \log (k)) \cdot w\left(T^{\star}(\mathcal{I})\right)
\end{aligned}
$$

where the first equality follows by Property (a) of Lemma 6.12, the second inequality from Theorem 5.1 applied to instance $\mathcal{I}^{\prime}$ (using that contraction of all pairs is 1 by Property (b) of Lemma 6.12), the third inequality from the fact that the optimum tree solution to instance $\mathcal{I}$ denoted $T^{\star}(\mathcal{I})$ is also a feasible solution to instance $\mathcal{I}^{\prime}$ (because by Property (d) of Lemma 6.12 the set of terminals does not change) and the last inequality from the fact that $k^{\prime}=O\left(k^{2}\right)$ by Property (c) of Lemma 6.12.

Next we prove Theorem 5.3. For this, we require some additional notation. Let $\mathcal{I}$ be an instance of Steiner Forest and $\mathcal{F}$ a feasible solution consisting of trees $T_{1}, \ldots, T_{m}$. Then the width of a tree $T_{j}$ with respect to an instance $\mathcal{I}$, denoted width $\left(T_{j}, \mathcal{I}\right)$, is defined to be the largest distance between any terminal pair in $T_{j}$ in the original metric, i.e.

$$
\operatorname{width}\left(T_{j}, \mathcal{I}\right)=\max _{u \in \mathcal{T} \cap T_{j}} d_{G}(u, \bar{u})
$$

where $\bar{u}$ is the terminal that $u$ should be connected to in instance $\mathcal{I}$, and $\mathcal{T}$ is the set of all terminals in instance $\mathcal{I}$. This is exactly the same definition of width as in [50]. For such a forest $\mathcal{F}=\left(T_{1}, \ldots, T_{q}\right)$ we define the potential

$$
\Phi(\mathcal{F}, \mathcal{I})=w(\mathcal{F})+\sum_{j=1}^{q} \operatorname{width}\left(T_{j}, \mathcal{I}\right)
$$

Where $w(\mathcal{F})$ is the cost of the forest. Note that $w(\mathcal{F}) \leqslant \Phi(\mathcal{F}, \mathcal{I}) \leqslant 2 w(\mathcal{F})$ for any forest $\mathcal{F}$ that is a feasible solution to instance $\mathcal{I}$. We will use this property to prove Theorem 5.3.

Theorem 5.3. We construct the instance $\mathcal{I}^{\prime}$ from instance $\mathcal{I}$ exactly as how we transformed the instance in Lemma 6.12. We show that if the costs of shortest paths are non-increasing over time, it must be that

$$
w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right)=O(w(\mathrm{OPT}(\mathcal{I}))
$$

Once this is established, the result follows as we have

$$
\begin{aligned}
\operatorname{cost}_{\mathrm{A}}(\mathcal{I})=\operatorname{cost}_{\mathrm{A}}\left(\mathcal{I}^{\prime}\right) & =O\left(\log \left(k^{\prime}\right) \cdot \log \log \left(k^{\prime}\right)\right) \cdot w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right) \\
& =O(\log (k) \cdot \log \log (k)) \cdot w(\mathrm{OPT}(\mathcal{I}))
\end{aligned}
$$

again by Lemma 6.12 and Theorem 5.1 (recall that subdividing pairs to create instance $\mathcal{I}^{\prime}$ as in Lemma 6.12 guarantees that the contraction is 1 for all pairs in the instance $\mathcal{I}^{\prime}$ ).

We argue that $w\left(\operatorname{OPT}\left(\mathcal{I}^{\prime}\right)\right)=O(w(\operatorname{OPT}(\mathcal{I}))$, using essentially the same potential function argument than that of [50]. The idea is to begin with the solution $\operatorname{OPT}(\mathcal{I})$ and add additional connections to produce a feasible solution $\mathcal{F}$ to $\mathcal{I}^{\prime}$, where $w(\mathcal{F}) \leqslant 2 w(\mathrm{OPT}(\mathcal{I}))$. Since $w\left(\mathrm{OPT}\left(\mathcal{I}^{\prime}\right)\right) \leqslant w(\mathcal{F})$ if we succeed the proof is complete.

As previously stated, we initialize $\mathcal{F}$ to be $\operatorname{OPT}(\mathcal{I})$. We will only add edges to $\mathcal{F}$ hence it will be clear that $\mathcal{F}$ will always be a feasible solution to instance $\mathcal{I}$. Now for each terminal pair $\left\{s_{i}, t_{i}\right\}$ that arrives in $\mathcal{I}$ we construct $\mathcal{F}^{\prime}$ from $\mathcal{F}$ as follows. Suppose that all the pairs in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ are already connected by the current solution $\mathcal{F}$, then nothing needs to be done.

Suppose otherwise that the solution $\mathcal{F}$ does not connect all pairs in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$. By re-indexing let $T_{1}, \ldots, T_{q}$ be the components of $\mathcal{F}$ that contain the terminals that appear in the pairs of $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$, ordered so that

$$
\operatorname{width}\left(T_{1}\right) \geqslant \ldots \geqslant \operatorname{width}\left(T_{q}\right)
$$

Construct $\mathcal{F}^{\prime}$ from $\mathcal{F}$ by adding the edges that A uses to connect the pairs in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$ in instance $\mathcal{I}^{\prime}$. We claim that the total cost of these edges (which is exactly equal to $\operatorname{cost}\left(\mathcal{I}, p_{i}\right)$ ) is at most the width of $T_{q}$ (with respect to instance $\mathcal{I}$ ), i.e.

$$
\operatorname{cost}\left(\mathcal{I}, p_{i}\right) \leqslant \operatorname{width}\left(T_{q}, \mathcal{I}\right)
$$

To see this first note that the tree $T_{q}$ either contains the pair $\left\{s_{i}, t_{i}\right\}$ or a pair that appeared even before $\left\{s_{i}, t_{i}\right\}$ in instance $\mathcal{I}$ (here we use that $\mathcal{F}$ is a feasible solution to instance $\mathcal{I}$ ). Hence if the sequence of shortest paths $d_{1}=d_{G}\left(s_{1}, t_{1}\right), \ldots, d_{k}=d_{G}\left(s_{k}, t_{k}\right)$ is non-increasing, then

$$
\operatorname{width}\left(T_{q}, \mathcal{I}\right) \geqslant \min _{i^{\prime} \leqslant i} d_{i^{\prime}}=d_{i} \geqslant c_{i}=\operatorname{cost}\left(\mathcal{I}, p_{i}\right)
$$

For the same reason if the sequence of costs $c_{1}=\operatorname{cost}\left(\mathcal{I}, p_{1}\right), \ldots, c_{k}=$ $\operatorname{cost}\left(\mathcal{I}, p_{k}\right)$ is non-increasing, then

$$
\operatorname{width}\left(T_{q}, \mathcal{I}\right) \geqslant \min _{i^{\prime} \leqslant i} d_{i^{\prime}} \geqslant \min _{i^{\prime} \leqslant i} c_{i^{\prime}}=c_{i}=\operatorname{cost}\left(\mathcal{I}, p_{i}\right) .
$$

We use here that $d_{i} \geqslant c_{i}$ for all $i$. This proves our claim.
We now note that, if we needed to add edges because the current solution $\mathcal{F}$ did not connect all the pairs in $\mathcal{P}\left(\left\{s_{i}, t_{i}\right\}\right)$, then we obtained a new solution $\mathcal{F}^{\prime}$ such that

$$
w\left(\mathcal{F}^{\prime}\right) \leqslant w(\mathcal{F})+\operatorname{cost}\left(\mathcal{I}, p_{i}\right)
$$

and at least two of the components in $T_{1}, T_{2}, \ldots, T_{q}$ were merged into one component. Hence

$$
\begin{aligned}
\operatorname{width}\left(\mathcal{F}^{\prime}, \mathcal{I}\right) \leqslant \operatorname{width}(\mathcal{F}, \mathcal{I})-\min _{1 \leqslant i \leqslant q} & \operatorname{width}\left(T_{i}, \mathcal{I}\right) \\
& \leqslant \operatorname{width}(\mathcal{F}, \mathcal{I})-\operatorname{width}\left(T_{q}, \mathcal{I}\right)
\end{aligned}
$$

In particular we obtain

$$
\Phi\left(\mathcal{F}^{\prime}, \mathcal{I}\right)-\Phi(\mathcal{F}, \mathcal{I}) \leqslant \operatorname{cost}\left(\mathcal{I}, p_{i}\right)-\operatorname{width}\left(T_{q}, \mathcal{I}\right) \leqslant 0
$$

by the above remarks. We then update $\mathcal{F}$ to be $\mathcal{F}^{\prime}$.
After completing the process we have that $\mathcal{F}$ connects all pairs in $\mathcal{I}^{\prime}$, as required and that the potential function did not increase. Therefore we have that

$$
w\left(\mathcal{F}^{\prime}\right) \leqslant \Phi\left(\mathcal{F}^{\prime}, \mathcal{I}\right) \leqslant \Phi(\mathcal{F}, \mathcal{I}) \leqslant 2 w(\mathcal{F})
$$

as desired.

This chapter contains the proof of Theorem $5 \cdot 4$, which is arguably much simpler than in other works. We start by a brief explanation of the overall idea in Section 7.1 before providing the complete proof in Section 7.2. The results presented in this chapter are based on the paper entitled "A Simple LPbased Approximation Algorithm for the Matching Augmentation Problem", which is a joint work with Marina Drygala and Ola Svensson. It appeared at the International Conference on Integer Programming and Combinatorial Optimization (IPCO '22) [10].

### 7.1 OUR PROOF TECHNIQUE

The proof of Theorem 5.4 relies on several crucial observations that we sketch here. The first observation is that the total cost of the DFS tree $T$ is always at most the cost of $x^{*}$ (denoted $c\left(x^{*}\right)$ ). This follows because $T$ must contain all the light edges since they are given priority over any other edge (note that since we assume that $M$ is a matching, it cannot happen that two distinct light edges want priority at the same time). Therefore, the total cost of the tree $T$ is exactly equal to $n-1-|M|$, while it is easy to show that $c\left(x^{*}\right) \geqslant(n-|M|)$.

Another interesting fact is that if one considers the LP solution $x^{*}$ restricted to the edges not in the tree $T$ (denote this solution by $x_{E \backslash T}^{*}$ ), then this is a feasible solution to the cut LP of the TAP instance with respect to the tree $T$ (i.e. $x_{E \backslash T}^{*}$ is a feasible solution to $L P(G, T)$ ). Hence, if we denote by $y^{*}$ the optimum fractional solution to $L P(G, T)$, we have that $c\left(y^{*}\right) \leqslant c\left(x_{E \backslash T}^{*}\right)$.

Because $T$ is a DFS tree, the TAP instance with respect to the tree $T$ contains only "uplinks" and therefore $L P(G, T)$ is known to be integral [1]. We note that this already gives a simple proof that the integrality gap of $L P(G, M)$ is at most 2 . To get better than 2 , we only need to show that

$$
(n-|M|-1)+c\left(y^{*}\right) \leqslant(2-\delta) c\left(x^{*}\right)
$$

for some $\delta>0$. Conceptually, we distinguish between two cases. If $c\left(x^{*}\right)>$ $(1+\delta)(n-|M|)$ (i.e. the LP solution is expensive), then the DFS tree is significantly cheaper than $c\left(x^{*}\right)$ and it is easy to conclude that the cost of our solution $T \cup A$ is better than $2 c\left(x^{*}\right)$. Otherwise, assume that the LP value is close to the trivial lower bound of $(n-|M|)$. In this case, we show that $c\left(y^{*}\right) \leqslant(1-\delta) c\left(x^{*}\right)$.

To show this, we consider two possibilities. We can prove that either we can scale down a significant portion of $x_{E \backslash T}^{*}$ to obtain a cheaper feasible solution to $L P(G, T)$, or that $c\left(x_{E \backslash T}^{*}\right)$ itself is already significantly smaller than $c\left(x^{*}\right)$. When a lot of the tree cuts in $T$ (i.e. the cuts defined by removing an edge from $T$ to obtain two trees and taking the edges with one endpoint in each tree) have some slack in the TAP solution $x_{E \backslash T}^{*}$ (that is when a lot of tree
cuts $S$ satisfy $\left.x_{E \backslash T}^{*}(\delta(S))>1+\delta\right)$, the first case is realized. Otherwise, when almost all of the tree cuts are nearly tight (i.e. satisfy $x_{E \backslash T}^{*}(\delta(S)) \leqslant 1+\delta$ ), we can show that the DFS must have captured a good fraction of the value of $c\left(x^{*}\right)$ inside the tree $T$. This step uses some crucial properties of extreme point solutions as well as our choice of DFS. Therefore the cost of $x_{E \backslash T}^{*}$ is significantly smaller than the cost of $x^{*}$ completing the argument.

Before proceeding to the proof, it is worthwhile to mention that we are not aware of any example on which our algorithm has a ratio worse than $4 / 3$ times the cost of $x^{*}$. It remains open to give a tighter analysis of this algorithm. We also note that [62] also makes use of DFS for the related problem of unweighted 2-ECSS. They obtain a ratio of $3 / 2$ for the unweighted 2-ECSS problem. However, their DFS is not LP-based and we remark that if we do not guide the DFS with the LP solution, the approximation ratio can be arbitrarily close to 2 . We give an example in Figure 7.1. One can see that the DFS tree (rooted at $r$ ) contains all the matching edges, and the tree augmentation problem requires us to take all but one of the back-edges. However, the optimum solution to the MAP instance is to take a Hamiltonian tour containing all the light edges. Generalizing the same example by simply increasing the depth of the tree leads to an approximation arbitrarily close to 2 .


Figure 7.1: An example of a bad DFS tree.

### 7.2 THE ANALYSIS OF THE LP-BASED ALGORITHM

In this section, we prove Theorem 5.4. It is organized as follows. In subsection 7.2.1, we introduce some basic definitions. In the subsequent subsection, we proceed via a case distinction to prove the theorem.

### 7.2.1 Preliminaries

We will use $T$ to refer to the DFS tree computed by the algorithm, and we will list edges in $G$ as $u v$, where $u$ is ancestor of $v$ in $T$. Since $T$ is a DFS tree, all edges in $G$ must have the property that one endpoint is an ancestor of the other in $T$. We will let $B=E \backslash T$ denote the set of back-edges of $G$.

