

A Constant Factor Approximation Algorithm for Unsplittable Flow on Paths

Paul Bonsma* Jens Schulz† Andreas Wiese‡

January 12, 2013

Abstract

In the unsplittable flow problem on a path, we are given a capacitated path P and n tasks, each task having a demand, a profit, and start and end vertices. The goal is to compute a maximum profit set of tasks, such that for each edge e of P , the total demand of selected tasks that use e does not exceed the capacity of e . This is a well-studied problem that has been studied under alternative names, such as resource allocation, bandwidth allocation, resource constrained scheduling, temporal knapsack and interval packing.

We present a polynomial time constant-factor approximation algorithm for this problem. This improves on the previous best known approximation ratio of $O(\log n)$. The approximation ratio of our algorithm is $7 + \epsilon$ for any $\epsilon > 0$.

We introduce several novel algorithmic techniques, which might be of independent interest: a framework which reduces the problem to instances with a bounded range of capacities, and a new geometrically inspired dynamic program which solves a special case of the maximum weight independent set of rectangles problem to optimality. In the setting of resource augmentation, wherein the capacities can be slightly violated, we give a $(2 + \epsilon)$ -approximation algorithm. In addition, we show that the problem is strongly NP-hard even if all edge capacities are equal and all demands are either 1, 2, or 3.

1 Introduction

In the Unsplittable Flow Problem on a Path (UFPP), we are given a path $P = (V, E)$ with an integral capacity u_e for each edge $e \in E$. In addition, we are given a set of n tasks T where each task $i \in T$ is characterized by a *start vertex* $s_i \in V$, an *end vertex* $t_i \in V$, a *demand* $d_i \in \mathbb{N}$, and a *profit* $w_i \in \mathbb{N}$. A task i uses an edge $e \in E$ if e lies on the path from s_i to t_i . The aim is to compute a set of tasks $F \subseteq T$ with maximum total profit $\sum_{i \in F} w_i$ such that for each edge, the sum of the demands of all tasks in F that use this edge does not exceed its capacity.

The name of this problem is motivated by an interpretation as a multicommodity flow problem, where each task corresponds to a commodity. The term “unsplittable” means that

*Humboldt Universität zu Berlin, Computer Science Department, Unter den Linden 6, 10099 Berlin, Germany, bonsma@informatik.hu-berlin.de. The author was supported by DFG grant BO 3391/1-1.

†Technische Universität Berlin, Institute of Mathematics, Straße des 17. Juni 136, 10623 Berlin, Germany, jschulz@math.tu-berlin.de. The author was supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

‡Università di Roma “La Sapienza”, Via Ariosto 25, 00185 Rome, Italy, wiese@dis.uniroma1.it. The author was supported by the German Academic Exchange Service (DAAD).

the total amount of flow d_i from each commodity i has to be routed completely along the path from the source s_i to the sink t_i or not at all. There are several settings and applications in which this problem occurs, and several other interpretations of the problem. Therefore, this problem, and close variants thereof, have been studied under the names *bandwidth allocation* [10, 21, 32], *admission control* [36], *interval packing* [22] *temporal knapsack* [11], *multicommodity demand flow* [20], *unsplittable flow problem* [6, 7, 17, 19], *scheduling with fixed start and end times* [4], and *resource allocation* [8, 14, 23, 36]. In many applications, the vertices correspond to time points, and tasks have fixed start and end times. Within this time interval they consume a given amount of a common resource, of which the available amount varies over time.

UFPP is easily seen to be (weakly) NP-hard, since it contains the Knapsack problem as a special case (in case the path is just a single edge). In addition, Darmann et al. [23] show that the special case where all profits and all capacities are uniform is also weakly NP-hard. Chrobak et al. [22] strengthen this result by showing strong NP-hardness in this case. In addition, they show that the case where the profits equal the demands is strongly NP-hard. These results show that the problem admits no polynomial time approximation scheme (PTAS) unless $P = NP$. On the other hand, the special case of a single edge (Knapsack) admits an FPTAS. When the number of edges is bounded by a constant, UFPP admits a PTAS since it is a special case of Multi-Dimensional Knapsack [27].

Most of the research on UFPP has focused on two restricted cases: firstly, the special case in which all capacities are equal has been well-studied, which is also known as the *Resource Allocation Problem (RAP)* [8, 14, 23, 24, 36]. A more general special case of UFPP is given by the *No-Bottleneck Assumption (NBA)*: in that case it is required that $\max_i d_i \leq \min_e u_e$ (this holds in particular for RAP). We will denote this restriction of the problem by *UFPP-NBA*. For UFPP-NBA, a $(2 + \epsilon)$ -approximation algorithm is known [20], which matches the earlier best approximation ratio for RAP [14].

Many previous papers on UFPP partition the tasks into small and large tasks, and use different algorithmic techniques for these two groups. For a task i , denote by $b(i)$ the minimum capacity among all edges used by task i . For δ with $0 < \delta \leq 1$, we say that a task i is δ -small if $d_i \leq \delta \cdot b(i)$ holds, and δ -large otherwise. The two main algorithmic techniques that have been used in previous results are dynamic programming (for large tasks) and rounding of solutions to the linear programming relaxation of the problem (for small tasks). These techniques work well when the NBA holds.

However, there are several important obstacles that prevent these techniques to be generalized to the general case of UFPP. For example, Chakrabarti et al. [17] show that under the NBA the natural LP-relaxation of UFPP has a constant integrality gap. However, without this assumption the integrality gap can be as large as $\Omega(n)$ [17]. Moreover, the NBA implies that if all tasks are δ -large, then in any solution there can be at most $2 \lfloor 1/\delta^2 \rfloor$ tasks which use each edge. This property is useful for setting up a dynamic program; see [17] and Section 3.2.2. Without the NBA this is no longer possible.

Despite these obstacles, there are a few breakthrough results for (general) UFPP: The best known polynomial time algorithm by Bansal et al. [7] achieves an approximation factor of $O(\log n)$, thus beating the integrality gap of the natural LP-relaxation. This result has been generalized to trees by Chekuri et al. [19]. In addition, they gave a linear programming relaxation for UFPP with integrality gap $O(\log^2 n)$ [19]. Finally, Bansal et al. [6] gave a $(1 + \epsilon)$ -approximation algorithm with quasi-polynomial running time, which additionally requires that the capacities and the demands are quasi-polynomial, i.e. bounded by $2^{\text{poly} \log n}$.

Nevertheless, it remained an open question whether UFPP admits a constant factor approximation algorithm (this was asked e.g. in [7, 19]).

1.1 Our Contribution and Outline

We present the first polynomial time constant-factor approximation algorithm for the general case of UFPP. The algorithm has an approximation ratio of $7 + \epsilon$, for arbitrary $\epsilon > 0$.