As in the introduction, we will call an edge of weight 1 a heavy edge and an edge of weight o a light edge. For every edge $e$ in the DFS tree $T$ computed, we let $T(e)$ denote the tree cut corresponding to the edge $e$ in the tree $T$. Formally, $T(e)=\delta\left(T_{v}\right)$, where $e=u v$ and $T_{v}$ is the sub-tree rooted at $v$. We call an edge $e \in T \alpha$-tight if we have

$$
x^{*}(T(e))-x_{e}^{*}<1+\alpha
$$

Implicitly, if we call an edge $e \alpha$-tight, this will mean that $e$ belongs to the tree $T$. In addition, we denote by $N_{t}^{(\alpha)}$ the number of $\alpha$-tight edges in the tree $T$. For a tree $T$, we denote by $x_{T}^{*}$ the restriction of $x^{*}$ to the edges in the tree $T$. We note that for any instance of the MAP, it must be that $c\left(x^{*}\right) \geqslant(n-|M|)$. This follows by a simple double counting argument on the fractional degree of each component (precisely we have $n-|M|$ components that must have fractional degree 2 each). It is also clear that the DFS tree $T$ must contain all the light edges in $M$ since they are given priority. Hence the cost of $T$ is at most $n-|M|-1 \leqslant c\left(x^{*}\right)$. In the following, we will fix two parameters $\varepsilon=10^{-1}, \gamma=10^{-3}$.

### 7.2.2 The analysis of the algorithm

We note that if $c\left(x^{*}\right) \geqslant(1+\gamma)(n-|M|)$, it is easy to show that the cost of the returned solution $T \cup A$ is at most

$$
\begin{equation*}
(n-|M|-1)+c\left(x^{*}\right) \leqslant \frac{c\left(x^{*}\right)}{1+\gamma}+c\left(x^{*}\right)=c\left(x^{*}\right)\left(2-\frac{\gamma}{1+\gamma}\right) \tag{7.1}
\end{equation*}
$$

However, if $c\left(x^{*}\right)<(1+\gamma)(n-|M|)$ and $N_{t}^{(\gamma)} \leqslant(1-\gamma)(n-|M|)$ (i.e. there are few $\gamma$-tight tree cuts), then we proceed as follows. We partition the set of back edges in our graph $B$ into $B_{t}^{(\gamma)} \cup B_{s}^{(\gamma)}$, where $B_{t}^{(\gamma)}$ contains all edges $e \in B$ that are contained in $T\left(e^{\prime}\right)$ for some $\gamma$-tight edge $e^{\prime} \in T$. Then $x^{\prime}$, defined by

$$
x^{\prime}(e)= \begin{cases}x^{*}(e) & e \in B_{t}^{(\gamma)} \\ \frac{x^{*}(e)}{1+\gamma} & e \in B_{s}^{(\gamma)} \\ 1 & \text { otherwise }\end{cases}
$$

is also a feasible solution to $L P(G, T)$. The total fractional value represented by edges in $B_{t}^{(\gamma)}$ is at most $(1+\gamma) N_{t}^{(\gamma)}$. Hence, $c\left(x^{\prime}\right)$ can be upper bounded as follows.

$$
c\left(x^{\prime}\right) \leqslant \frac{c\left(x^{*}\right)-(1+\gamma) N_{t}^{(\gamma)}}{1+\gamma}+(1+\gamma) N_{t}^{(\gamma)}=\frac{c\left(x^{*}\right)}{1+\gamma}+\gamma N_{t}^{(\gamma)}
$$

Since the cost of $T \cup A$ is at most $c\left(x^{*}\right)+c\left(x^{\prime}\right)$ and we assume that

$$
N_{t}^{(\gamma)} \leqslant(1-\gamma)(n-|M|)
$$

it is easy to get the upper-bound of

$$
c\left(x^{*}\right)\left(1+\frac{1}{1+\gamma}\right)+\gamma(1-\gamma)(n-|M|) \leqslant c\left(x^{*}\right)\left(2-\frac{\gamma^{3}}{1+\gamma}\right)
$$

where the last inequality follows because $n-|M| \leqslant c\left(x^{*}\right)$. Since these two cases clearly give a better-than-2 approximation, we assume in the rest of the analysis that

$$
\begin{equation*}
(n-|M|) \leqslant c\left(x^{*}\right)<(1+\gamma)(n-|M|), \tag{7.3}
\end{equation*}
$$

and

$$
\begin{equation*}
N_{t}^{(\gamma)}>(1-\gamma)(n-|M|) \tag{7.4}
\end{equation*}
$$

We will show that $c\left(x_{T}^{*}\right)$ is at least a constant fraction times $c\left(x^{*}\right)$. Since the cost of the returned solution $T \cup A$ is at most $2 c\left(x^{*}\right)-c\left(x_{T}^{*}\right)$, this will conclude the proof. First, we partition the $\gamma$-tight tree cuts into two sets of cuts $\mathcal{S}_{0}$ and $\mathcal{S}_{1}$ containing the tight tree cuts associated with light edges and heavy edges, respectively. We can then distinguish between two sub-cases. For each edge $e=u v \in T$, we say that $e$ is a leaf edge if $v$ is a leaf in the tree $T$ (recall that we always write an edge $e$ as $e=u v$ such that $v$ is a descendant of $u$ in $T$ ). We denote $\mathcal{S}_{0}^{+}$the non-leaf edges in $\mathcal{S}_{0}$ and $\mathcal{S}_{0}^{-}$the leaf edges in $\mathcal{S}_{0}$. We have two main cases.
7.2.2.1 $\quad$ Suppose that $\left|\mathcal{S}_{1}\right| \geqslant \gamma(n-|M|)$ or that $\left|\mathcal{S}_{0}^{+}\right| \geqslant \gamma(n-|M|)$.

By feasibility of $x^{*}$ at least 2 units of $x^{*}$ must cross any tree cut. Hence $x^{*}\left(\delta\left(T_{v}\right)\right) \geqslant 2$, for any $v \in V$. By definition of $\gamma$-tightness we know that for any $\gamma$-tight edge $e=u v$ we have $x^{*}(e) \geqslant x^{*}\left(\delta\left(T_{v}\right)\right)-(1+\gamma) \geqslant 1-\gamma$.

Hence if $\left|\mathcal{S}_{1}\right| \geqslant \gamma(n-|M|)$, we have that

$$
c\left(x_{T}^{*}\right) \geqslant \gamma(1-\gamma)(n-|M|) \geqslant \frac{\gamma(1-\gamma)}{1+\gamma} c\left(x^{*}\right)
$$

which concludes the case when $\left|\mathcal{S}_{1}\right|$ is large. In the following we use some properties of extreme point solutions. We say that an edge $e$ is fractional (with respect to the fractional solution $x^{*}$ ) if $0<x_{e}^{*}<1$. A vertex $v$ is said to be $\alpha$-fractional if it has more than $1 / \alpha$ incident fractional edges in the support of $x^{*}$ (for any $\alpha>0$ ). We claim the following lemma, the proof of which relies on standard techniques that we repeat here for completeness. We note that a similar result was used in [69].

Lemma 7.1. If $x^{*}$ is an extreme point solution of the cut $L P$, then there are at most $2 n-1$ fractional edges in $G$. Moreover, for any $\alpha>0$, there are at most $4 \alpha n \alpha$-fractional vertices with respect to $x^{*}$.

Proof. An extreme point $x^{*}$ can be defined as the unique solution to the following system of $|E|$ linearly independent constraints, for some $\mathcal{S} \subseteq 2^{V}$ and $E_{0} \cup E_{1} \subseteq E$.

$$
\begin{array}{ll}
\sum_{e \in \delta(S)} x_{e}=2, & \text { for all } S \in \mathcal{S} \\
x_{e}=0 & \forall e \in E_{0} \\
x_{e}=1 & \forall e \in E_{1}
\end{array}
$$

We then use Theorem 4.9 from [33] that proves that we can select $\mathcal{S}$ to be a laminar family.

Lemma 7.2 (Theorem 4.9 in [33]). Let $x^{*}$ be an extreme point of the MAP cut $L P$ then the family of equations $\mathcal{S}$ can be chosen to be a laminar family.

It is well known that any laminar family has size at most $2 n-1$. Using this fact together with Lemma 7.2, we obtain that the number of fractional edges is at most $|E|-\left|E_{0}\right|-\left|E_{1}\right|=|\mathcal{S}| \leqslant 2 n-1$. To finish the proof of Lemma 7.1, fix some $\alpha>0$. Each vertex that is $\alpha$-fractional accounts for at least $1 / \alpha$ fractional edges. By a simple double-counting argument, there can be at most $\alpha \cdot(4 n-2) \leqslant 4 \alpha n$ vertices that are $\alpha$-fractional.

Using Lemma 7.1 with $\alpha=\gamma / 16$, we get that if $\left|\mathcal{S}_{0}^{+}\right| \geqslant \gamma(n-|M|)$, then (recall that $n-|M| \geqslant n / 2$ ) there are at least

$$
\gamma(n-|M|)-(\gamma / 4) n \geqslant(\gamma / 2)(n-|M|)
$$

edges $u v \in \mathcal{S}_{0}^{+}$such that $v$ is not $\gamma / 16$-fractional. We then claim the following simple lemma.

Lemma 7.3. Fix any $\alpha, \alpha^{\prime}>0$. Suppose that $e=u v$ is an $\alpha$-tight light edge, such that $v$ is not a leaf in $T$. Then if $v$ is not $\alpha^{\prime}$-fractional there exists some edge $e^{\prime}=v w$ in $T$ such that $x^{*}\left(e^{\prime}\right) \geqslant(1-\alpha) \alpha^{\prime}$.

Proof. By feasibility of $x^{*}$ we know that $x^{*}\left(\delta\left(T_{v} \backslash v\right)\right) \geqslant 2$. Since $e$ is $\alpha$-tight and $T$ is a DFS tree we have that $x^{*}\left(\delta\left(T_{v}\right)\right)-x^{*}(e) \leqslant 1+\alpha$. We know that $E\left(T_{v} \backslash v, v\right)=\delta\left(T_{v} \backslash v\right) \backslash \delta\left(T_{v}\right)$, and as a result $x^{*}\left(E\left(T_{v} \backslash v, v\right)\right) \geqslant 1-\alpha$. Since $v$ is not $\alpha^{\prime}$-fractional there must be an edge $e^{\prime} \in E\left(T_{v} \backslash v, v\right)$ with value at least $x^{*}\left(E\left(T_{v} \backslash v, v\right)\right) \alpha^{\prime} \geqslant(1-\alpha) \alpha^{\prime}$. Since our DFS selects always the highest possible fractional value if there is no light edge to explore, the first edge selected after exploring $v$ must be of fractional value at least $(1-\alpha) \alpha^{\prime}$.

Combining Lemma 7.3 with the previous observation, if $\left|\mathcal{S}_{0}^{+}\right| \geqslant \gamma(n-|M|)$ we get that

$$
c\left(x_{T}^{*}\right) \geqslant(\gamma / 2)(n-|M|)(1-\gamma)(\gamma / 16) \geqslant c\left(x^{*}\right) \frac{\gamma^{2}(1-\gamma)}{32(1+\gamma)} .
$$

Combining these two cases we get that if $\left|\mathcal{S}_{1}\right| \geqslant \gamma(n-|M|)$ or $\left|\mathcal{S}_{0}^{+}\right| \geqslant$ $\gamma(n-|M|)$ then

$$
c\left(x_{T}^{*}\right) \geqslant c\left(x^{*}\right) \cdot \min \left(\frac{\gamma^{2}(1-\gamma)}{32(1+\gamma)}, \frac{\gamma(1-\gamma)}{1+\gamma}\right)
$$

hence the cost of $T \cup A$ is upper bounded by

$$
\begin{equation*}
2 c\left(x^{*}\right)-c\left(x_{T}^{*}\right) \leqslant c\left(x^{*}\right)\left(2-\frac{\gamma^{2}(1-\gamma)}{32(1+\gamma)}\right) \tag{7.5}
\end{equation*}
$$

which is clearly better than 2 . Hence we are left with the last case, in which

$$
\left|\mathcal{S}_{0}^{-}\right|>(1-\gamma)(n-|M|)-\left|\mathcal{S}_{1}\right|-\left|\mathcal{S}_{0}^{+}\right|>(1-3 \gamma)(n-|M|) .
$$

7.2.2.2 $\quad$ Suppose $\left|\mathcal{S}_{0}^{-}\right|>(1-3 \gamma)(n-|M|)$.

This is the most interesting case. Note that for each edge $e=u v \in \mathcal{S}_{0}^{-}$, the fractional degree of $v$ restricted to heavy edges must be at least 1 , and all of this fractional degree is carried by backedges in $T$. Denote by $B^{\prime}$ this subset of backedges. Next we define $B^{\prime \prime} \subseteq B^{\prime}$ to be the subset of $B^{\prime}$ containing only edges with fractional value at least $\varepsilon=10^{-1}$. We claim that

$$
\begin{equation*}
\left|B^{\prime \prime}\right| \geqslant n / 10 \tag{7.6}
\end{equation*}
$$

Assume the contrary, since the fractional value of any edge is at most 1 then the total value carried by edges in $B^{\prime} \backslash B^{\prime \prime}$ must be at least

$$
\left|\mathcal{S}_{0}^{-}\right|-(n / 10)>(1-3 \gamma)(n-|M|)-n / 10>3 n / 8-n / 10
$$

(Recall that $(n-|M|) \geqslant n / 2$ and $(1-3 \gamma)>3 / 4)$. Since all the edges in $B^{\prime} \backslash B^{\prime \prime}$ have fractional value at most $\varepsilon$, there must be at least ( $3 n / 8-$ $n / 10) / \varepsilon=30 n / 8-n>2 n-1$ such edges, contradicting Lemma 7.1. Hence $\left|B^{\prime \prime}\right| \geqslant n / 10$.

For completeness we consider the case when $E$ contains heavy edges that are parallel to light edges. Partition $B^{\prime \prime}$ into $B_{1}^{\prime \prime} \cup B_{2}^{\prime \prime}$, where $B_{2}^{\prime \prime}$ is the set of edges in $B^{\prime \prime}$ parallel to an edge in $\mathcal{S}_{0}^{-}$. We define $B_{1}^{\prime \prime}$ to be the remaining edges in $B^{\prime \prime}$.

We claim that $\left|B_{2}^{\prime \prime}\right| \leqslant n / 100$, and thus loosely $\left|B_{1}^{\prime \prime}\right| \geqslant n / 20$.
To see this note that,

$$
\begin{equation*}
c\left(x^{*}\right)-(n-|M|) \geqslant\left|B_{2}^{\prime \prime}\right| \varepsilon \tag{7.7}
\end{equation*}
$$

Equation (7.7) holds as the lower bound of $(n-|M|)$ on $c\left(x^{*}\right)$ is obtained only by counting the fractional degree of each component in $M$. Since those parallel edges are not counted in this bound (they are only within a single component), they directly count in the value of $c\left(x^{*}\right)-(n-|M|)$, which counts the surplus of $c\left(x^{*}\right)$ above $(n-|M|)$.

Then as $c\left(x^{*}\right)-(n-|M|) \leqslant \gamma(n-|M|)$, by choice of $\varepsilon$ and $\gamma$ we obtain that $\left|B_{2}^{\prime \prime}\right| \leqslant n / 100$.