To obtain this result we introduce several new algorithmic techniques, which are interesting in their own right. We develop a useful viewpoint which allows us to reduce the problem to a special case of the maximum weight independent set of rectangles problem. In addition, we design a framework which reduces the problem to solving instances where essentially the edge capacities are within a constant factor of each other. The techniques can be applied and combined in various ways. For instance, for practical purposes, we also show how our results can be used to obtain a constant factor approximation algorithm with a reasonable running time of only $O(n^4 \log n)$. We now go into more detail about these results, the new techniques we introduce, and give an outline of the paper.

Similar to many previous papers, for our main algorithm we partition the tasks into ‘small’ and ‘large’ tasks. For the small tasks our main result is as follows: For any $\gamma > 0$ and $\epsilon > 0$, we present a $(3 + \epsilon)$ -approximation algorithm for UFPP in Section 3, for the case where each task is $(1 - \gamma)$ -small. We remark that a similar result was given by Chekuri et al. [19], who gave an $O(\log(1/\gamma)/\gamma^3)$ -approximation algorithm if each task is $(1 - \gamma)$ -small. Their result also applies to trees. To prove our $(3 + \epsilon)$ -approximation, we introduce a novel framework in which the tasks are first grouped into smaller sets, according to their $b(i)$ values, such that the techniques for the NBA case can be applied. So the resulting sets can be solved via relatively standard dynamic programming, LP-rounding, and network flow techniques. (This is similar to e.g. [14, 20].) Solutions to these smaller sets leave a small amount of the capacity of each edge unused. In our framework we recombine these solutions into a feasible solution for all tasks.

Using the techniques developed for the $(3 + \epsilon)$ -approximation, in Section 3.3 we also give a result for UFPP in the setting of *resource augmentation*. We give an algorithm that computes a $(2 + \epsilon)$ -approximative solution which is feasible if we increase the capacity of each edge by a factor of $1 + \beta$, for arbitrarily small $\epsilon > 0$ and $\beta > 0$. Note that this algorithm works with arbitrary task sets and does not require the tasks to be small.

For our main approximation algorithm, it remains to handle the large tasks. For these, we present the following main result: for any integer $k \geq 2$, if all tasks are $\frac{1}{k}$ -large, we give a $2k$ -approximation algorithm in Section 4. This is based on a geometric viewpoint of the problem: we represent UFPP instances by drawing a curve in the plane determined by the edge capacities, and representing tasks by axis-parallel rectangles, that are drawn as high as possible under this curve. The demand of a task determines the height of its rectangle, and the profit of a task determines the weight of the rectangle. Using a novel geometrically inspired dynamic program, we show that in polynomial time, a maximum weight set of pairwise non-intersecting rectangles can be found. Such a set corresponds to a feasible UFPP solution. In addition, we show that when every task is $\frac{1}{k}$ -large, this solution yields a $2k$ -approximative solution for UFPP. With this dynamic program we contribute towards the well-studied problem of finding a Maximum Weight Independent Set of Rectangles (MWISR) [1, 30, 34]. Below we discuss this problem in more detail.

For our main result, we partition the tasks into $\frac{1}{2}$ -small tasks and $\frac{1}{2}$ -large tasks. For the

first group, we apply the aforementioned $(3 + \epsilon)$ -approximation algorithm. For the second group, our second algorithm gives a 4-approximation. Returning the best solution of the two yields the $(7 + \epsilon)$ -approximation algorithm. The main algorithm is summarized in Section 5. In addition, in Section 5 we show how our results can be combined to obtain a $O(n^4 \log n)$ time constant factor approximation algorithm for UFPP, and we discuss how our results carry over to the generalization from a path to cycle networks, where we obtain a $(8 + \epsilon)$ -approximation algorithm.

Finally, we give an alternative proof of the strong NP-hardness of UFPP, which shows that a different restriction also remains strongly NP-hard. In the existing NP-hardness proofs [22, 23], arbitrarily large demands are used in the reductions. In Section 6, we prove that the problem is strongly NP-hard even for the restricted case where all demands are chosen from $\{1, 2, 3\}$ and capacities are uniform (RAP). Note that in contrast to our hardness result, it is known that in the slightly more restricted case where the capacities *and demands* are uniform, the problem admits a polynomial time algorithm: In that case, Arkin and Silverberg [4] have shown that the problem can be solved in time $O(n^2 \log n)$ by minimum-cost flow computations.

We end in Section 7 with a discussion. In Section 2, notation and terminology are introduced. First, in the next subsection, we give more background on the many variants of UFPP that have been studied in the literature.

1.2 Related Results

As mentioned above, the restrictions of UFPP where demands are uniform (RAP) and where the No-Bottleneck Assumption $\max_i d_i \leq \min_e u_e$ holds (UFPP-NBA) have been well-studied, with the current best approximation algorithm for both problems being the $(2 + \epsilon)$ -approximation of Chekuri et al. [20]. Both RAP and UFPP-NBA have been generalized in various ways: in a scheduling context, one may allow more freedom for choosing the start time s_i and end time t_i of a given task i . Philips et al. [36] obtain a 6-approximation algorithm for such a generalization of RAP, by using LP-rounding techniques. For a similar generalization, where for each task one out of a set of alternatives needs to be selected, Bar-Noy et al. [8] provide a constant factor approximation algorithm using the local ratio technique. RAP and UFPP-NBA can be generalized in a different way, more related to network flows, by considering graphs other than a path. In graphs other than trees, there may be different paths possible between a terminal pair s_i, t_i . However, a *single* path has to be chosen for each selected terminal pair. This generalization of UFPP to general graphs is called the *Unsplittable Flow Problem (UFP)*, or *UFP-NBA* if the NBA applies. Baveja and Srinivasan [12] provide an $O(\sqrt{|E|})$ -approximation algorithm for UFP-NBA (on all graphs), improving on various earlier results. A simpler combinatorial algorithm with the same guarantee was subsequently given by Azar and Regev [5]. Chakrabarti et al. [17] also give an approximation algorithm for all graphs. In addition, they observed that an α -approximation algorithm for UFPP-NBA gives an $(\alpha + 1 + \epsilon)$ -approximation for UFP-NBA on cycles, for any $\epsilon > 0$. This way they gave the first constant factor approximation algorithm for both UFPP-NBA and UFP-NBA on cycles. Chekuri et al. [20] obtain a 48-approximation for UFP-NBA on trees.

In addition, many hardness results are known for UFPP (not necessarily under the NBA) generalized to various graphs: for general graphs, it is hard to approximate within a factor of $\Omega(|E|^{1-\epsilon})$ unless $P = NP$ [5], and for depth-3 trees the problem is APX-hard [29]. Hardness-of-approximation results are known even for the special case with unit demands and capacities (the *Edge Disjoint Path Problem*); see [2, 3].

When viewing UFPP as a packing problem, the corresponding *covering problem* has also been studied [15, 16]. In that case, tasks have *costs* instead of profits, and the objective is to find a minimum cost set of tasks F , such that for each edge, the sum of the demands of all tasks in F that use this edge is *at least* its capacity. Recently, Chakaravarthy et al. [15] designed a primal-dual 4-approximation algorithm for this problem.

Recall that we reduced UFPP for large tasks to a special case of the Maximum (Weight) Independent Set of Rectangles (M(W)ISR) problem. In this problem, a collection of n axis-parallel rectangles is given and the task is to find a maximum (weight) subset of disjoint rectangles. For the unweighted case of this problem, a randomized $O(\log \log n)$ -approximation is known [18]. For the weighted case, there are several $O(\log n)$ -approximation algorithms [1, 30, 34]. The algorithm by Erlebach et al. [26] gives a PTAS for the case that the ratio between height and width of rectangles is bounded (note that this does not apply in the special case that we need here for approximating UFPP). Our new dynamic programming technique might be useful for further research on this problem.

We remark that this approach for the large tasks is closely related to another well-studied variant of RAP: In adjacent resource scheduling problems, one wants to schedule a job on several machines in parallel which must be *contiguous*, that is, adjacent to each other. In other words, this is a variant of MWISR where rectangles are allowed to move vertically, within a given range. Duin and van Sluis [24] prove the decision variant of scheduling tasks on contiguous machines to be strongly NP-complete. RAP on contiguous machines has been considered under the name *storage allocation problem* (SAP), in which tasks are axis-aligned rectangles that are only allowed to move vertically. Leonardi et al. [32] provide a 12-approximation algorithm for SAP and Bar-Yehuda et al. [9] present a $(2 + \epsilon + 1/(e - 1))$ -approximation algorithm.

In this paper, we study UFPP in the setting of *resource augmentation*. This means that we find a solution which is feasible if we increase the capacity of each edge by a modest factor of $(1 + \epsilon)$. The paradigm of resource augmentation is very popular in real-time scheduling. There, the augmented resource is the speed of the machines. For instance, it is known that the natural earliest deadline first policy (EDF) is guaranteed to work on m machines with speed $2 - 1/m$ if the instance can be feasibly scheduled on m machines with unit speed [35]. In addition, a matching feasibility test is known [13]. For further examples of resource augmentation results in real-time scheduling see [25, 31].

2 Preliminaries

We assume that the vertices of the path $P = (V, E)$ are numbered $V = \{0, \dots, m\}$, and $E = \{\{i, i + 1\} \mid 0 \leq i \leq m - 1\}$. We assume that the tasks are numbered $T = \{1, \dots, n\}$. Recall that tasks are characterized by two vertices s_i and t_i with $s_i < t_i$, and positive integer demand d_i and profit w_i .

For each task $i \in T$ we denote by $P_i \subseteq E$ the edge set of the subpath of P from s_i to t_i . If $e \in P_i$, then task i is said to *use* e . For each edge e we denote by $T_e \subseteq T$ the set of tasks which use e . For a set of tasks F we define its profit by $w(F) := \sum_{i \in F} w_i$. Our objective is to find a set of tasks F with maximum profit such that $\sum_{i \in F \cap T_e} d_i \leq u_e$ for each edge e . For each task i we define its *bottleneck capacity* $b(i)$ by $b(i) := \min_{e \in P_i} u_e$. An edge e is called a *bottleneck edge* for the task i if $e \in P_i$ and $u_e = b(i)$. In addition, we define for every task i that $\ell(i) := b(i) - d_i$. The value $\ell(i)$ can be interpreted as the remaining capacity of

a bottleneck edge of i when i is selected in a solution. Consider a vertex $v \in V$ and an edge $e \in E$ with $e = \{x, x + 1\}$. We write $v < e$ (or $v > e$) if $v \leq x$ (resp. $v \geq x + 1$). For two edges $e = \{x, x + 1\}$ and $e' = \{x', x' + 1\}$ in E , we write $e < e'$ if $x < x'$ and $e \leq e'$ if $x \leq x'$. In other words, we interpret an edge $\{x, x + 1\}$ simply as a number between x and $x + 1$. Without loss of generality, we will assume throughout this paper that $u_e \geq 1$ for all edges e and $d_i \geq 1$ for all tasks i ; zero demands and capacities can easily be handled in a preprocessing step. Moreover, observe that one can easily adjust any given instance to an equivalent instance in which each vertex is either a start or an end vertex of a task. Such an adjustment can be implemented in linear time and it hence does not dominate the running times of the algorithms presented in this paper. Therefore, we will henceforth assume that $m < 2n$. Throughout this paper, we will use the notations defined above to refer to the UFPP instance currently under consideration. In the few cases where we consider multiple instances, the notations will be clear from the context.

We define an α -approximation algorithm for a maximization problem to be a polynomial time algorithm which computes a feasible solution for a given instance such that its objective value is at least $\frac{1}{\alpha}$ times the optimal value. Throughout this paper, for a subset of the tasks $F \subseteq T$, $OPT(F)$ denotes an optimal solution for the UFPP instance restricted to the task set F . The following simple fact shows how we can combine our approximation algorithms for different task subsets into one algorithm for all tasks.

Fact 2.1 *Consider a UFPP instance with task set T , and a partition $\{T_1, T_2\}$ of T . If for $i = 1, 2$, there exists an α_i -approximation algorithm for the instance restricted to the tasks in T_i , then there exists an $(\alpha_1 + \alpha_2)$ -approximation algorithm for the entire instance.*

Proof: For $i = 1, 2$, let ALG_i denote the solution returned by the approximation algorithm for the instance restricted to the tasks in T_i , and $OPT_i = OPT(T_i)$ an optimal solution. Let $OPT = OPT(T)$ denote an optimal solution for all tasks. So for $i = 1, 2$, $w(ALG_i) \geq \frac{1}{\alpha_i} w(OPT_i)$. The algorithm that returns the maximum profit solution of ALG_1 and ALG_2 has an approximation ratio of $\alpha_1 + \alpha_2$, since

$$\begin{aligned} w(OPT) &= w(OPT \cap T_1) + w(OPT \cap T_2) \leq w(OPT_1) + w(OPT_2) \\ &\leq \alpha_1 w(ALG_1) + \alpha_2 w(ALG_2) \\ &\leq \alpha_1 \max\{w(ALG_1), w(ALG_2)\} + \alpha_2 \max\{w(ALG_1), w(ALG_2)\} \\ &\leq (\alpha_1 + \alpha_2) \max\{w(ALG_1), w(ALG_2)\}. \end{aligned}$$

□

3 Small Tasks

In this section we present a $(3 + \epsilon)$ -approximation algorithm for any set of tasks which are $(1 - \gamma)$ -small (for arbitrarily small $\epsilon > 0$ and $\gamma > 0$). In our main algorithm (for a general set of tasks) we will invoke this algorithm as a subroutine for all tasks which are $\frac{1}{2}$ -small. Moreover, with a slight adjustment of introduced techniques we construct a polynomial time algorithm computing a $(2 + \epsilon)$ -approximative solution for the entire instance (not only small tasks) which is feasible if the capacities of the edges are increased by a factor $1 + \epsilon$ (resource augmentation), see Section 3.3.