Consider the set of vertices $X$ that contains the ancestor vertices of the edges in $B_{1}^{\prime \prime}$. We claim that

$$
\begin{equation*}
|X| \geqslant n / 500 \tag{7.8}
\end{equation*}
$$

To prove this, we first claim that

$$
\begin{equation*}
c\left(x^{*}\right)-(n-|M|) \geqslant\left|B_{1}^{\prime \prime}\right| \varepsilon-2|X| . \tag{7.9}
\end{equation*}
$$

To see this, note again that the value $c\left(x^{*}\right)-(n-|M|)$ represents the surplus value of $c\left(x^{*}\right)$ above the lower bound that gives fractional degree 2 to every vertex. This trivial lower bound gives a fractional value-which is the fractional degree restricted to heavy edges-of at most 2 to every vertex, hence a fractional value of at most $2|X|$ to the set of vertices $X$. Since every edge in $B_{1}^{\prime \prime}$ has fractional value of at least $\varepsilon$ and is adjacent to a single vertex in $X$, we get that the surplus value of $c\left(x^{*}\right)$ above the trivial lower bound is at least $\left|B_{1}^{\prime \prime}\right| \varepsilon-2|X|$ which proves Equation (7.9).


Figure 7.2: On the left side, the case when the first edge selected out of $u$ is heavy. On the right the case when the first edge selected out of $u$ is light.

Since by assumption we have $c\left(x^{*}\right)<(1+\gamma)(n-|M|)$ we conclude with Equation (7.9) that

$$
\gamma(n-|M|)>c\left(x^{*}\right)-(n-|M|) \geqslant\left|B_{1}^{\prime \prime}\right| \varepsilon-2|X|
$$

which implies, by our lower bound on $\left|B_{1}^{\prime \prime}\right|$ and our choice of $\gamma$ and $\varepsilon$,

$$
\begin{equation*}
|X| \geqslant \frac{\left|B_{1}^{\prime \prime}\right| \varepsilon-\gamma(n-|M|)}{2} \geqslant \frac{n / 200-n / 10^{3}}{2}=n / 500 \tag{7.10}
\end{equation*}
$$

For each vertex $u \in X$, denote by $e_{u}$ the first edge selected by the DFS after reaching $u$. Denote $X^{\prime} \subseteq X$ the subset of $X$ containing only vertices $u \in X$ such that $e_{u}$ does not belong to $\mathcal{S}_{0}^{+}$. Then, we have by assumption,

$$
\left|X^{\prime}\right| \geqslant|X|-\left|\mathcal{S}_{0}^{+}\right| \geqslant n / 500-\gamma(n-|M|) \geqslant n / 500-n / 10^{3}=n / 10^{3}
$$

We finally claim the following, which crucially uses how the DFS selects the edges to explore in priority.

## Claim 7.4.

$$
c\left(x_{T}^{*}\right) \geqslant \varepsilon\left|X^{\prime}\right|
$$

Proof. There are two cases to consider (depicted in Figure 7.2).
If $u \in X^{\prime}$ is such that $e_{u}=u v$ is a heavy edge, by definition of $X^{\prime}$ there must be an edge $e^{\prime}=u f$ coming from a leaf $f$ in the tree $T$ to $u$ of fractional value $x_{e^{\prime}} \geqslant \varepsilon$. At the first time the DFS visits the vertex $u$, the leaf $f$ was not explored yet hence the edge $e^{\prime}$ was a valid choice of edge to explore. Since our DFS always takes the highest fractional value, it must be that

$$
x_{e_{u}} \geqslant x_{e^{\prime}} \geqslant \varepsilon
$$

If $u \in X^{\prime}$ is such that $e_{u}=u v$ is a light edge, recall that by definition of $X^{\prime}, v$ must be a leaf in $T$. Then when the DFS arrived at $v$, it must be that all reachable vertices from $v$ were already visited. Hence the DFS must have backtracked to $u$. Now note that by our construction of $B_{1}^{\prime \prime}$, we know that there must be another leaf $f$ such that $e^{\prime}=u f$ is a back-edge in the tree of fractional value $x_{u f} \geqslant \varepsilon$ (recall that $u f$ is not parallel to the edge $e_{u}$ ). Since
$f$ is a leaf, $u$ must have been explored before $f$ therefore, after backtracking from $v$ to $u$ the edge $u f$ was a valid edge to take. Therefore the DFS must have selected a second edge $e^{\prime \prime}$ in the tree from $u$ such that

$$
x_{e^{\prime}} \geqslant x_{u f} \geqslant \varepsilon
$$

Hence we proved that all vertices $u$ in $X$ must be adjacent to at least one heavy edge of fractional value $\varepsilon$ that belongs to the tree $T$ and goes to a child of $u$. Hence the proof of the claim.

By the previous claim, we have $c\left(x_{T}^{*}\right) \geqslant \varepsilon\left|X^{\prime}\right|$ hence the cost of the returned solution $T \cup A$ is at most

$$
\begin{equation*}
2 c\left(x^{*}\right)-c\left(x_{T}^{*}\right) \leqslant 2 c\left(x^{*}\right)-n / 10^{4} \leqslant c\left(x^{*}\right)\left(2-10^{-4}\right) \tag{7.11}
\end{equation*}
$$

which ends the proof of Theorem 5.4.

In the first part of this thesis, we have given improved approximation algorithms for the Santa Claus problem in various settings. For the MaxMin Arborescences problem, we gave an $O$ (poly $(\log \log n))$-approximation running in quasi-polynomial time. As we highlighted in the introduction of this thesis, the ultimate goal would be to design a constant factor approximation for the Santa Claus problem with linear valuation functions (if it exists). Here we give three open problems that are necessary intermediate steps towards this goal.

Open Question 8.1. Can the $O(\operatorname{poly}(\log \log n))$-approximation guarantee be generalized to the Santa Claus problem with linear valuations?

A big hurdle to extend our proof to the general case seems to be our use of the Lovasz Local Lemma. Indeed, in the general version of the problem, one needs to deal with a slightly more general version of the MaxMin Arborescences problem, in which two vertices in a arborescence can be connected by a long path instead of a single edge. This seems to introduce more dependencies between bad events. Nevertheless, we believe that many of our ideas will be very useful to generalize to the Santa Claus with linear valuations.

Open Question 8.2. Is there an $O(1)$-approximation for MaxMin Arborescences (possibly in quasi-polynomial time)?

It is worthwhile to recall that the MaxMin Arborescences problem is a special case of the Santa Claus problem. Hence obtaining a constant factor for MaxMin Arborescences is necessary if one aims at a constant factor for the Santa Claus problem. We mention that even in the special case of a single source in a layered graph, we do not know how to solve Open Question 8.2. This seems to be an interesting starting point.

Open Question 8.3. Is an $O(\operatorname{poly}(\log n))$-approximation in polynomial time possible for MaxMin Arborescences?

The best guarantee that is known to be possible in polynomial time is only polynomial [15, 21]. This state of affairs in reminiscent of the Directed Steiner Tree problem, for which an $O\left(n^{\varepsilon}\right)$-approximation is achievable in polynomial time for every fixed $\varepsilon>0$ [22]. Designing an $O(\operatorname{poly}(\log n))$-approximation in polynomial time for the Directed Steiner Tree problem is a notorious open problem in approximation algorithms. Hence we believe Open Question 8.3 to be very challenging.

As our second result, we showed an $O(\log \log n)$-approximation algorithm (in polynomial time) for the Restricted Submodular Santa Claus. The main open question here is the following.

Open Question 8.4. Is there an $O(1)$-approximation for the Restricted Submodular Santa Claus?

In the second part of this thesis, we provided new results for fundamental network design problems. For the Steiner Forest problem, we show that offline greedy is an $O(\log (k) \cdot \log \log (k))$-approximation. Even if our result has important consequences also for the online version, it is still an interesting open question to generalize our result to online greedy. Furthermore, improving the bound to a simple $O(\log (k))$ would be very interesting. Both these questions can be summarized in a single open question.

Open Question 8.5. Is the greedy algorithm an $O(\log (k))$-approximation in the online setting?

As a starting point, one might assume that all pairs have contraction 1 and the cost of each pair belongs to the interval [OPT/ $\log (k)$, OPT]. As far as we know, this case seems to retain all the difficulty and we are not aware of any better bounds in this setting that the ones in this thesis.

Finally, for the Matching Augmentation Problem, we gave a simple polynomial time algorithm that guarantees a better-than-2 approximation when compared to the cut LP. Tha Matching Augmentation problem can be seen as an instance of the cheapest 2 -edge connected subgraph problem in a graph with edge-weights 0 or 1 , where the edges of cost 0 form a matching. A natural direction would be to drop the assumption on the structure of edges of cost 0 , in which case we obtain the Forest Augmentation problem. A natural research direction is the following question.

Open Question 8.6. Is the integrality gap of the cut LP for the Forest Augmentation problem strictly less than 2 ?

It was shown that it is better than 2 for the (unweighted) Tree Augmentation problem ([74]), which is the case where we assume the edges of cost 0 to form a spanning tree. Our result shows a similar statement for the Matching Augmentation problem. Since these two cases are the two "extreme" cases of the Forest Augmentation problem, it is believable that the cut LP is also better than 2 in the more general Forest Augmentation case.

Part III
APPENDIX

## A. 1 A CHALLENGING INSTANCE FOR RANDOMIZED ROUNDING

In this section, we discuss a specific limitation of the rounding algorithm in previous works (see [21] and [15]). Essentially, we explain why their techniques cannot give a better than $\Theta\left(\frac{\log (n)}{\log \log (n)}\right)$ approximation. This argument is based on an instance that we build below, which also illustrates why the bottom-to-top pruning cannot fix the issue by itself. We reuse here the notation defined in Section 3.2. For clarity, we restate here the LP relaxation of the problem (for a single source $s$ ), called the path LP.

$$
\begin{aligned}
\sum_{q \in C(p)} x(q) & =k \cdot x(p) & \forall p \in P \backslash T \\
\sum_{q \in I(v) \cap D(p)} x(q) & \leqslant x(p) & \forall p \in P, \forall v \in V \\
x((s)) & =1 & \forall s \in S \\
x & \geqslant 0 &
\end{aligned}
$$

Below we give a construction of a random instance inspired by [51]. The graph will have $h+1$ layers labeled from o to $h$ and a single source $s$ located in $L_{0}$. $n$ will be the number of vertices in the last layer $L_{h}$ and we choose $h=\Theta\left(\frac{\log n}{\log \log n}\right)$ so that $\left(\log ^{10} n\right)^{h}=n$. The set of sinks corresponds exactly to the vertices in the last layer $L_{h}$. The graph until layer $h-1$ is a complete and regular tree where every internal vertex has exactly $\log ^{20} n$ children. Vertices in layer $L_{h-1}$ will be called leaves. Next, we describe how we connect layer $h-1$ to layer $h$.

Each vertex $v \in L_{\ell}$ (note that $v$ is a sink) runs the following sampling process, independently of other sinks. Start at the source $s$. From the source select each vertex $u$ in $L_{1}$ independently at random with probability $1 / \log ^{10} n$. All the vertices in $L_{2}$ whose parent was selected is then selected independently at random with probability $1 / \log ^{10} n$. We repeat layer by layer, each vertex in layer $L_{i}$ whose parent was selected is then selected independently at random with probability $1 / \log ^{10} n$.

Finally, connect the $\operatorname{sink} v \in L_{h}$ to all vertices in layer $L_{h-1}$ that were selected by the above random process initiated by $v$.

In the following we denote by $T_{u}$ for any $u \in \cup_{0 \leqslant i \leqslant h-1} L_{i}$ the subtree rooted at vertex $u$ that is contained in layers $\cup_{0 \leqslant i \leqslant h-1} L_{i}$. We first claim the following.

Claim A.1. With high probability, for any $j \leqslant h-1$ no $\operatorname{sink} v \in L_{h}$ is connected to more than

$$
\left(1+\frac{1}{\log (n)}\right) \cdot\left(\log ^{10}(n)\right)^{h-1-j}
$$

leaves in layer $\ell-1$ that belong to the subtree rooted at a vertex $u \in L_{j}$.
Proof. To see this, note that during the random process initiated by the sink $v$, every internal node $u$ that was selected selects in expectation $\log ^{20}(n) / \log ^{10}(n)=\log ^{10}(n)$ of its children. By a standard Chernoff bound the probability that this internal nodes selects more than

$$
\left(1+\frac{1}{\log ^{2}(n)}\right) \cdot \log ^{10}(n)
$$

of its children is at most

$$
\exp \left(-\frac{1}{3 \log ^{4}(n)} \cdot \log ^{10}(n)\right) \leqslant \exp \left(-\log ^{6}(n) / 3\right)
$$

Therefore with high probability, no internal node selects more than $\left(1+\frac{1}{\log ^{2} n}\right)$. $\log ^{10}(n)$ of its children which implies that the total number of selected nodes at layer $h-1$ that belong to the subtree $T_{u}$ (which are exactly the neighbors of $v$ in that subtree) is at most

$$
\left(1+\frac{1}{\log ^{2} n}\right)^{h-1-j}\left(\log ^{10} n\right)^{h-1-j} \leqslant\left(1+\frac{1}{\log (n)}\right)\left(\log ^{10} n\right)^{h-1-j}
$$

Let $u$ be a vertex in layer $L_{h-1}$. Then $u$ is connected to each vertex $v \in L_{h}$ independently with probability $1 /\left(\log ^{10} n\right)^{h-1}$, therefore with high probability each vertex $u \in L_{h-1}$ has at least

$$
\frac{n}{\left(1+\frac{1}{\log (n)}\right) \cdot\left(\log ^{10}(n)\right)^{h-1}}=\frac{\log ^{10}(n)}{1+\frac{1}{\log (n)}}
$$

many neighbors in layer $h$.
With this observation and Claim A.1, it is easy to prove the following.
Claim A.2. With high probability, the path LP is feasible on this instance for

$$
k=\frac{\log ^{10}(n)}{1+\frac{1}{\log (n)}}
$$

Proof. We set the following fractional values.

$$
x(p)= \begin{cases}(1+1 / \log n)^{-(h-1)} \cdot\left(\log ^{10} n\right)^{-(h-1)} & \text { if } p \text { ends at layer } h, \text { and } \\ (1+1 / \log n)^{-i} \cdot\left(\log ^{10} n\right)^{-i} & \text { if } p \text { ends at layer } i \leqslant h-1\end{cases}
$$

Clearly, $x((s))=1$. We can easily verify the demand constraints: let $p$ be a path that ends at layer $i \leqslant h-2$, then

$$
\begin{aligned}
& \sum_{q \in C(p)} x(q)=\frac{\log ^{20} n}{(1+1 / \log n)^{i+1} \cdot\left(\log ^{10} n\right)^{i+1}} \\
&=\frac{k}{(1+1 / \log n)^{i} \cdot\left(\log ^{10} n\right)^{i}}=k \cdot x(p)
\end{aligned}
$$

If $p$ ends at layer $h-1$, then with high probability we have

$$
\sum_{q \in C(p)} x(q) \geqslant \frac{\log ^{10}(n)}{1+\frac{1}{\log n}} \cdot \frac{1}{(1+1 / \log n)^{h-1} \cdot\left(\log ^{10} n\right)^{h-1}}=k \cdot x(p)
$$

The capacity constraints are also easily verified with high probability. Let $u$ be a vertex in layer $i \leqslant h-1$ and $p$ a path ending in layer $j \leqslant i$ then

$$
\begin{aligned}
& \sum_{q \in I(v) \cap D(p)} x(q) \leqslant \frac{1}{(1+1 / \log n)^{i} \cdot\left(\log ^{10} n\right)^{i}} \\
& \leqslant \frac{1}{(1+1 / \log n)^{j} \cdot\left(\log ^{10} n\right)^{j}}=x(p)
\end{aligned}
$$

If $v$ is a vertex in layer $h$ and $p$ a path ending in layer $i \leqslant h-2$ then, using Claim A.1, we obtain

$$
\begin{aligned}
& \sum_{q \in I(v) \cap D(p)} x(q) \\
& \quad \leqslant\left(1+\frac{1}{\log n}\right) \cdot\left(\log ^{10} n\right)^{h-1-i} \cdot \frac{1}{(1+1 / \log n)^{h-1} \cdot\left(\log ^{10} n\right)^{h-1}} \\
& \quad \leqslant \frac{1}{(1+1 / \log n)^{h-2} \cdot\left(\log ^{10} n\right)^{i}} \\
& \quad \leqslant x(p)
\end{aligned}
$$

If $v$ is a vertex in layer $h$ and $p$ a path ending in layer $i=h-1$ then

$$
\sum_{q \in I(v) \cap D(p)} x(q) \leqslant \frac{1}{(1+1 / \log n)^{h-1} \cdot\left(\log ^{10} n\right)^{h-1}}=x(p)
$$

Given Claim A.2, one might be tempted to run the intuitive randomized rounding as in previous works. It is easy to show that this rounding will select a tree in which every internal node will have between ( $1-$ $2 / \log n) \cdot \log ^{10} n$ and $(1+2 / \log n) \cdot \log ^{10} n$ children with high probability. Note that this implies that this tree will have at least $(1-2 / \log n)^{h-1}$. $\left(\log ^{10} n\right)^{h-1}=\Omega\left(\log ^{10(h-1)} n\right)$ leaves in layer $L_{h-1}$. These leaves will each demand $k=\Omega\left(\log ^{10}(n)\right)$ sinks. Hence the total number of sinks needed will be $\Omega\left(\log ^{10 h} n\right)=\Omega(n)$. We prove the following last claim.

Claim A.3. With high probability, any subtree $T$ rooted at $s$ with leaves at layer $h-1$ and in which every internal node has between

$$
\left(1-\frac{2}{\log n}\right) \log ^{10}(n)
$$

and

$$
\left(1+\frac{2}{\log n}\right) \log ^{10}(n)
$$

many children is connected to at most

$$
O\left(\frac{\left(\log ^{10} n\right)^{h}}{h}\right)=O\left(\frac{n}{h}\right)
$$

many sinks in layer $L_{h}$.
Before giving a proof sketch of this, we explain why this shows that some kind of pruning is needed. As explained above, the intuitive randomized rounding will select a tree in which every node has roughly $\log ^{10}(n)$ children with high probability. This tree needs to be connected to $\Omega(n)$ sinks in layer $L_{h}$ in order to give $\Omega(k)$ children to every leaf in $L_{h-1}$. By Claim A.3, only a $O(1 / h)$ fraction of these sinks are available to these leaves, yielding an approximation factor of $\Omega(h)=\Omega(\log n / \log \log n)$ on the vertices in $L_{h-1}$. Hence without any post-processing that can modify the tree, it seems impossible to break through this $\tilde{\Omega}(\log n)$ approximation factor. Moreover, this instance shows that the post-processing cannot be only local: a simple post-processing one could think of is to remove vertices with high congestion and hope that this does not remove too many children of any other vertex. However, it does not suffice to remove only children of $L_{h-1}$ in this example. Instead one needs to remove vertices from many layers. A careful reader may notice that if instead of sampling $k$ children for every node (as it is the case in previous works) we allow ourselves to sample only $k / 2$ children (which is equivalent to sample $k$ children and then perform our top-to-bottom pruning); then the necessary number of sinks in $L_{h}$ drops by a factor of $2^{-h}$. This factor seems to be enough to overcome the issue highlighted by our construction.

We finish this discussion by giving a proof sketch of Claim A.3.
Proof sketch for Claim A.3. Let us call $d$ the number of children of each node in a fixed tree $T$ and assume for simplicity that $d=\log ^{10}(n)$ for all internal nodes in the tree. For any vertex $u \in L_{j} \cap T$ with $j \leqslant h-1$, we will denote by $T_{u}$ the subtree of $T$ rooted at $u$. We say that a vertex $v \in L_{h}$ survives in $T_{u}$ if at least one leaf of $T_{u}$ is connected to $v$. We define

$$
p_{j}:=\mathbb{P}\left(v \text { survives in } T_{u} \mid v \text { selected } u\right)
$$

for any $u \in L_{j} \cap T$. Note that by symmetry this probability does not depend on which vertex $u \in T \cap L_{j}$ and which $\operatorname{sink} v$ we select, hence the simplified notation. We will now upper bound $p_{j}$ depending on the depth $j$. We note that for a vertex $v$ to survive in $T_{u}$, it must be that there is at least one child $u^{\prime}$ of $u$ such that $v$ selects $u^{\prime}$ and $v$ survives in $T_{u^{\prime}}$. The probability that $v$ survives in $T_{u}$ conditioned by the fact that $v$ selected $u$ can be written as

$$
\begin{aligned}
& \mathbb{P}\left[v \text { survives in } T_{u} \mid v \text { selected } u\right] \\
& =1-\prod_{u^{\prime} \text { child of } u \text { in the tree } T_{u}}\left(1-\frac{P\left[v \text { survives in } T_{u^{\prime}} \mid v \text { selected } u^{\prime}\right]}{\log ^{10} n}\right)
\end{aligned}
$$

which by symmetry is equivalent to

$$
p_{j}=1-\left(1-\frac{p_{j+1}}{\log ^{10} n}\right)^{d}
$$

When $n$ goes to infinity, this is roughly equal to

$$
1-e^{-\frac{d p_{j+1}}{\log ^{10} n}} \approx 1-e^{-p_{j+1}}
$$

From there, it is easy to prove by induction that

$$
p_{j} \leqslant \frac{O(1)}{h-j}
$$

Hence the expected number of vertices of the last layer that survive in any fixed tree $T$ rooted at the source is at most

$$
p_{0} \cdot n=O\left(\frac{n}{h}\right)
$$

By independence and a standard Chernoff bound, one can show that with probability at most

$$
\exp \left(-\Omega\left(\frac{n}{h}\right)\right)
$$

more than $\frac{10 n}{h}$ sinks survive in a fixed tree $T$. To upper bound the number of possible trees, note that it suffices so select the set of $\left(\log ^{10}(n)\right)^{h-1}$ leaves among the $\left(\log ^{20}(n)\right)^{h-1}$ possible leaves. Hence there are at most

$$
\left(\log ^{20}(n)\right)^{(h-1) \cdot\left(\log ^{10} n\right)^{h-1}}=\exp \left(\Theta\left(\frac{n \cdot h \cdot \log \log (n)}{\log ^{10} n}\right)\right)=\exp \left(o\left(\frac{n}{h}\right)\right)
$$

possible trees with internal degree $\log ^{10}(n)$. By a standard union bound, this happens for no such tree.

## A. 2 PREPROCESSING THE PATH LP

We recall the goal is to obtain, from a feasible fractional solution $x$ to the path LP a multiset of paths $P^{\prime}$ containing $(s)$ for each $s \in S$ and satisfying

$$
\begin{align*}
\sum_{q \in C_{P^{\prime}}(p)} y(q) & =\frac{k}{4} \cdot y(p) & \forall p \in P^{\prime} \backslash T^{\prime}  \tag{A.1}\\
\sum_{q \in I_{P^{\prime}}(v) \cap D_{P^{\prime}}(p)} y(q) & \leqslant 2 y(p) & \forall p \in P^{\prime}, v \in V  \tag{A.2}\\
y(p) & =\frac{1}{\left(4 \log ^{2} n\right)^{i}} & \forall i \leqslant h, \forall p \in L_{i} \tag{А.3}
\end{align*}
$$

The value of $y$ is fixed by the distance to the sources; hence, the only challenge in this part is in fact to select the multiset $P^{\prime}$.

Towards this, consider a randomized rounding algorithm, that proceeds layer by layer starting from the sources. We will produce partial solutions $P^{(i)}$ that correspond to our final solution for all paths in $L_{\leqslant i}$. Recall that this is a multiset and there might be several times the same copy of a path $p \in P$, but every copy will have a unique parent. Initially, we add a single copy of the trivial path $(s)$ for each $s \in S$ to $P^{(0)}$. Assume now that we defined the solutions up to $P^{(i)}$, that is, we sampled up to layer $i$ for some $i \geqslant 0$. For
each path $p$ that is open and belongs to $P^{(i)} \cap L_{i}$, we sample $(k / 4) \cdot\left(4 \log ^{2} n\right)$ many children paths where the $j$-th child equals path $q \in C_{P}(p) \cap L_{i+1}$ with probability

$$
\frac{x(q)}{\sum_{q^{\prime} \in C_{P}(p)} x\left(q^{\prime}\right)} .
$$

For each path $q$ that is sampled, we set the parent of $q$ to be $p$ (which is a unique copy). Let $Q_{q}$ be the multiset of all copies of a path $q \in P$ that were sampled in this way. Then we set

$$
P^{(i+1)} \leftarrow P^{(i)} \cup\left(\cup_{q \in L_{i+1}} Q_{q}\right)
$$

We repeat this process until reaching the last layer $h$ and we set $P^{\prime}=P^{(h)}$. We argue below that we obtain the desired properties with high probability. To this end, we notice that each open path $p \in P^{(i)} \cap L_{i}$ samples exactly $(k / 4) \cdot\left(4 \log ^{2} n\right)$ children paths $q$ and hence we obtain

$$
\sum_{q \in C_{P^{\prime}}(p)} y(q)=\frac{(k / 4) \cdot\left(4 \log ^{2} n\right)}{\left(4 \log ^{2} n\right)^{i+1}}=\frac{k}{4} y(p)
$$

Therefore constraint (A.1) is satisfied with probability 1 . It remains to verify that constraint (A.2) holds with high probability. For this, we consider the following quantity, that keeps track of the (possibly fractional) congestion induced by the descendants of a path $p \in P^{(i)}$ on a vertex $v$. If $v \in L_{\leqslant i}$, then this induced congestion cannot change anymore so we will only keep track of this congestion for $v \in L_{\geqslant i+1}$. For ease of notation, let us define for all $q \in P$,

$$
\operatorname{cong}(q, v):=\sum_{q^{\prime} \in D_{P}(q) \cap I_{P}(v)} x\left(q^{\prime}\right) .
$$

Then, we define for any $p \in P^{(i)}$

$$
\operatorname{cong}\left(p, v \mid P^{(i)}\right):=\sum_{q \in D_{P^{(i)}}(p) \cap L_{i}} \frac{y(q)}{x(q)} \operatorname{cong}(q, v)
$$

which intuitively is the fractional congestion induced by descendants of $p$ on vertex $v$ but where we condition by what happened until layer $i$. Note that if $v \in L_{j}$ the quantity cong $\left(p, v \mid P^{(j)}\right)$ is the final quantity we need to bound in order to ensure constraint (A.2). Indeed, in this case we obtain that

$$
\operatorname{cong}\left(p, v \mid P^{(j)}\right)=\sum_{q \in D_{P^{(j)}}(p) \cap I_{P^{(j)}}(v)} \frac{y(q)}{x(q)} \cdot x(q)=\sum_{q \in D_{P^{(j)}}(p) \cap I_{P^{(j)}}(v)} y(q)
$$

Claim A.4. For any $p \in P^{(i)} \cap L_{i}$, we have that

$$
\operatorname{cong}\left(p, v \mid P^{(i)}\right) \leqslant y(p)
$$

with probability 1.

Proof. If $p \in P^{(i)} \cap L_{i}$, then we have that

$$
\operatorname{cong}\left(p, v \mid P^{(i)}\right)=y(p) \frac{\operatorname{cong}(p, v)}{x(p)} \leqslant y(p)
$$

using constraint (3.3) of the original path LP.
Claim A.5. For any open $p \in P^{(i)}$ and any $i \geqslant 0$, we have that

$$
\mathbb{P}\left[\operatorname{cong}\left(p, v \mid P^{(i+1)}\right) \geqslant \frac{\operatorname{cong}\left(p, v \mid P^{(i)}\right)}{2}+y(p)\right] \leqslant \frac{1}{n^{2 \log (n)}}
$$

where the probability is taken over the randomness to round layer $L_{i+1}$.
Proof. Denote by $N_{q}$ the number of copies of $q \in L_{i+1} \cap D_{P}(p)$ that are sampled as descendants of $p$. We also denote $P(q)$ the parent of $q$ in the set $P$. Accordingly, $N_{P(q)}$ is the number of copies of the path $P(q)$ that are chosen in $P^{(i)}$ as descendants of $p$. First, we compute

$$
\begin{aligned}
\mathbb{E}[ & \left.\operatorname{cong}\left(p, v \mid P^{(i+1)}\right)\right]=\sum_{q \in D_{P}(p) \cap L_{i+1}} \mathbb{E}\left[N_{q}\right] \cdot \frac{1}{\left(4 \log ^{2} n\right)^{i+1} \cdot x(q)} \cdot \operatorname{cong}(q, v) \\
& =\sum_{q \in D_{P}(p) \cap L_{i+1}} N_{P(q)} \cdot \frac{(k / 4) \cdot\left(4 \log ^{2} n\right)}{\left(4 \log ^{2} n\right)^{i+1}} \cdot \frac{x(q)}{\sum_{q^{\prime} \in C_{P}(P(q))} x\left(q^{\prime}\right)} \cdot \frac{\operatorname{cong}(q, v)}{x(q)} \\
& =\sum_{q \in D_{P}(p) \cap L_{i+1}} N_{P(q)} \cdot(k / 4) \cdot \frac{1}{\left(4 \log ^{2} n\right)^{i}} \cdot \frac{1}{\sum_{q^{\prime} \in C_{P}(P(q))} x\left(q^{\prime}\right)} \cdot \operatorname{cong}(q, v) \\
& =\sum_{q \in D_{P}(p) \cap L_{i+1}} N_{P(q)} \cdot(k / 4) \cdot y(P(q)) \cdot \frac{1}{k \cdot x(P(q))} \cdot \operatorname{cong}(q, v) \\
& =\frac{1}{4} \sum_{q \in D_{P}(p) \cap L_{i+1}} N_{P(q)} \cdot \frac{y(P(q))}{x(P(q))} \cdot \operatorname{cong}(q, v) \\
& =\frac{1}{4} \sum_{q^{\prime} \in D_{P}(p) \cap L_{i}} N_{q^{\prime}} \cdot \frac{y\left(q^{\prime}\right)}{x\left(q^{\prime}\right)} \cdot \sum_{q \in C_{P}\left(q^{\prime}\right)} \operatorname{cong}(q, v) \\
& =\frac{1}{4} \sum_{q^{\prime} \in D_{P}(p) \cap L_{i}} N_{q^{\prime}} \cdot \frac{y\left(q^{\prime}\right)}{x\left(q^{\prime}\right)} \cdot \operatorname{cong}\left(q^{\prime}, v\right)=\frac{\operatorname{cong}\left(p, v \mid P^{(i)}\right)}{4} .
\end{aligned}
$$

where we used constraint (3.2) of the original path LP to obtain the fourth line, and standard algebraic manipulations for the rest. Second, we note that $\operatorname{cong}\left(p, v \mid P^{(i+1)}\right)$ can be written as a sum of independent random variables of value

$$
\frac{y(q)}{x(q)} \cdot \operatorname{cong}(q, v)
$$

for some $q \in L_{i+1}$. By constraint (3.3) of the path LP, the absolute value of these variables is never more than $1 /\left(100 \log ^{2} n\right)^{i+1}$. By a standard Chernoff bound, we can conlude that

$$
\begin{aligned}
\mathbb{P}\left[\operatorname{cong}\left(p, v \mid P^{(i+1)}\right) \geqslant \frac{\operatorname{cong}\left(p, v \mid P^{(i)}\right)}{2}+y(p)\right] & \leqslant \exp \left(-\frac{y(p)\left(4 \log ^{2} n\right)^{i+1}}{2}\right) \\
& \leqslant n^{-2 \log (n)}
\end{aligned}
$$