Our strategy is to define groups of tasks such that the bottleneck capacities of all tasks in one group are within a certain range. This allows us to compute a feasible solution for each group, whose profit is at most a factor $3 + \epsilon'$ smaller than the profit of an optimal solution for the group. In addition, each computed solution leaves a certain amount of capacity of every edge unused. We devise a framework which combines solutions for a selection of groups into a feasible solution for the entire instance, in a way that yields a $(3 + \epsilon)$ -approximation (with an appropriate choice of ϵ' for the given ϵ).

3.1 Framework

We define the framework sketched above. We group the tasks into sets according to their bottleneck capacities. Let $\ell \in \mathbb{N}$ be a constant. We define $F^{k,\ell} := \{i \in T \mid 2^k \leq b(i) < 2^{k+\ell}\}$ for each integer k . Note that this includes negative values for k , and that at most $\ell \cdot n$ sets are non-empty (only those will be relevant later). In the sequel, we will present an algorithm which computes feasible solutions $ALG(F^{k,\ell}) \subseteq F^{k,\ell}$. These solutions will satisfy the following properties.

Definition 3.1 Consider a set $F^{k,\ell}$ and let $\alpha > 0$ and $\beta \geq 0$. A set $F \subseteq F^{k,\ell}$ is called (α, β) -approximative if

- $w(F) \geq \frac{1}{\alpha} \cdot w(OPT(F^{k,\ell}))$, and
- $\sum_{i \in F \cap T_e} d_i \leq u_e - \beta \cdot 2^k$ for each edge e such that $T_e \cap F^{k,\ell} \neq \emptyset$. (In particular, it is a feasible solution.)

An algorithm which computes (α, β) -approximative sets in polynomial time is called an (α, β) -approximation algorithm. We call the second condition the modified capacity constraint.

Our framework consists of a procedure that turns an (α, β) -approximation algorithm for each set $F^{k,\ell}$ into a $\left(\frac{\ell+q}{\ell} \cdot \alpha\right)$ -approximation algorithm for all given tasks, where q and β are chosen such that $\beta = 2^{1-q} > 0$. Later, for our resource augmentation result (see Section 3.3) we will work with $(\alpha, 0)$ -approximative sets and therefore, some of the claims below will be proven more generally, allowing β to be zero.

Lemma 3.2 (Framework) Let $\ell \in \mathbb{N}$ and $q \in \mathbb{N}$ be constants and let $\beta := 2^{1-q}$. Let the sets $F^{k,\ell}$ be defined as stated above for an instance of UFPP. Assume we are given an (α, β) -approximation algorithm for each set $F^{k,\ell}$ with running time $O(p(n))$ for a polynomial p . Then there is a $\left(\frac{\ell+q}{\ell} \cdot \alpha\right)$ -approximation algorithm with running time $O(m \cdot p(n))$ for the set of all tasks.

Now we describe the algorithm that yields Lemma 3.2. Assume that we are given an (α, β) -approximation algorithm which computes solutions $ALG(F^{k,\ell}) \subseteq F^{k,\ell}$. The key idea is that due to the unused edge-capacities of the sets $ALG(F^{k,\ell})$, the union of several of these sets still yields a feasible solution. With an averaging argument we will show further that the indices k for the sets $ALG(F^{k,\ell})$ that we want to combine can be chosen such that the resulting set is an $\left(\frac{\ell+q}{\ell} \cdot \alpha\right)$ -approximation. Formally, for each offset $c \in \{0, \dots, \ell + q - 1\}$ we define an index set $\eta(c) = \{c + i \cdot (\ell + q) \mid i \in \mathbb{Z}\}$ (the values ℓ and q that it depends on

will always be clear from the context). For each $c \in \{0, \dots, \ell + q - 1\}$ we compute the set $ALG(c) = \bigcup_{k \in \eta(c)} ALG(F^{k, \ell})$. We output the set $ALG(c^*)$ with maximum profit among all sets $ALG(c)$. In Lemma 3.5 we will prove that the resulting set is an $(\frac{\ell+q}{\ell} \cdot \alpha)$ -approximation. First, in Lemma 3.4 we will prove that each set $ALG(c)$ is feasible (using that $\beta \geq 2^{1-q}$). This requires the following property.

Proposition 3.3 *Let F be a feasible UFPP solution such that $b(i) < u$ for all $i \in F$. Then for every edge e , $\sum_{i \in F \cap T_e} d_i < 2u$.*

Proof: Let e be an edge. If $u_e < 2u$, then the claim follows immediately. Now suppose that $u_e \geq 2u$. Any task $i \in F \cap T_e$ must use an edge whose capacity is less than u . In particular, it must use either the closest edge e_L to the left of e or the closest edge e_R to the right of e whose capacity is less than u . The total demand of tasks in F using e_L is less than u , and the same holds for e_R . It follows that the total capacity used by tasks in $F \cap T_e$ must be less than $2u$. \square

Lemma 3.4 *Let $\ell \in \mathbb{N}$ and $q \in \mathbb{N}$ be constants. For each set $F^{k, \ell}$ let $ALG(F^{k, \ell})$ be a (α, β) -approximative set with $\beta \geq 2^{1-q} > 0$. Then for each $c \in \{0, \dots, \ell + q - 1\}$ the set $ALG(c) = \bigcup_{k \in \eta(c)} ALG(F^{k, \ell})$ is feasible.*

Proof: Consider a set $ALG(F^{k, \ell})$. By definition of the (α, β) -approximative sets, $ALG(F^{k, \ell})$ leaves $\beta \cdot 2^k \geq 2^{k+1-q}$ units of the capacity of every used edge free. Observe that this is at least twice the maximum bottleneck capacity of tasks in $F^{k-(\ell+q), \ell}$. Therefore, by Proposition 3.3, the set $ALG(F^{k, \ell}) \cup ALG(F^{k-(\ell+q), \ell})$ is feasible. In fact, it again leaves a fraction of the capacities free, which makes it possible to continue this argument for further sets $ALG(F^{k-i(\ell+q)})$ with $i \geq 2$, and prove that their union is feasible.