There are at most $n^{h} \leqslant n^{\log (n)+1}$ constraints to maintain to ensure constraint (A.2) over at most $h$ rounds of rounding. Hence by Claim A.5, we have $\operatorname{cong}\left(p, v \mid P^{(i+1)}\right) \leqslant \operatorname{cong}\left(p, v \mid P^{(i)}\right) / 2+y(p)$ for all paths $p$, vertices $v$, and rounds $i$ with probability at least $1-n^{-\log (n)+1}$. Using Claim A. 4 with this fact we obtain that for any path $p$ and vertex $v$

$$
\sum_{q \in I_{P^{\prime}}(v) \cap D_{P^{\prime}}(p)} y(q) \leqslant y(p) \sum_{j=0}^{\infty} \frac{1}{2^{j}} \leqslant 2 y(p)
$$

This proves the desired result.

## B. 1 REDUCTION TO HYPERGRAPH MATCHING PROBLEM

## B.1.1 Solving the Configuration LP

The goal of this section is to prove Theorem 4.1. We consider the dual of the configuration LP (after adding an artificial minimization direction $\min 0^{T} x$ ).

$$
\begin{aligned}
\max & \sum_{i \in P} y_{i}
\end{aligned}-\sum_{j \in R} z_{j} .
$$

Observe that the optimum of the dual is either 0 obtained by $y_{i}=0$ and $z_{j}=0$ for all $i, j$ or it is unbounded: If it has any solution with $\sum_{i \in P} y_{i}-\sum_{j \in R} z_{j}>$ 0 , the variables can be scaled by an arbitrary common factor to obtain any objective value. If it is unbounded, this can therefore be certified by providing a feasible solution $y, z$ with

$$
\begin{equation*}
\sum_{i \in P} y_{i}-\sum_{j \in R} z_{j} \geqslant 1 \tag{*}
\end{equation*}
$$

We approximate the dual in the variant with constraint $(*)$ instead of a maximization direction using the ellipsoid method. The separation problem of the dual is as follows. Given $z_{j}$, $y_{i}$ find a player $i$ and set $C$ with $g\left(C \cap \Gamma_{i}\right) \geqslant T$ such that $\sum_{j \in C} z_{j}<y_{i}$.

To this end, consider the related problem of maximizing a monotone submodular function subject to knapsack constraints. In this problem we are given a monotone submodular function $g$ over a ground set $E$ and the goal is to maximize $g\left(E^{\prime}\right)$ over all $E^{\prime} \subseteq E$ with $\sum_{j \in E^{\prime}} a_{j} \leqslant b$. Here $a_{j} \geqslant 0$ is a weight associated with $j \in E$ and $b$ is a capacity. For this problem Srividenko gave a polynomial time ( $1-1 / e$ )-approximation algorithm [80]. It is not hard to see that this can be used to give a constant approximation for the variation where strict inequality is required in the knapsack constraint: Assume w.l.o.g. that $0<a_{j}<b$ for all $j$. Then run Srivideko's algorithm to find a set $E^{\prime}$ with $\sum_{j \in E^{\prime}} a_{j} \leqslant b$. Notice that $g\left(E^{\prime}\right)$ is at least $(1-1 / e) \mathrm{OPT}$, also when OPT is the optimal value with respect to strict inequality. If $E^{\prime}$ contains only one element then equality in the knapsack constraint cannot hold and we are done. Otherwise, split $E^{\prime}$ into two arbitrary non-empty parts $E^{\prime \prime}$ and $E^{\prime \prime \prime}$. It follows that $\sum_{j \in E^{\prime \prime}} a_{j}<b$ and $\sum_{j \in E^{\prime \prime \prime}} a_{j}<b$. Moreover, either $g\left(E^{\prime \prime}\right) \geqslant g\left(E^{\prime}\right) / 2$ or $g\left(E^{\prime \prime}\right) \geqslant g\left(E^{\prime}\right) / 2$. Hence, this method yields a $c$-approximation for $c=(1-1 / e) / 2$. We now demonstrate how to use this to find a $c$-approximation to the configuration LP.

Let OPT be the optimum of the configuration LP. It suffices to solve the problem of finding for a given $T$ either a solution of value $c T$ or deciding that $T>$ OPT. This can then be embedded into a standard dual approximation framework. We run the ellipsoid method on the dual of the configuration LP with objective value $c T$ and constraint $(*)$. This means we have to solve the separation problem. Let $z, y$ be the variables at some state. We first check whether ( $*$ ) is satisfied, that is $\sum_{i \in P} y_{i}-\sum_{j \in R} z_{j} \geqslant 1$. If not, we return this inequality as a separating hyperplane. Hence, assume $(*)$ is satisfied and our goal is to find a violated constraint of the form $\sum_{j \in C} z_{j}<y_{i}$ for some $i \in P$ and $C \in \mathcal{C}(i, T)$. For each player $i$ we maximize $f$ over all $S \subseteq \Gamma_{i}$ with $\sum_{j \in S} z_{j}<y_{i}$. We use the variant of Srividenko's algorithm described above to obtain a $c$-approximation for each player. If for one player $i$ the resulting set $S$ satisfies $f(S) \geqslant c T$, then we have found a separating hyperplane to provide to the ellipsoid method. Otherwise, we know that $f(S)<T$ for all players $i$ and $S \subseteq \Gamma_{i}$ with $\sum_{j \in S} z_{j}<y_{i}$. In other words, for all players $i$ and all $C \in \mathcal{C}(i, T)$ it holds that $\sum_{j \in C} z_{j} \geqslant y_{i}$, i.e., $z, y$ is feasible for objective value $T$ and hence OPT $>T$. If the ellipsoid method terminates without concluding that OPT $>T$, we can derive a feasible primal solution with objective value $c T$ : The configurations constructed for separating hyperplanes suffice to prove that the dual is bounded. These configurations can only be polynomially many by the polynomial running time of the ellipsoid method. Hence, when restricting the primal to these configurations it must remain feasible. To obtain the primal solution we now only need to solve a polynomial size linear program. This concludes the proof of Theorem 4.1.

## B.1.2 Clusters

This section is devoted to prove Lemma 4.2. The arguments are similar to those used in [14].

Lemma B.1. Let $x^{*}$ be a solution to the configuration $L P$ of value $T^{*}$. Then $x^{*}$ can be transformed into some $x_{i, C}^{\prime} \geqslant 0$ for $i \in P, C \in \mathcal{C}_{t}\left(i, T^{*}\right)$ which satisfies the following. There is a partition of the players into clusters $K_{1} \cup \cdots \cup K_{k} \cup Q=$ $P$ that satisfy the following.

1. any thin resource $j$ is fractionally assigned at most once, that is,

$$
\sum_{i \in P} \sum_{C \in \mathcal{C}_{t}\left(i, T^{*}\right): j \in C} x_{i, C}^{\prime} \leqslant 1
$$

We say that the congestion on resource $j$ is at most 1 .
2. every cluster $K_{j}$ gets at least $1 / 2$ thin configurations in $x^{\prime}$, that is,

$$
\sum_{i \in K_{j}} \sum_{C \in \mathcal{C}_{t}\left(i, T^{*}\right)} x_{i, C}^{\prime} \geqslant 1 / 2
$$

3. given any $i_{1} \in K_{1}, i_{2} \in K_{2}, \ldots, i_{k} \in K_{k}$ there is a matching of fat resources to players $P \backslash\left\{i_{1}, \ldots, i_{k}\right\}$ such that each of these players $i$ gets a unique fat resource $j \in \Gamma_{i}$.

The role of the set of players $Q$ in the lemma above is that each of them gets one fat resource for certain.

Proof. We first transform the solution $x^{*}$ as follows. For every configuration $C$ (for player $i$ ) that contains at least one fat resource and such that $x_{i, C}^{*}>0$, we select arbitrarily one of these fat resources $j$ and we set $x_{i,\{j\}}^{*}=x_{i, C}^{*}$ and then we set $x_{i, C}^{*}=0$. It is clear that this does not increase the congestion on resources and now every configuration that has non-zero value is either a thin configuration or a singleton containing one fat resource. Therefore we can consider the bipartite graph $G$ formed between the players and the fat resources where there is an edge between player $i$ and fat resource $j$ if the corresponding configuration $C=\{j\}$ is of non zero value (i.e. $x_{i, C}^{*}>0$ ). The value of such an edge will be exactly the value $x_{i, C}^{*}$. We now make $G$ acyclic by doing the following operation until there exists no cycle anymore. Pick any cycle (which must have even length since the graph is bipartite) and increase the coordinate of $x^{*}$ corresponding to every other edge in the cycle by a small constant. Decrease the value corresponding to the remaining edges of the cycle by the same constant. This ensures that fat resources are still (fractionally) taken at most once and that the players still have one unit of configurations fractionally assigned to them. We continue this until one of the edge value becomes 0 or 1 . If an edge becomes 0 , delete that edge and if it becomes 1 , assign the corresponding resource to the corresponding player forever. Then delete the player and the resource from the graph and add the player to the cluster $Q$. By construction, every added player to $Q$ is assigned a unique fat resource. Notice that when we stop, each remaining player still has at least 1 unit of configurations assigned to him and every fat resource is still (fractionally) taken at most once. Hence we get a new assignment vector where the assignments of fat resources to players form a forest. We also note that the congestion on thin resources did not increase during this process (it actually only decreased either when we replace fat configurations by a singleton and when players are put into the set $Q$ and deleted from the instance). We show below how to get the clusters for any tree in the forest.

1. If the tree consists of a single player, then it trivially forms its own cluster. By feasibility of the original solution $x^{*}$, condition 2 of the lemma holds.
2. If there is a fat resource that has degree 1 , assign it to its player, add the player to $Q$ and delete both the player and resource. Continue this until every resource has a degree of at least 2 . This step adds players to cluster $Q$. By construction, every added player is assigned a unique fat resource.
3. While there is a resource of degree at least 3, we perform the following operation. Root the tree containing such a resource at an arbitrary player. Consider a resource $j$ of degree at least 3 such that the subtree rooted at this resource contains only resources of degree 2 . Because this resource must have at least 2 children in the tree $i_{1}, i_{2}, \ldots$ (which are players) and because

$$
\sum_{i \in P} \sum_{C: j \in C} x_{i, C}^{*} \leqslant 1
$$

it must be that one of the children (say $i_{1}$ ) satisfies $x_{i_{1},\{j\}}^{*} \leqslant 1 / 2$. We then delete the edge $\left(j, i_{1}\right)$ in the tree and set $x_{i_{1},\{j\}}^{*}$ to 0 .
4. Every resource now has degree exactly 2 . We form a cluster for each tree in the forest. The cluster will contain the players and fat resources in the tree. We note that in every tree, only the player at the root lost at most $1 / 2$ unit of a fat resource by the previous step in the construction. By the degree property of resources and because the graph contains no cycle, it must be that in each cluster $K$ we have $|R(K)|=|P(K)|-1$ where $|R(K)|$ is the number of resources in the cluster and $|P(K)|$ the number of players. Because each resource is assigned at most once, and because only one player in the cluster lost at most $1 / 2$ unit of a fat resource, it must be that the cumulative amount of thin configurations assigned to players in $K$ is at least

$$
|P(K)|-|R(K)|-1 / 2=1 / 2
$$

This gives the second property of the lemma. For the third property, notice that for any choice of player $i \in K$, we can root the tree corresponding to the cluster $K$ at the player $i$ and assign all the fat resources in $K$ to their only child in the tree (they all have degree 2 ). This gives the third property of the lemma.
As each of these steps individually maintained maintained a congestion of at most 1 on every thin resource, we indeed get a new solution $x^{\prime}$ and the associated clusters with the required properties.

Lemma B. 1 implies that for each cluster we need to cover only one player with a thin configuration. Then the remaining players can be covered with fat resources. We will now replace $x^{\prime}$ by a solution $x^{\prime \prime}$ which takes slightly worse configurations $\mathcal{C}_{t}\left(i, T^{*} / 5\right)$, but satisfies (2) in Lemma B. 1 with 2 instead of $1 / 2$. This can be achieved by splitting each configuration $C \in \mathcal{C}_{t}\left(i, T^{*}\right)$ in 4 disjoint parts $C_{1}, C_{2}, C_{3}, C_{4} \in \mathcal{C}_{t}\left(i, T^{*} / 5\right)$. Let $C_{1} \subseteq C$ with $f\left(C_{1}\right) \geqslant T^{*} / 5$ minimal in the sense that $f\left(C_{1} \backslash\{j\}\right)<T^{*} / 5$ for all $j \in C_{1}$. Let $j_{1} \in C_{1}$. By submodularity and because $j_{1}$ is thin it holds that

$$
f\left(C \backslash C_{1}\right) \geqslant f(C)-f\left(C_{1} \backslash\left\{j_{1}\right\}\right)-f\left(\left\{j_{1}\right\}\right) \geqslant 4 T^{*} / 5-T^{*} / 100
$$

Hence, in the same way we can select $C_{2} \subseteq C \backslash C_{1}, C_{3} \subseteq C \backslash\left(C_{1} \cup C_{2}\right)$ and $C_{4} \subseteq C \backslash\left(C_{1} \cup C_{2} \cup C_{3}\right)$. We now augment $x^{\prime}$ to $x^{\prime \prime}$ by initializing $x^{\prime \prime}$ with 0 and then for each $i$ and $C \in \mathcal{C}\left(i, T^{*}\right)$ increasing $x_{i, C_{1}}^{\prime \prime}, x_{i, C_{2}}^{\prime \prime}, x_{i, C_{3}}^{\prime \prime}$, and $x_{i, C_{4}}^{\prime \prime}$ by $x_{i, C}^{\prime}$. Here $C_{1}, C_{2}, C_{3}, C_{4} \in \mathcal{C}\left(i, T^{*} / 5\right)$ are the configurations derived from $C$ by splitting it as described above.