Formally, let $c \in \{0, \dots, \ell + q - 1\}$ and let e be an edge. Denote $P := \ell + q$. Let \bar{k} be the largest integer in $\eta(c)$ such that $2^{\bar{k}} \leq u_e$. For every $x \in \eta(c)$, denote

$$U_e^x = \sum_{j \in T_e \cap ALG(F^{x, \ell})} d_j.$$

Since $ALG(F^{\bar{k}, \ell})$ is an (α, β) -approximation and $2^{1-q} \leq \beta$, we have that

$$U_e^{\bar{k}} \leq u_e - \beta \cdot 2^{\bar{k}} \leq u_e - 2^{\bar{k}+1-q}.$$

For every $i \geq 1$, applying Proposition 3.3 for the modified capacities $u'_e = u_e - \beta \cdot 2^{\bar{k}-i \cdot P}$ shows that

$$U_e^{\bar{k}-i \cdot P} \leq 2(2^{\bar{k}-i \cdot P+q} - \beta \cdot 2^{\bar{k}-i \cdot P}) \leq 2^{\bar{k}+1-i \cdot P+q} - 2^{\bar{k}+2-i \cdot P-q}.$$

Summarizing, we have that

$$\begin{aligned} \sum_{j \in T_e \cap ALG(c)} d_j &= \sum_{i=0}^{\infty} U_e^{\bar{k}-i \cdot P} \\ &\leq u_e - 2^{\bar{k}+1-q} + \sum_{i=1}^{\infty} (2^{\bar{k}+1-i \cdot P+q} - 2^{\bar{k}+2-i \cdot P-q}) \\ &< u_e + \sum_{i=1}^{\infty} 2^{\bar{k}+1-(i-1) \cdot P-q} - \sum_{i=0}^{\infty} 2^{\bar{k}+1-i \cdot P-q} = u_e. \end{aligned}$$

□

Lemma 3.5 *Let $\ell \in \mathbb{N}$ and $q \in \mathbb{N}$ be constants. For each set $F^{k,\ell}$ let $ALG(F^{k,\ell})$ be a (α, β) -approximative set with $\beta \geq 0$. Then for the offset c^* which maximizes $w(ALG(c)) = \sum_{k \in \eta(c)} w(ALG(F^{k,\ell}))$ it holds that $w(ALG(c^*)) \geq \frac{\ell}{\ell+q} \cdot \frac{1}{\alpha} \cdot w(OPT)$, where OPT denotes an optimal solution of the given instance of UFPP.*

Proof: Every task is included in ℓ different sets $F^{k,\ell}$. Using this fact, we calculate that

$$\begin{aligned} \sum_{c=0}^{\ell+q-1} w(ALG(c)) &\geq \sum_{c=0}^{\ell+q-1} \sum_{k \in \eta(c)} \frac{1}{\alpha} \cdot w(OPT(F^{k,\ell})) = \\ \frac{1}{\alpha} \cdot \sum_{k \in \mathbb{Z}} w(OPT(F^{k,\ell})) &\geq \frac{1}{\alpha} \cdot \sum_{k \in \mathbb{Z}} w(OPT \cap F^{k,\ell}) = \frac{\ell}{\alpha} \cdot w(OPT). \end{aligned}$$

So there must be a value c such that $w(ALG(c)) \geq \frac{\ell}{\ell+q} \cdot \frac{1}{\alpha} \cdot w(OPT)$. In particular, this holds for c^* . □

Proof of Lemma 3.2: Lemma 3.4 shows that $ALG(c^*)$ is feasible and Lemma 3.5 shows that $ALG(c^*)$ is a $(\frac{\ell+q}{\ell} \cdot \alpha)$ -approximation. For computing $ALG(c^*)$ we need to compute the set $ALG(F^{k,\ell})$ for each relevant value k . There are at most $m\ell \in O(m)$ relevant values k . Finding the optimal offset c^* can be done in $O(m)$ steps. This yields an overall running time of $O(m \cdot p(n))$ (recall that $p(n)$ is the polynomial bounding the running time needed to compute the sets $ALG(F^{k,\ell})$). □

3.2 An Approximation Algorithm for Small Tasks

Now that we have developed the framework to translate (α, β) -approximation algorithms for the sets $F^{k,\ell}$ into an approximation algorithm for the entire instance (Lemma 3.2), it remains to present such an (α, β) -approximation algorithm. In this section, we present a $(2 + \frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm for sets $F^{k,\ell}$ in which all tasks are $(1 - 2\beta)$ -small (for arbitrarily small $\beta \geq 0$). Together with our framework of Lemma 3.2 this yields a $(3 + \epsilon)$ -approximation algorithm for UFPP for the case that all tasks are $(1 - \gamma)$ -small, for arbitrary $\epsilon > 0$ and $\gamma > 0$. To get some intuition, the reader can think of β being equal to $\min\{\epsilon, \gamma/2\}$.

Suppose we are given a set $F^{k,\ell}$ with only $(1 - 2\beta)$ -small tasks. In order to derive the mentioned $(2 + \frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm, we choose a value $\delta > 0$ and split the set into δ -small tasks (*tiny* tasks) and tasks which are δ -large but $(1 - 2\beta)$ -small (*medium* tasks). We define δ such that for the tiny tasks there is a $(\frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm, presented in the following subsection. For the medium tasks, we give a $(2, \beta)$ -approximation algorithm in Section 3.2.2.

3.2.1 An Approximation Algorithm for Tiny Tasks

We show that for given $\epsilon > 0$ and $\beta \geq 0$, there is a $\delta > 0$ such that if all tasks are δ -small, then for each set $F^{k,\ell}$ there is a $(\frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm. The key idea is to use linear programming techniques and a result by Chekuri et al. [20] about the integrality gap

of the canonical LP-relaxation of UFPP under the no-bottleneck assumption (NBA). UFPP with a task set T can be formulated in a straightforward way as an integer linear program:

$$\begin{aligned}
 IP : \quad & \max \quad \sum_{i \in T} w_i \cdot x_i \\
 & \text{s.t.} \quad \sum_{i \in T_e} x_i \cdot d_i \leq u_e \quad \forall e \in E \\
 & \quad \quad \quad x_i \in \{0, 1\} \quad \forall i \in T
 \end{aligned}$$

The LP relaxation is obtained by replacing the constraint $x_i \in \{0, 1\}$ by $0 \leq x_i \leq 1$. Chekuri et al. [20] gave an algorithm for UFPP instances which satisfy the NBA (i.e. $\max_i d_i \leq \min_e u_e$), in which all tasks are δ -small, which returns a UFPP solution that is at most a factor $f(\delta)$ worse than the optimum of the LP relaxation. Here $f(\delta)$ is a function for which the limit is 1 as δ approaches zero. In other words: The integrality gap of the canonical LP-relaxation is $1 + \epsilon$ if the NBA holds and all tasks are sufficiently small.