Finally, we sample for each cluster some $\ell \geqslant 12 \log (n)$ many configurations with the distribution of $x^{\prime \prime}$ to obtain the statement of Lemma 4.2 which we restate for convenience.

Lemma (Lemma 4.2 restated). Let $\ell \geqslant 12 \log (n)$. Given a solution of value $T^{*}$ for the configuration LP in randomized polynomial time we can find a partition of the players into clusters $K_{1} \cup \cdots \cup K_{k} \cup Q=P$ and multisets of configurations $\mathcal{C}_{h} \subseteq \bigcup_{i \in K_{h}} \mathcal{C}_{T}\left(i, T^{*} / 5\right), h=1, \ldots, k$, such that

1. $\left|\mathcal{C}_{h}\right|=\ell$ for all $h=1, \ldots, k$ and
2. Each small resource appears in at most $\ell$ configurations of $\bigcup_{h} \mathcal{C}_{h}$.
3. given any $i_{1} \in K_{1}, i_{2} \in K_{2}, \ldots, i_{k} \in K_{k}$ there is a matching of fat resources to players $P \backslash\left\{i_{1}, \ldots, i_{k}\right\}$ such that each of these players $i$ gets a unique fat resource $j \in \Gamma_{i}$.

Proof. We start with the clusters obtained with Lemma B. 1 and the solution $x^{\prime \prime}$ described above. Recall that

$$
\sum_{i \in K_{h}} \sum_{C \in C_{t}\left(i, T^{*} / 5\right)} x_{i, C}^{\prime \prime} \geqslant 2
$$

for each cluster $K_{h}$. We assume w.l.o.g. that equality holds by reducing some variables $x_{i, C}^{\prime \prime}$. Clearly then each resource is still contained in at most one configuration in total.

For each cluster $K_{h}$, we sample a configuration that contains a player in this cluster according to the probability distribution given by the values $\left\{x_{i, C}^{\prime \prime} / 2\right\}_{i \in K_{h}, C \in \mathcal{C}_{t}\left(i, T^{*} / 5\right)}$. By the assumption of equality stated above this indeed defines a probability distribution. We repeat this process $\ell$ times. We first note that for one iteration, each resource is in expectation contained in

$$
\sum_{i \in P} \sum_{C \in \mathcal{C}\left(i, T^{*} / 5\right): j \in C} x_{i, C}^{\prime \prime} / 2 \leqslant 1 / 2
$$

selected configurations. Hence in expectation all the resource are contained in $\ell / 2$ selected configurations after $\ell$ iterations. By a standard Chernoff bound (see Proposition 2.4), we have that with probability at most

$$
\exp (-\ell / 6) \leqslant 1 / n^{2}
$$

a resource is contained in more than $\ell$ configurations. By a union bound, it holds that all resources are contained in at most $\ell$ selected configurations with high probability.

## B. 2 PROPERTIES OF RESOURCE SETS

Lemma B. 2 (Lemma 4.7 restated). Consider Random Experiment 4.6 with $\ell \geqslant 300.000 \log ^{3}(n)$. For any $k \geqslant 0$ and any $C \in \mathcal{C}(\geqslant k)$ we have

$$
\frac{1}{2} \ell^{-k}|C| \leqslant\left|\mathcal{I}_{k} \cap C\right| \leqslant \frac{3}{2} \ell^{-k}|C|
$$

with probability at least $1-1 / n^{10}$.
Proof. The lemma trivially holds for $k=0$. For $k>0$, by assumption $C \in \mathcal{C}^{(\geqslant k)}$ hence $|C| \geqslant \ell^{k+3}$. Since each resource of $R=R_{0}$ survives in $R_{k}$ with probability $\ell^{-k}$ we clearly have that in expectation

$$
\mathbb{E}\left(\left|R_{k} \cap C\right|\right)=\ell^{-k}|C|
$$

Hence the random variable $X=\left|R_{k} \cap C\right|$ is a sum of independent variables of value either 0 or 1 and such that $\mathbb{E}(X) \geqslant \ell^{3}$. By a standard Chernoff bound (see Proposition 2.4), we get

$$
\begin{aligned}
\mathbb{P}\left(X \notin\left[\frac{\mathbb{E}(X)}{2}, \frac{3 \mathbb{E}(X)}{2}\right]\right) & \leqslant 2 \exp \left(-\frac{\mathbb{E}(X)}{12}\right) \\
& \leqslant 2 \exp \left(-\frac{300.000 \log ^{3}(n)}{12}\right) \leqslant \frac{1}{n^{10}}
\end{aligned}
$$

since by assumption $\ell \geqslant 300.000 \log ^{3}(n)$.
Lemma B. 3 (Lemma 4.8 restated). Consider Random Experiment 4.6 with $\ell \geqslant 300.000 \log ^{3}(n)$. For any $k \geqslant 0$ and any $C \in \mathcal{C}^{(\geqslant k)}$ we have

$$
\sum_{C^{\prime} \in \mathcal{C}^{(k)}}\left|C^{\prime} \cap C \cap R_{k}\right| \leqslant \frac{10}{\ell^{k}}\left(|C|+\sum_{C^{\prime} \in \mathcal{C}^{(k)}}\left|C^{\prime} \cap C\right|\right)
$$

with probability at least $1-1 / n^{10}$.
Proof. The expected value of the random variable
$X=\sum_{C^{\prime} \in \mathcal{C}^{(k)}}\left|C^{\prime} \cap C \cap R_{k}\right|$ is equal to

$$
\mathbb{E}(X)=\frac{1}{\ell^{k}} \sum_{C^{\prime} \in \mathcal{C}^{(k)}}\left|C^{\prime} \cap C\right|
$$

Since each resource is in at most $\ell$ configurations, $X$ is a sum of independent random variables that take value in a range $[0, \ell]$. Then by a standard Chernoff bound (see Lemma 2.4), we get

$$
\mathbb{P}\left(X \geqslant 10\left(\frac{|C|}{\ell^{k}}+\mathbb{E}(X)\right)\right) \leqslant \exp \left(-\frac{3|C|}{\ell^{k+1}}\right) \leqslant \frac{1}{n^{10}}
$$

since by assumption, $|C| \geqslant \ell^{k+3}$ and $\ell \geqslant 300.000 \log ^{3}(n)$.

We finish by the proof of the last property. As mentioned in Chapter 4, this statement is a generalization of some ideas that already appeared in [14]. However, in [14], the situation is simpler since they need to sample down the resource set only once (i.e. there are only two sets $R_{1} \subseteq R$ and not a full hierarchy of resource sets $R_{d} \subseteq R_{d-1} \subseteq \cdots \subseteq R_{1} \subseteq R$ ). Given the resource set $R_{1}$, they want to select configurations and give to each selected configuration $K$ all of its resource set $\left|K \cap R_{1}\right|$ so that no resource is assigned too many times. In our case the situation is also more complex than that since at every step the selected configurations receive only a fraction of their current resource set. Nevertheless, we extend the ideas of Bansal and Srividenko to our more general setting. We recall the main statement before proceeding to its proof.

Lemma B. $\mathbf{4}$ (Lemma 4.9 restated). Consider Random Experiment 4.6 with $\ell \geqslant 300.000 \log ^{3}(n)$. Fix $k \geqslant 0$. Conditioned on the event that the bounds in Lemma 4.7 hold for $k$, then with probability at least $1-1 / n^{10}$ the following holds for all $\mathcal{F} \subseteq \mathcal{C}(\geqslant k+1), \alpha: \mathcal{F} \rightarrow \mathbb{N}$, and $\gamma \in \mathbb{N}$ such that $\ell^{3} / 1000 \leqslant \alpha(C) \leqslant n$ for all $C \in \mathcal{F}$ and $\gamma \in\{1, \ldots, \ell\}:$ If there is a $(\alpha, \gamma)$-good assignment of $R_{k+1}$ to $\mathcal{F}$, then there is $a\left(\alpha^{\prime}, \gamma\right)$-good assignment of $R_{k}$ to $\mathcal{F}$ where

$$
\begin{equation*}
\alpha^{\prime}(C) \geqslant \ell\left(1-\frac{1}{\log (n)}\right) \alpha(C) \tag{B.1}
\end{equation*}
$$

for all $C \in \mathcal{F}$. Moreover, this assignment can be found in polynomial time.
We first provide the definitions of a flow network that allows us to state a clean condition whether a good assignment of resources exists or not. We then provide the high probability statements that imply the lemma.

For any subset of configurations $\mathcal{F} \subseteq \mathcal{C}^{(\geqslant k+1)}$, resource set $R_{k}, \alpha: \mathcal{F} \rightarrow \mathbb{N}$, and any integer $\gamma$, consider the following directed network (denoted by $\left.\mathcal{N}\left(\mathcal{F}, R_{k}, \alpha, \gamma\right)\right)$. Create a vertex for each configuration in $\mathcal{F}$ as well as a vertex for each resource. Add a source $s$ and $\operatorname{sink} t$. Then add a directed $\operatorname{arc}$ from $s$ to the vertex $C \in \mathcal{F}$ with capacity $\alpha(C)$. For every pair of a configuration $C$ and a resource $i$ such that $i \in C$ add a directed arc from $C$ to $i$ with capacity 1 . Finally, add a directed arc from every resource to the sink of capacity $\gamma$. See Figure B.1 for an illustration.


Figure B.1: The directed network and an $s-t$ cut.
We denote by

$$
\text { maxflow }\left(\mathcal{N}\left(\mathcal{F}, R_{k}, \alpha, \gamma\right)\right)
$$

the value of the maximum $s-t$ flow in $\mathcal{N}\left(\mathcal{F}, R_{k}, \alpha, \gamma\right)$.
Before delving into the technical lemmas, we provide a brief road map for the proof. First, we argue that for any subset of configurations, in the two networks induced on this subset and the consecutive resource sets (which are $R_{k}$ and $R_{k+1}$ ), the value of the maximum flow differs by approximately a factor $\ell$ (this is Lemma B. 6 stated below). Then by a union bound over all possible subsets of configurations, we say that the above argument consecutively holds with good probability. This helps us conclude that a good
assignment of the resource set $R_{k+1}$ implies that there is a good assignment of the resource set $R_{k}$. Notice that if one does not have the above argument with respect to all subsets of configurations at once, it is not necessary that a good assignment of resources must exist. In particular, we need Lemma B. 5 to show that if on all subsets of configurations the maximum flow is multiplied by approximately $\ell$ when we expand the resource set from $R_{k+1}$ to $R_{k}$, then an $(\alpha, \gamma)$-good assignment of $R_{k+1}$ implies an ( $\left.\alpha^{\prime}, \gamma\right)$-good assignment of $R_{k}$, where $\alpha^{\prime}$ is almost equal to $\ell \alpha$.

Lemma B.5. Let $\mathcal{F}$ be a set of configurations, $R^{\prime} \subseteq R, \alpha: \mathcal{F} \rightarrow \mathbb{N}$ a set of resources, $\gamma \in \mathbb{N}$, and $\varepsilon \geqslant 0$. Define

$$
\alpha^{\prime}(C)=\lfloor(1-\varepsilon) \alpha(C)\rfloor .
$$

There is an $\left(\alpha^{\prime}, \gamma\right)$-good assignment of $R^{\prime}$ to $\mathcal{F}$ if and only if for every $\mathcal{F}^{\prime} \subseteq$ $\mathcal{F}$, the maximum flow in the network $\mathcal{N}\left(\mathcal{F}^{\prime}, R^{\prime}, \alpha, \gamma\right)$ is of value at least $\sum_{C \in \mathcal{F}^{\prime}} \alpha^{\prime}(C)$. Moreover, this assignment can be found in polynomial time.

Proof. First assume there is such an $\left(\alpha^{\prime}, \gamma\right)$-good assignment. Then send a flow of $\alpha^{\prime}(C)$ from $s$ to each $C \in \mathcal{F}$. If resource $i$ is assigned to $C$, send a flow of 1 from $C$ to $i$. Finally ensure that flow is preserved at every vertex corresponding to a resource by sending the correct amount of flow to $t$. Since no resource is taken more than $\gamma$ times, this flow is feasible.
We prove the other direction by contradiction. Denote by $\mathcal{N}$ the network $\mathcal{N}\left(\mathcal{F}, R^{\prime}, \alpha^{\prime}, \gamma\right)$. If there is no good assignment satisfying the condition of the lemma then the maximum flow in $\mathcal{N}$ must be strictly less than $\sum_{C \in \mathcal{F}} \alpha^{\prime}(C)$ (otherwise consider the maximum flow, which can be taken to be integral, and give to every configuration $C$ all the resources to which they send a flow of 1 ). Then by the max-flow min-cut theorem, there exists an $s-t$ cut $S$ that has value strictly less than $\sum_{C \in \mathcal{F}} \alpha^{\prime}(C)$. Let $\mathcal{C}^{\prime}$ be the set of configurations on the side of the source in $S$. Notice that $\mathcal{C}^{\prime}$ cannot be empty by assumption on the value of the cut.
Consider the induced network $\mathcal{N}\left(\mathcal{C}^{\prime}, R^{\prime}, \alpha^{\prime}, \gamma\right)$ and the cut $S$ in it. It has a value strictly lower than $\sum_{C \in \mathcal{C}^{\prime}} \alpha^{\prime}(C)$. This, in turn implies that the cut $S$ in $\mathcal{N}\left(\mathcal{C}^{\prime}, R^{\prime}, \alpha, \gamma\right)$ has a value strictly lower than $\sum_{C \in \mathcal{C}^{\prime}} \alpha^{\prime}(C)$, since this cut does not contain any edge from the source $s$ to some configuration. Hence the maximum flow in $\mathcal{N}\left(\mathcal{C}^{\prime}, R^{\prime}, \alpha, \gamma\right)$ has a value strictly less than $\sum_{C \in \mathcal{C}^{\prime}} \alpha^{\prime}(C)$, a contradiction to the assumption in the premise.

Lemma B.6. Let $\mathcal{F} \subseteq \mathcal{C}^{\geqslant(k+1)}, \alpha: \mathcal{F} \rightarrow \mathbb{N}$ such that $\ell^{3} / 1000 \leqslant \alpha(C) \leqslant n$ for all $C \in \mathcal{F}$, and $1 \leqslant \gamma \leqslant \ell$. Denote by $\mathcal{N}$ the network $\mathcal{N}\left(\mathcal{F}, R_{k}, \ell \cdot \alpha, \gamma\right)$ and by $\tilde{\mathcal{N}}$ the network $\mathcal{N}\left(\mathcal{F}, R_{k+1}, \alpha, \gamma\right)$. Then

$$
\operatorname{maxflow}(\mathcal{N}) \geqslant \frac{\ell}{1+0.5 / \log (n)} \operatorname{maxflow}(\tilde{\mathcal{N}})
$$

with probability at least $1-1 /(n \ell)^{20|\mathcal{F}|}$.
Proof. We use the max-flow min-cut theorem that asserts that the value of the maximum flow in a network is equal to the value of the minimum $s-t$ cut in the network. Consider a minimum cut $S$ of network $\mathcal{N}$ with $s \in S$
and $t \notin S$. Denote by $c(S)$ the value of the cut. We will argue that with high probability this cut induces a cut of value at most $c(S) / \ell \cdot(1+0.5 / \log (n))$ in the network $\tilde{\mathcal{N}}$. This directly implies the lemma.

Denote by $\mathcal{C}^{\prime}$ the set of configurations of $\mathcal{F}$ that are in $S$, i.e., on the source side of the cut, and $\mathcal{C}^{\prime \prime}=\mathcal{F} \backslash \mathcal{C}^{\prime}$. Similarly consider $R^{\prime}$ the set of resources in the $s$ side of the cut and $R^{\prime \prime}=R_{k} \backslash R^{\prime}$. With a similar notation, we denote $\tilde{R}^{\prime}=R^{\prime} \cap R_{k+1}$ the set of resources of $R^{\prime}$ surviving in $R_{k+1}$; and $\tilde{R}^{\prime \prime}=R^{\prime \prime} \cap R_{k+1}$. Finally, denote by $\tilde{S}$ the cut in $\tilde{\mathcal{N}}$ obtained by removing resources of $R^{\prime}$ that do not survive in $R_{k+1}$ from $S$, i.e., $\tilde{S}=\{s\} \cup \mathcal{C}^{\prime} \cup R^{\prime}$. The value of the cut $S$ of $\mathcal{N}$ is

$$
c(S)=\sum_{C \in \mathcal{C}^{\prime \prime}} \ell \cdot \alpha(C)+e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right)+\gamma\left|R^{\prime}\right|
$$

where $e(X, Y)$ denotes the number of edges from $X$ to $Y$. The value of the cut $\tilde{S}$ in $\tilde{\mathcal{N}}$ is

$$
c(\tilde{S})=\sum_{C \in \mathcal{C}^{\prime \prime}} \alpha(C)+e\left(\mathcal{C}^{\prime}, \tilde{R}^{\prime \prime}\right)+\gamma\left|\tilde{R}^{\prime}\right|
$$

We claim the following properties.
Claim B.7. For every $C \in \mathcal{F}$, the outdegree of the vertex corresponding to $C$ in $\mathcal{N}$ is at least $\ell^{4} / 2$.

Since $C \in \mathcal{C}^{(\geqslant k+1)}$ and by Lemma $4 \cdot 7$, we clearly have that $\left|C \cap R_{k}\right| \geqslant$ $\ell^{4} / 2$.
Claim B.8. It holds that

$$
c(S) \geqslant \frac{|\mathcal{F}| \ell^{3}}{1000}
$$

We have by assumption on $\alpha(C)$

$$
\begin{aligned}
c(S)=\sum_{C \in \mathcal{C}^{\prime \prime}} \ell \cdot \alpha(C)+e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right)+\gamma\left|R^{\prime}\right| & \geqslant \sum_{C \in \mathcal{C}^{\prime \prime}} \frac{\ell^{3}}{1000}+e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right)+\gamma\left|R^{\prime}\right| \\
& \geqslant \frac{\left|\mathcal{C}^{\prime \prime}\right| \ell^{3}}{1000}+e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right)+\gamma\left|R^{\prime}\right|
\end{aligned}
$$

Now consider the case where $e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right) \leqslant\left|\mathcal{C}^{\prime}\right| \ell^{3} / 1000$. Since each vertex in $\mathcal{C}^{\prime}$ has outdegree at least $\ell^{4} / 2$ in the network $\mathcal{N}$ (by Claim B.7) it must be that $e\left(\mathcal{C}^{\prime}, R^{\prime}\right) \geqslant\left|\mathcal{C}^{\prime}\right| \ell^{4} / 2-\left|\mathcal{C}^{\prime}\right| \ell^{3} / 1000>\left|\mathcal{C}^{\prime}\right| \ell^{4} / 3$. Using that each vertex in $R^{\prime}$ has indegree at most $\ell$ (each resource is in at most $\ell$ configurations), this implies $\left|R^{\prime}\right| \geqslant\left|\mathcal{C}^{\prime}\right| \ell^{3} / 3$. Since $\gamma \geqslant 1$ we have in all cases that $e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right)+$ $\gamma\left|R^{\prime}\right| \geqslant\left|\mathcal{C}^{\prime}\right| \ell^{3} / 1000$. Hence

$$
c(S) \geqslant \frac{\left|\mathcal{C}^{\prime \prime}\right| \ell^{3}}{1000}+\frac{\left|\mathcal{C}^{\prime}\right| \ell^{3}}{1000}=\frac{|\mathcal{F}| \ell^{3}}{1000}
$$

This proves Claim B.8. We can now finish the proof of the lemma. Denote by $X$ the value of the random variable $e\left(\mathcal{C}^{\prime}, \tilde{R^{\prime \prime}}\right)+\gamma\left|\tilde{R}^{\prime}\right|$. We have that

$$
\mathbb{E}[X]=\frac{1}{\ell}\left(e\left(\mathcal{C}^{\prime}, R^{\prime \prime}\right)+\gamma\left|R^{\prime}\right|\right)
$$

Moreover, $X$ can be written as a sum of independent variables in the range $[0, \ell]$ since each vertex is in at most $\ell$ configurations and $\gamma \leqslant \ell$ by assumption. By a Chernoff bound (see Lemma 2.4) with

$$
\delta=\frac{0.5 c(S)}{\log (n) \cdot\left(c(S)-\sum_{C \in \mathcal{C}^{\prime \prime}} \alpha(C)\right)} \geqslant \frac{0.5}{\log (n)}
$$

we have that

$$
\begin{aligned}
& \mathbb{P}\left(X \geqslant \mathbb{E}(X)+\frac{0.5 c(S)}{\ell \log (n)}\right) \leqslant \exp \left(-\frac{\min \left\{\delta, \delta^{2}\right\} \mathbb{E}(X)}{3 \ell}\right) \\
& \leqslant \exp \left(-\frac{c(S)}{12 \ell^{2} \log ^{2}(n)}\right) \leqslant \exp \left(-\frac{|\mathcal{F}| \ell^{3}}{12.000 \ell^{2} \log ^{2}(n)}\right) \leqslant \frac{1}{(n \ell)^{20|\mathcal{F}|}}
\end{aligned}
$$

where the third inequality comes from Claim B. 8 and the last one from the assumption that $\ell \geqslant 300.000 \log ^{3}(n)$. Hence with probability at least $1-1 /(n \ell)^{20|\mathcal{F}|}$, we have that

$$
c(\tilde{S})=\sum_{C \in \mathcal{C}^{\prime \prime}} \alpha(C)+e\left(\mathcal{C}^{\prime}, \tilde{R}^{\prime \prime}\right)+\gamma\left|\tilde{R}^{\prime}\right| \leqslant \frac{1}{\ell} c(S)+\frac{0.5}{\ell \log (n)} c(S)
$$

We are now ready to prove Lemma 4.9. Note that Lemma B. 6 holds with probability at least $1-1 /(n \ell)^{20|\mathcal{F}|}$. Given the resource set $R_{k}$ and a cardinality $s=|\mathcal{F}|$ there are $O\left((n \ell)^{2 s}\right)$ ways of defining a network satisfying the conditions from Lemma B. $6\left((m \ell)^{s} \leqslant(n \ell)^{s}\right.$ choices of $\mathcal{F}, n^{s}$ choices for $\alpha$ and $\ell$ choices for $\gamma$ ). By a union bound, we can assume that the properties of Lemma B. 6 hold for every possible network with probability at least $1-1 / n^{10}$. Assume now there is a $(\alpha, \gamma)$-good assignment of $R_{k+1}$ to some family $\mathcal{F}$. Then by Lemma B. 5 the maxflow $\left(\mathcal{N}\left(\mathcal{F}^{\prime}, R_{k+1}, \alpha, \gamma\right)\right)$ is exactly $\sum_{C \in \mathcal{F}^{\prime}} \alpha(C)$ for any $\mathcal{F}^{\prime} \subseteq \mathcal{F}$. By Lemma B.6, this implies that $\operatorname{maxflow}\left(\mathcal{N}\left(\mathcal{F}^{\prime}, R_{k}, \ell \cdot \alpha, \gamma\right)\right)$ is at least $\ell /(1+0.5 / \log (n)) \sum_{C \in \mathcal{F}^{\prime}} \alpha(C)$. By Lemma B.5, this implies a $\left(\alpha^{\prime}, \gamma\right)$-good assignment from $R_{k}$ to $\mathcal{F}$, where

$$
\begin{aligned}
\alpha^{\prime}(C)=\lfloor\ell /(1+0.5 / \log (n))\rfloor \alpha(C) \geqslant & \ell /(1+1 / \log (n)) \alpha(C) \\
& \geqslant \ell(1-1 / \log (n)) \alpha(C) .
\end{aligned}
$$