This result can be used for sets $F^{k,\ell}$: By definition, tasks in $F^{k,\ell}$ use only edges e with $u_e \geq 2^k$. Call these *relevant edges*. It is therefore possible to choose the value δ small enough to ensure that the NBA holds, when considering only the relevant edges and δ -small tasks in $F^{k,\ell}$. Furthermore, modifying the capacities by choosing $u'_e = u_e - \beta \cdot 2^k$ decreases the capacities of relevant edges at most by a factor $1 - \beta$. Therefore, the optimal value of the LP relaxation also becomes at most a factor $1 - \beta$ smaller. These are the key ideas to prove the following lemma:

Lemma 3.6 *For every combination of constants $\epsilon > 0$, $0 \leq \beta < 1$, and $\ell \in \mathbb{N}$, there exists a $\delta > 0$ such that if all tasks are δ -small, then for each set $F^{k,\ell}$ there is a $(\frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm.*

In the remainder of this subsection we prove Lemma 3.6. Assume we are given constants $\epsilon > 0$, $\ell \in \mathbb{N}$, and β with $0 \leq \beta < 1$. For an instance \mathcal{I} of UFPP, we denote by $LP(\mathcal{I})$ the natural LP-relaxation of the IP-formulation given above, where each constraint $x_i \in \{0, 1\}$ is replaced by $0 \leq x_i \leq 1$. By $LP_{OPT}(\mathcal{I})$ we denote the optimum value of the LP. We define $f(\delta') = \frac{1+\sqrt{\delta'}}{1-\sqrt{\delta'}-\delta'}$. The following result is proved by Chekuri et al. [20], although an exact analysis of the running time is not given. We observe that their algorithm admits a $O(n^3 \log n)$ implementation.

Lemma 3.7 *Consider a UFPP instance \mathcal{I} for which the NBA holds, in which all tasks are δ' -small, with $\delta' \leq \frac{3-\sqrt{5}}{2}$. Then in time $O(n^3 \log n)$, a feasible UFPP solution ALG for \mathcal{I} can be computed with $w(ALG) \geq f(\delta')^{-1} \cdot LP_{OPT}(\mathcal{I})$.*

Proof: The algorithm of Chekuri et al. works as follows. The tasks are partitioned into at most n groups, depending on their demands. The demands and the capacity are scaled such that a problem with uniform demands and uniform capacities is obtained. Since the demands and capacities are uniform, this can be solved optimally in time $O(n^2 \log n)$, using the algorithm by Arkin and Silverberg [4]. Then Chekuri et al. show that combining the solutions of each group yields a feasible solution. Furthermore, they have shown in [20, Corollary 3.4] that the obtained solution is at most a factor $f(\delta')$ worse than the optimal LP solution. \square

Lemma 3.8 *Let $\delta > 0$, $\delta' > 0$, $\beta \geq 0$ and $\ell \in \mathbb{N}$ be constants such that $\delta \leq \frac{1-\beta}{2^\ell}$ and $\delta' = \frac{\delta}{1-\beta} \leq \frac{3-\sqrt{5}}{2}$. If all tasks in $F^{k,\ell}$ are δ -small, then in time $O(n^3 \log n)$, a solution ALG for $F^{k,\ell}$ can be computed that is $(\frac{f(\delta')}{1-\beta}, \beta)$ -approximative.*

Proof: Consider the UFPP instance \mathcal{I} that is obtained by restricting the instance to the tasks in $F^{k,\ell}$, and only considering edges used by these tasks. So for every task i , it holds that $2^k \leq b(i) \leq 2^{k+\ell}$, and $d_i \leq \delta \cdot b(i)$. For every edge e , we have $u_e \geq 2^k$.

Now construct the instance \mathcal{I}' from \mathcal{I} by modifying the capacities as follows: $u'_e := u_e - \beta \cdot 2^k$, for each edge e . The instance \mathcal{I}' contains the same task set as \mathcal{I} , without modifications. For a task i , by $b(i)$ and $b'(i)$ we denote its bottleneck capacity in \mathcal{I} and \mathcal{I}' , respectively. Hence $b'(i) = b(i) - \beta \cdot 2^k$.

We will first argue that the algorithm from Lemma 3.7 may be applied to \mathcal{I}' . First, we note that since in each task is δ -small with respect to the original capacities, under the modified capacities each task is still δ' -small:

$$d_i \leq \delta \cdot b(i) = \frac{\delta}{(1-\beta)} \cdot (1-\beta) \cdot b(i) \leq \delta' \cdot (b(i) - \beta \cdot 2^k) = \delta' \cdot b'(i).$$

Recall that we required that $\delta' \leq \frac{3-\sqrt{5}}{2}$, so this condition of Lemma 3.7 is satisfied. Now we argue that for the modified capacities, the NBA holds. Since $\delta \leq (1-\beta)/2^\ell$, for all tasks i and edges e it holds that $d_i \leq \delta \cdot 2^{k+\ell} \leq 2^k - \beta \cdot 2^k \leq u_e - \beta \cdot 2^k = u'_e$. This shows that the NBA also holds with respect to the modified capacities.

Hence we may apply the algorithm from Lemma 3.7 to obtain a feasible solution ALG for \mathcal{I}' , with $w(ALG) \geq f(\delta')^{-1} LP_{OPT}(\mathcal{I}')$. We argue that $LP_{OPT}(\mathcal{I}') \geq (1-\beta) \cdot LP_{OPT}(\mathcal{I})$: Since for every edge e it holds that $u_e \geq 2^k$, it follows that $u'_e = u_e - \beta \cdot 2^k \geq (1-\beta)u_e$. Thus, if we take a feasible solution to $LP(\mathcal{I})$, and scale all the variables x_i by a factor $(1-\beta)$, we obtain a feasible solution to $LP(\mathcal{I}')$, in which the objective value has also been scaled by a factor $(1-\beta)$. This gives that

$$\begin{aligned} w(ALG) &\stackrel{\text{Lemma 3.7}}{\geq} f(\delta')^{-1} \cdot LP_{OPT}(\mathcal{I}') \\ &\geq f(\delta')^{-1} \cdot (1-\beta) \cdot LP_{OPT}(\mathcal{I}) \\ &\geq \left(f(\delta') \cdot \frac{1}{1-\beta} \right)^{-1} \cdot OPT(\mathcal{I}), \end{aligned}$$

where $OPT(\mathcal{I})$ is an optimal (integer) UFPP solution for the instance \mathcal{I} . \square

The proof of Lemma 3.6 now easily follows:

Proof of Lemma 3.6: Observe that the limit of $f(\delta')$ is 1 as δ' approaches zero. Hence for every $\epsilon > 0$, $0 \leq \beta < 1$ and $\ell \in \mathbb{N}$, we can choose a $\delta' > 0$ and $\delta := (1-\beta)\delta'$ such that $f(\delta') \leq 1+\epsilon$, $\delta \leq \frac{1-\beta}{2^\ell}$, and $\delta' \leq \frac{3-\sqrt{5}}{2}$. Then Lemma 3.8 shows that if all tasks are δ -small, then in time $O(n^3 \log n)$, a solution ALG for $F^{k,\ell}$ can be computed that is $(\frac{1+\epsilon}{1-\beta}, \beta)$ -approximative. \square

3.2.2 An Approximation Algorithm for Medium Tasks

It remains now to find a $(2, \beta)$ -approximation algorithm for tasks in $F^{k,\ell}$ that are both δ -large and $(1-2\beta)$ -small (for the δ we obtained from Lemma 3.6), for arbitrarily small $\beta \geq 0$. When

restricting to sets $F^{k,\ell}$ with only δ -large tasks, the essential property is that for any edge e and any feasible solution F , there are at most $O\left(\frac{2^\ell}{\delta}\right)$ tasks in F that use e . This property allows for a straightforward dynamic program to be used to compute an optimal solution (see e.g. [14, 17]). This can be turned into a $(2, \beta)$ -approximate solution: since tasks are $(1 - 2\beta)$ -small, it can be shown that in polynomial time, any solution F can be partitioned into two sets which are both feasible for the modified capacities. Using these ideas, we will prove the following lemma in the remainder of this subsection.

Lemma 3.9 *Let $\beta \geq 0$, $\delta > 0$ and $\ell \in \mathbb{N}$ be constants and assume we are given an instance of UFPP in which all tasks are both $(1 - 2\beta)$ -small and δ -large. Then, for each set $F^{k,\ell}$ there is a $(2, \beta)$ -approximation algorithm.*

Suppose we are given a set $F^{k,\ell}$ whose tasks are all δ -large and $(1 - 2\beta)$ -small.

Proposition 3.10 *Let $F \subseteq F^{k,\ell}$ be a feasible solution in which all tasks are δ -large. Then for any edge e , at most $2^{\ell+1}/\delta$ tasks in F use e .*

Proof: Let F be a feasible solution. For each task $i \in F^{k,\ell}$ it holds that $b(i) \geq 2^k$. In addition, all tasks are δ -large. Therefore, for all tasks $i \in F$ it holds that $d_i \geq \delta \cdot 2^k$. For every task $i \in F$, $b(i) < 2^{k+\ell}$. So according to Proposition 3.3, for every edge e it holds that $\sum_{i \in F \cap T_e} d_i < 2^{k+\ell+1}$. Therefore, at most $\frac{2^{k+\ell+1}}{\delta \cdot 2^k} = \frac{2^{\ell+1}}{\delta}$ tasks in F use e . \square

Due to Proposition 3.10, there is a straight forward dynamic programming algorithm for computing an optimal solution [14, 17]. (We remark that the running time is $n^{O(2^\ell/\delta)}$.)

Proposition 3.11 *For constant $\delta > 0$ and $\ell \in \mathbb{N}$, if all tasks are δ -large, an optimal solution for $F^{k,\ell}$ can be computed in polynomial time.*

Proof sketch: By Proposition 3.10, each edge can be used by at most $2^\ell/\delta$ tasks. Hence, for each edge e there can be at most $n^{O(2^\ell/\delta)}$ combinations of tasks which use e in an optimal solution. We enumerate all these combinations for each edge e . For each of these, we establish a dynamic programming cell which stores the maximum profit one can obtain from tasks i with $s_i < e$, given the respective combination of tasks that use e . The correct values for these cells can be computed by iterating through the edges of the path from left to right. See [14, 17] for details. \square

Lemma 3.12 *Let $\beta \geq 0$ and let F be a feasible solution for $F^{k,\ell}$ in which all tasks are $(1 - 2\beta)$ -small. Then in time $O(n^2)$, F can be partitioned into two sets H^1 and H^2 which are both feasible for the modified capacity constraints $u'_e := u_e - \beta \cdot 2^k$.*

Proof: Note that the claim is trivially true if $\beta = 0$ by setting $H^1 := F$ and $H^2 := \emptyset$, so now assume $\beta > 0$. We initialize two sets $H^1 := \emptyset$, $H^2 := \emptyset$. Assume that the tasks in F are ordered such that the start vertices are non-decreasing. We consider the tasks in F in this order. In the i -th iteration we take the task i . We add i to a set H^ℓ with $\ell \in \{1, 2\}$ such that $H^\ell \cup \{i\}$ obeys the modified capacity constraint, i.e., it leaves a free capacity of $\beta \cdot 2^k$ in each edge.

It remains to show that indeed either $H^1 \cup \{i\}$ or $H^2 \cup \{i\}$ obeys the modified capacity constraint. Assume to the contrary that neither $H^1 \cup \{i\}$ nor $H^2 \cup \{i\}$ obey the modified capacity constraint. Then there are edges e, e' on the path of i such that

$$\sum_{i' \in (H^1 \cup \{i\}) \cap T_e} d_{i'} > u_e - \beta \cdot 2^k \quad (3.1)$$

and

$$\sum_{i' \in (H^2 \cup \{i\}) \cap T_{e'}} d_{i'} > u_{e'} - \beta \cdot 2^k. \quad (3.2)$$

Inequality (3.1) implies that $\sum_{i' \in H^1 \cap T_e} d_{i'} > u_e - \beta \cdot 2^k - d_i$. Inequality (3.2) gives that $\sum_{i' \in H^2 \cap T_{e'}} d_{i'} > u_{e'} - \beta \cdot 2^k - d_i$. Assume w.l.o.g. that $e < e'$ or $e = e'$. Recall that we considered the tasks by non-decreasing start index. Hence, all tasks in $(H^2 \cup \{i\}) \cap T_{e'}$ use e as well. For the next calculation we need that

$$d_i \leq (1 - 2\beta)b(i) \leq u_{e'}(1 - 2\beta)$$

and hence $u_{e'} - d_i \geq 2\beta \cdot u_{e'}$. Also, note that $u_{e'} \geq 2^k$ since $i \in F \subseteq F^{k,\ell}$. We calculate that

$$\begin{aligned} u_e &\geq \left(\sum_{i' \in H^1 \cap T_e} d_{i'} \right) + \left(\sum_{i' \in H^2 \cap T_{e'}} d_{i'} \right) + d_i \\ &> \left(u_e - \beta \cdot 2^k - d_i \right) + \left(u_{e'} - \beta \cdot 2^k - d_i \right) + d_i \\ &= u_e + u_{e'} - d_i - 2\beta \cdot 2^k \\ &\geq u_e + 2\beta u_{e'} - 2\beta \cdot 2^k \\ &\geq u_e. \end{aligned}$$

This is a contradiction. Hence, task i can be added to one of the sets H^ℓ such that $H^\ell \cup \{i\}$ still obeys the capacity constraint. When computing H^1 and H^2 we need to check for each task in F whether adding it to one of the sets violates the modified capacity constraint in one of the edges. Since w.l.o.g. $m < 2n$, this check can be done in $O(n)$ time for each task. There are n tasks in total, and hence the entire procedure can be implemented in $O(n^2)$. \square

Now we can prove Lemma 3.9.