## BIBLIOGRAPHY

[1] David Adjiashvili. "Beating approximation factor two for weighted tree augmentation with bounded costs." In: ACM Transactions on Algorithms 15.2 (2018), pp. 1-26.
[2] Anthony Alexander, Sylvia Boyd, and Paul Elliott-Magwood. On the integrality gap of the 2 -edge connected subgraph problem. Tech. rep. Citeseer, 2006.
[3] Noga Alon and Yossi Azar. "On-line Steiner trees in the Euclidean plane." In: Proceedings of SoCG (1992), pp. 337-343.
[4] Noga Alon, Shlomo Hoory, and Nathan Linial. "The Moore bound for irregular graphs." In: Graphs and Combinatorics 18.1 (2002), pp. 53-57.
[5] Chidambaram Annamalai, Christos Kalaitzis, and Ola Svensson. "Combinatorial algorithm for restricted max-min fair allocation." In: $A C M$ Transactions on Algorithms 13.3 (2017), pp. 1-28.
[6] Arash Asadpour, Uriel Feige, and Amin Saberi. "Santa claus meets hypergraph matchings." In: ACM Transactions on Algorithms 8.3 (2012), 24:1-24:9.
[7] Arash Asadpour and Amin Saberi. "An approximation algorithm for max-min fair allocation of indivisible goods." In: Proceedings of STOC (2007), pp. 114-121.
[8] Baruch Awerbuch, Yossi Azar, and Yair Bartal. "On-line generalized Steiner problem." In: Theoretical Computer Science 324.2-3 (2004), pp. 313324.
[9] Étienne Bamas, Marina Drygala, and Andreas Maggiori. "An Improved Analysis of Greedy for Online Steiner Forest." In: Proceedings of SODA (2022), pp. 3202-3229.
[10] Étienne Bamas, Marina Drygala, and Ola Svensson. "A Simple LP-Based Approximation Algorithm for the Matching Augmentation Problem." In: Proceedings of IPCO (2022), pp. 57-69.
[11] Etienne Bamas, Paritosh Garg, and Lars Rohwedder. "The Submodular Santa Claus Problem in the Restricted Assignment Case." In: Proceedings of ICALP (2021), 22:1-22:18.
[12] Étienne Bamas and Lars Rohwedder. "Better Trees for Santa Claus." In: arXiv preprint arXiv:2211.14259 (2022).
[13] Nikhil Bansal. Scheduling: Open problems old and new. Presentation at MAPSP (2017).
[14] Nikhil Bansal and Maxim Sviridenko. "The santa claus problem." In: Proceedings of STOC (2006), pp. 31-40.
[15] MohammadHossein Bateni, Moses Charikar, and Venkatesan Guruswami. "Maxmin allocation via degree lower-bounded arborescences." In: Proceedings of STOC (2009), pp. 543-552.
[16] Piotr Berman and Chris Coulston. "On-line algorithms for Steiner tree problems." In: Proceedings of STOC (1997), pp. 344-353.
[17] Ivona Bezáková and Varsha Dani. "Allocating indivisible goods." In: ACM SIGecom Exchanges 5.3 (2005), pp. 11-18.
[18] Béla Bollobás. Extremal graph theory. Courier Corporation, 2004.
[19] Jarosław Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. "Steiner tree approximation via iterative randomized rounding." In: Journal of the ACM ( $7 A C M) 60.1$ (2013), pp. 1-33.
[20] Federica Cecchetto, Vera Traub, and Rico Zenklusen. "Bridging the gap between tree and connectivity augmentation: unified and stronger approaches." In: Proceedings of STOC (2021), pp. 370-383.
[21] Deeparnab Chakrabarty, Julia Chuzhoy, and Sanjeev Khanna. "On allocating goods to maximize fairness." In: Proceedings of FOCS (2009), pp. 107-116.
[22] Moses Charikar, Chandra Chekuri, To-Yat Cheung, Zuo Dai, Ashish Goel, Sudipto Guha, and Ming Li. "Approximation algorithms for directed Steiner problems." In: Journal of Algorithms 33.1 (1999), pp. 7391.
[23] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. "Designing network protocols for good equilibria." In: SIAM fournal on Computing 39.5 (2010), pp. 1799-1832.
[24] Siu-Wing Cheng and Yuchen Mao. "Restricted Max-Min Fair Allocation." In: Proceedings of ICALP 107 (2018), 37:1-37:13.
[25] Siu-Wing Cheng and Yuchen Mao. "Restricted Max-Min Allocation: Approximation and Integrality Gap." In: Proceedings of ICALP (2019), 38:1-38:13.
[26] Joe Cheriyan, Jack Dippel, Fabrizio Grandoni, Arindam Khan, and Vishnu V Narayan. "The matching augmentation problem: a 7/4-approximation algorithm." In: Mathematical Programming 182.1 (2020), pp. 315-354.
[27] Joseph Cheriyan, Robert Cummings, Jack Dippel, and Jasper Zhu. "An improved approximation algorithm for the matching augmentation problem." In: SIAM fournal on Discrete Mathematics 37.1 (2023), pp. 163190.
[28] Joseph Cheriyan and Zhihan Gao. "Approximating (unweighted) tree augmentation via lift-and-project, part I: stemless TAP." In: Algorithmica 80.2 (2018), pp. 530-559.
[29] Joseph Cheriyan and Zhihan Gao. "Approximating (unweighted) tree augmentation via lift-and-project, part II." In: Algorithmica 80.2 (2018), pp. 608-651.
[30] Joseph Cheriyan, Howard Karloff, Rohit Khandekar, and Jochen Könemann. "On the integrality ratio for tree augmentation." In: Operations Research Letters 36.4 (2008), pp. 399-401.
[31] Herman Chernoff. "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations." In: The Annals of Mathematical Statistics (1952), pp. 493-507.
[32] Nachshon Cohen and Zeev Nutov. "A (1+ ln2)-approximation algorithm for minimum-cost 2-edge-connectivity augmentation of trees with constant radius." In: Theoretical Computer Science 489 (2013), pp. 67-74.
[33] Gérard Cornuéjols, Jean Fonlupt, and Denis Naddef. "The traveling salesman problem on a graph and some related integer polyhedra." In: Mathematical programming 33.1 (1985), pp. 1-27.
[34] Artur Czumaj and Andrzej Lingas. "On approximability of the minimumcost k-connected spanning subgraph problem." In: Proceedings of SODA (1999), pp. 281-290.
[35] Sami Davies, Thomas Rothvoss, and Yihao Zhang. "A tale of Santa Claus, hypergraphs and matroids." In: Proceedings of SODA (2020), pp. 2748-2757.
[36] Sina Dehghani, Soheil Ehsani, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Saeed Seddighin. "Greedy algorithms for online survivable network design." In: Proceedings of ICALP (2018), 152:1-152:14.
[37] Guy Even, Jon Feldman, Guy Kortsarz, and Zeev Nutov. "A 1.8 approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2." In: ACM Transactions on Algorithms 5.2 (2009), pp. 1-17.
[38] Uriel Feige. "A threshold of $\ln \mathrm{n}$ for approximating set cover." In: Journal of the ACM ( $7 A C M) 45.4$ (1998), pp. 634-652.
[39] Uriel Feige. "On allocations that maximize fairness." In: Proceedings of SODA (2008), pp. 287-293.
[40] Amos Fiat and Gerhard J Woeginger. Online algorithms: The state of the art. Vol. 1442. Springer, 1998.
[41] Samuel Fiorini, Martin Groß, Jochen Könemann, and Laura Sanità. "Approximating weighted tree augmentation via Chvátal-Gomory cuts." In: Proceedings of SODA (2018), pp. 817-831.
[42] Greg N Frederickson and Joseph Ja'ja. "On the relationship between the biconnectivity augmentation and travelling salesman problems." In: Theoretical Computer Science 19.2 (1982), pp. 189-201.
[43] Greg N Frederickson and Joseph Ja'Ja'. "Approximation algorithms for several graph augmentation problems." In: SIAM fournal on Computing 10.2 (1981), pp. 270-283.
[44] Mohit Garg, Fabrizio Grandoni, and Afrouz Jabal Ameli. "Improved Approximation for Two-Edge-Connectivity." In: Proceedings of SODA (2023), pp. 2368-2410.
[45] Mohit Garg, Felix Hommelsheim, and Nicole Megow. "Matching Augmentation via Simultaneous Contractions." In: arXiv preprint arXiv:2211.01912 (2022).
[46] Michel X Goemans, Nicholas JA Harvey, Satoru Iwata, and Vahab Mirrokni. "Approximating submodular functions everywhere." In: Proceedings of SODA (2009), pp. 535-544.
[47] Michel X Goemans and David P Williamson. "A general approximation technique for constrained forest problems." In: SIAM Journal on Computing 24.2 (1995), pp. 296-317.
[48] Fabrizio Grandoni, Afrouz Jabal Ameli, and Vera Traub. "Breaching the 2 -approximation barrier for the forest augmentation problem." In: Proceedings of STOC (2022), pp. 1598-1611.
[49] Fabrizio Grandoni, Christos Kalaitzis, and Rico Zenklusen. "Improved approximation for tree augmentation: saving by rewiring." In: Proceedings of STOC (2018), pp. 632-645.
[50] Anupam Gupta and Amit Kumar. "Greedy algorithms for steiner forest." In: Proceedings of STOC (2015), pp. 871-878.
[51] Eran Halperin, Guy Kortsarz, Robert Krauthgamer, Aravind Srinivasan, and Nan Wang. "Integrality ratio for group Steiner trees and directed Steiner trees." In: SIAM fournal on Computing 36.5 (2007), pp. 14941511.
[52] Eran Halperin and Robert Krauthgamer. "Polylogarithmic inapproximability." In: Proceedings of STOC (2003), pp. 585-594.
[53] Penny Haxell and Tibor Szabó. "Improved Integrality Gap in Max-Min Allocation: or Topology at the North Pole." In: Proceedings of SODA (2023), pp. 2875-2897.
[54] Dorit S Hochbaum and David B Shmoys. "Using dual approximation algorithms for scheduling problems theoretical and practical results." In: Journal of the ACM (FACM) 34.1 (1987), pp. 144-162.
[55] Christoph Hunkenschröder, Santosh Vempala, and Adrian Vetta. "A 4/3-Approximation Algorithm for the Minimum 2-Edge Connected Subgraph Problem." In: ACM Transactions on Algorithms 15.4 (2019). ISSN: 1549-6325.
[56] Makoto Imase and Bernard M Waxman. "Dynamic Steiner tree problem." In: SIAM Journal on Discrete Mathematics 4.3 (1991), pp. 369384.
[57] AI Impacts. Global computing capacity. 2015. URL: https://aiimpacts. org/global-computing-capacity (visited on 12/12/2022).
[58] Kamal Jain. "A factor 2 approximation algorithm for the generalized Steiner network problem." In: Combinatorica 21.1 (2001), pp. 39-60.
[59] Richard M Karp. "Reducibility among combinatorial problems." In: Complexity of computer computations. Springer, 1972, pp. 85-103.
[60] L.G. Khachiyan. "Polynomial algorithms in linear programming." In: USSR Computational Mathematics and Mathematical Physics 20.1 (1980), pp. 53-72.
[61] Samir Khuller and Ramakrishna Thurimella. "Approximation algorithms for graph augmentation." In: Journal of algorithms 14.2 (1993), pp. 214-225.
[62] Samir Khuller and Uzi Vishkin. "Biconnectivity Approximations and Graph Carvings." In: Journal of the ACM (7ACM) 41.2 (1994), 214-235.
[63] Guy Kortsarz and Zeev Nutov. "A simplified 1.5-approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2 ." In: ACM Transactions on Algorithms 12.2 (2015), pp. 1-20.
[64] Guy Kortsarz and Zeev Nutov. "LP-relaxations for tree augmentation." In: Discrete Applied Mathematics 239 (2018), pp. 94-105.
[65] Andreas Krause, Ram Rajagopal, Anupam Gupta, and Carlos Guestrin. "Simultaneous placement and scheduling of sensors." In: Proceedings of IPSN (2009), pp. 181-192.
[66] Lap Chi Lau, Ramamoorthi Ravi, and Mohit Singh. Iterative methods in combinatorial optimization. Vol. 46. Cambridge University Press, 2011.
[67] Benny Lehmann, Daniel Lehmann, and Noam Nisan. "Combinatorial auctions with decreasing marginal utilities." In: Games Econ. Behav. 55.2 (2006), pp. 270-296.
[68] Jan Karel Lenstra, David B Shmoys, and Éva Tardos. "Approximation algorithms for scheduling unrelated parallel machines." In: Mathematical programming 46.1 (1990), pp. 259-271.
[69] Tobias Mömke and Ola Svensson. "Removing and adding edges for the traveling salesman problem." In: Journal of the ACM ( $7 A C M) 63.1$ (2016), pp. 1-28.
[70] Robin A Moser and Gábor Tardos. "A constructive proof of the general Lovász local lemma." In: Journal of the ACM (7ACM) 57.2 (2010), pp. 115.
[71] Marcin Mucha. "13/9-Approximation for Graphic TSP." In: Theory of Computing Systems 55.4 (2014), pp. 640-657.
[72] Hiroshi Nagamochi. "An approximation for finding a smallest 2-edgeconnected subgraph containing a specified spanning tree." In: Discrete Applied Mathematics 126.1 (2003), pp. 83-113.
[73] Alantha Newman. "An Improved Analysis of the Mömke-Svensson Algorithm for Graph-TSP on Subquartic Graphs." In: SIAM fournal on Discrete Mathematics 34.1 (2020), pp. 865-884.
[74] Zeev Nutov. "On the tree augmentation problem." In: Algorithmica 83.2 (2021), pp. 553-575.
[75] Debmalya Panigrahi. COMPSCI 638: Graph Algorithms. 2019. URL: https : //www2.cs.duke.edu/courses/fall19/compsci638/ fall19_notes/lecture16.pdf (visited on 02/20/2023).
[76] Lukáš Poláček and Ola Svensson. "Quasi-polynomial local search for restricted max-min fair allocation." In: ACM Transactions on Algorithms 12.2 (2015), pp. 1-13.
[77] Gabriel Robins and Alexander Zelikovsky. "Improved Steiner tree approximation in graphs." In: Proceedings of SODA (2000), pp. 770-779.
[78] Petra Schuurman and Gerhard J Woeginger. "Polynomial time approximation algorithms for machine scheduling: Ten open problems." In: fournal of Scheduling 2.5 (1999), pp. 203-213.
[79] András Sebö and Jens Vygen. "Shorter tours by nicer ears: 7/5-Approximation for the graph-TSP, $3 / 2$ for the path version, and $4 / 3$ for two-edgeconnected subgraphs." In: Combinatorica 34.5 (2014), pp. 597-629.
[8o] Maxim Sviridenko. "A note on maximizing a submodular set function subject to a knapsack constraint." In: Oper. Res. Lett. 32.1 (2004), pp. 4143.
[81] V Traub and R Zenklusen. "A better-than-2 approximation for weighted tree augmentation." In: Proceedings of FOCS (2021), pp. 1-12.
[82] Vera Traub and Rico Zenklusen. "Local search for weighted tree augmentation and Steiner tree." In: Proceedings of SODA (2022), pp. 32533272.
[83] Jan Vondrák. "Optimal approximation for the submodular welfare problem in the value oracle model." In: Proceedings of STOC (2008). Ed. by Cynthia Dwork, pp. 67-74
[84] Jeffery Westbrook and Dicky C. K. Yan. "The performance of greedy algorithms for the on-line Steiner tree and related problems." In: Mathematical systems theory 28.5 (1995), pp. 451-468.
[85] David P Williamson and David B Shmoys. The design of approximation algorithms. Cambridge university press, 2011.
[86] Alexander Z Zelikovsky. "An 11/6-approximation algorithm for the network Steiner problem." In: Algorithmica 9.5 (1993), pp. 463-470.

## EDUCATION

- (2018-2023) EPFL, Lausanne, Switzerland.

PhD student in Theoretical Computer Science, Advisor: Ola Svensson.

- (2017-2018) Université Paris Diderot - École Normale Supérieure, Paris, France. M.Sc. in Theoretical Computer Science (MPRI) (summa cum laude); GPA: 18.81/20 (rank: 2/63).
- (2014-2018) École polytechnique (Engineering diploma), Paris, France.
- (2012-2014) Lycée Louis-Le-Grand, classe préparatoire MPSI/MP*, Paris, France.
- (2012) French "Baccalauréat série S" (summa cum laude), Paris, France.


## WORK EXPERIENCE

- (Nov.-Dec. 2022) Academic visit at the Simons Institute, Berkeley, USA.
- (Apr.-Aug. 2018) Research internship at the G-SCOP lab, Grenoble, France.
- (Mar.-Jul. 2017) Research and software engineering internship, INRIA, Grenoble, France.
- (Summer 2016) Internship at Surrey Satellite Technology Ltd, Guildford, England.
- (Sep. 2014-Apr. 2015) Military Service.

AWARDS AND HONORS

- (2020) NeurIPS oral presentation ( $\approx 1 \%$ acceptance rate).
- (2020) NeurIPS spotlight presentation $(\approx 4 \%$ acceptance rate).
- (2018) EPFL IC School 1-year Fellowship.
- (2015) Citation for outstanding service in my unit during military service.
- Etienne Bamas and Lars Rohwedder, Better Trees for Santa Claus, to appear in STOC '23.
- Etienne Bamas, Marina Drygala, and Ola Svensson, A Simple LP-Based Approximation Algorithm for the Matching Augmentation Problem, IPCO ' 22.
- Etienne Bamas, Marina Drygala, and Andreas Maggiori, An Improved Analysis of Greedy for Online Steiner Forest, SODA '22.
- Etienne Bamas, Paritosh Garg, and Lars Rohwedder, The Submodular Santa Claus Problem in the Restricted Assignment Case, ICALP ' 21.
- Etienne Bamas, Andreas Maggiori, and Ola Svensson, The Primal-Dual method for Learning Augmented Algorithms, NeurIPS '20.
- Etienne Bamas, Andreas Maggiori, Lars Rohwedder, and Ola Svensson, Learning Augmented Energy Minimization via Speed Scaling, NeurIPS ' 20.
- Etienne Bamas and Louis Esperet, Local Approximation of the Maximum Cut in Regular Graphs, WG ' 19.
- Etienne Bamas and Louis Esperet, Distributed Coloring of Graphs with an Optimal Number of Colors, STACS '19.

TEACHING

- (since 2021) Supervision of semester projects at EPFL. Supervised students: Alexandre Reynaud (master student), Taha El Ghazi (master student).
- (2019-2023) Teaching Assistant at EPFL ("Algorithms" (head TA), "Theory of Computation" (head TA), "Information, Calcul, Communication").
- (2017-2018) Teaching Assistant at Lycée Janson-de-Sailly for oral exams in mathematics.


## ACADEMIC SERVICE AND TALKS

- I co-organized the workshop ALPS 2022 on algorithms with predictions ( $\sim 45$ international participants).
- Reviewer for the conferences: MFCS ‘ 19, WAOA ' 20 , ITCS ' 20 , NeurIPS '21, SODA '22, ICML '22, NeurIPS ‘22, SODA '23, STOC '23, ICALP ‘ 23.
- Reviewer for the journals: Algorithmica (2021).
- I gave a talk at the following conferences: WG ' 19 , STACS ' 19 , NeurIPS '20, ICALP '21, Operations Research Bern '21, SODA ' 22.


[^0]:    2 Quasi-polynomial refers here to a running time of $n^{\text {poly }(\log n)}$, which is only slightly worse than polynomial time.

[^1]:    1 To be precise, the exact approximation ratio is $O(\log \log (m) / \log \log \log (m))$.

[^2]:    2 In fact they get a slightly better ratio of $\alpha=O(\log \log (m) / \log \log \log (m))$.

[^3]:    1 In the literature on online algorithms, we say that an algorithm is $\alpha$-competitive if it returns a solution that is never more than $\alpha$ times the optimum. For the purpose of this thesis, one might think of an $\alpha$-competitive algorithm as an $\alpha$-approximation algorithm that works in an online setting.