Proof of Lemma 3.9: Since the tasks are δ -large, we can compute an optimal solution $OPT(F^{k,\ell})$ in polynomial time (Proposition 3.11). Since tasks are $(1 - 2\beta)$ -small, in time $O(n^2)$, the solution $OPT(F^{k,\ell})$ can be partitioned into two solutions H^1 and H^2 that obey the modified capacity constraint (Lemma 3.12). Returning the solution of these two with maximum profit then yields a $(2, \beta)$ -approximation for $F^{k,\ell}$. \square

3.2.3 An Approximation Algorithm for $(1 - \gamma)$ -Small Tasks

By combining Lemmas 3.6 and 3.9, we obtain an (α, β) -approximation algorithm for small tasks.

Lemma 3.13 *Let $\epsilon > 0$, $\beta \geq 0$, and $\ell \in \mathbb{N}$ be constants with $\beta < 1$ and assume we are given an instance of UFPP in which all tasks are $(1 - 2\beta)$ -small. Then, for each set $F^{k,\ell}$ there is a $(2 + \frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm.*

Proof: Given ϵ , β and ℓ , Lemma 3.6 shows that there exists a $\delta > 0$ such that for the δ -small tasks we have an $(\frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm (for each set $F^{k,\ell}$). The remaining tasks are both δ -large and $(1 - 2\beta)$ -small. If $\delta \geq 1 - 2\beta$, we are done. Otherwise, Lemma 3.9 shows that we have an $(2, \beta)$ -approximation algorithm for these (for each set $F^{k,\ell}$). Together this gives an $(2 + \frac{1+\epsilon}{1-\beta}, \beta)$ -approximation algorithm for each $F^{k,\ell}$ (observe that Fact 2.1 also applies to (α, β) -approximation algorithms). \square

Using the above lemma and our framework (Lemma 3.2), we obtain our main result for small tasks.

Theorem 3.14 *For any pair of constants $\epsilon > 0$ and $\gamma > 0$, there is a polynomial time $(3 + \epsilon)$ -approximation algorithm for UFPP instances in which all tasks are $(1 - \gamma)$ -small.*

Proof: Choose $\ell \in \mathbb{N}$, $q \in \mathbb{N}$, $\beta := 2^{1-q}$ and $\epsilon' > 0$ such that $1 - 2\beta \geq 1 - \gamma$ (hence all tasks are $(1 - 2\beta)$ -small), and $\frac{\ell+q}{\ell} \cdot (2 + \frac{1+\epsilon'}{1-\beta}) \leq 3 + \epsilon$, which is always possible¹. Now, combining the framework using the chosen ℓ and q (Lemma 3.2) with $(2 + \frac{1+\epsilon'}{1-\beta}, \beta)$ -approximation algorithms for every $F^{k,\ell}$ (Lemma 3.13) yields a $(3 + \epsilon)$ -approximation algorithm. \square

3.3 Resource Augmentation

Using the techniques derived above, we now describe a polynomial time algorithm that computes a set of tasks $T' \subseteq T$ such that $w(T') \geq (2 + \epsilon)^{-1} \cdot w(OPT)$, and T' is feasible if the capacity of every edge is increased by a factor $1 + \epsilon$. Note that for this result we do not require that the tasks T are $(1 - \gamma)$ -small.

The main idea is the following: from the results in Sections 3.2.1 and 3.2.2 we will conclude that there are $(1 + \epsilon, 0)$ and $(1, 0)$ -approximation algorithms for the tiny tasks and all remaining tasks, respectively. Combining these, we obtain a $(2 + \epsilon, 0)$ -approximation algorithm for each set $F^{k,\ell}$ (without any further conditions on its tasks) in Proposition 3.15. We apply our framework with the sets computed by this algorithm. In Lemma 3.16 we show that the union $ALG(c) = \bigcup_{k \in \eta(c)} ALG(F^{k,\ell})$ is feasible when the capacities of the edges are increased by a factor $1 + 2^{2-q}$ (this lemma takes the role of Lemma 3.4 from the original framework).

The first step is to establish the approximation algorithms for the sets $ALG(F^{k,\ell})$.

Proposition 3.15 *Let $\epsilon > 0$. There is a $(2 + \epsilon, 0)$ -approximation algorithm for each set $F^{k,\ell}$ which runs in polynomial time.*

Proof: Using Lemma 3.6 with $\beta = 0$ yields that for all $\epsilon > 0$ there is a $\delta > 0$ such that there is a $(1 + \epsilon, 0)$ -approximation algorithm for each set $F^{k,\ell}$ which consists only of δ -small

¹ For instance, choose β small enough such that $1/(1 - \beta) \leq 1 + \epsilon'$, $1 - 2\beta \geq 1 - \gamma$, and that there is an integer $q \geq 2$ with $\beta = 2^{1-q}$. Now choose $\ell \in \mathbb{N}$ such that $\frac{\ell+q}{\ell} \leq 1 + \epsilon'$. We obtain an approximation factor of $(2 + \frac{1+\epsilon'}{1-\beta}) \cdot \frac{\ell+q}{\ell} \leq (2 + (1 + \epsilon')^2) \cdot (1 + \epsilon')$ which is at most $3 + \epsilon$ if ϵ' is sufficiently small.

tasks. Proposition 3.11 implies that for each fixed $\delta > 0$ there is a polynomial time $(1, 0)$ -approximation algorithm (i.e., an optimal algorithm in the usual sense) for sets $F^{k,\ell}$ which consist only of δ -large tasks. Using Fact 2.1 this yields a $(2 + \epsilon, 0)$ -approximation algorithm for each set $F^{k,\ell}$. \square

The next lemma is an adjusted version of Lemma 3.4. In contrast to the latter, here the increased capacity allows us to combine the computed sets $ALG(F^{k,\ell})$ to a globally feasible solution, without requiring that every solution should leave a fraction of the capacity free. Recall that $\eta(c) = \{c + i \cdot (\ell + q) \mid i \in \mathbb{Z}\}$.

Lemma 3.16 *Let $q \in \mathbb{N}$, $\ell \in \mathbb{N}$, and $c \in \{0, \dots, \ell + q - 1\}$. For each k , let $ALG^{k,\ell} \subseteq F^{k,\ell}$ be a feasible solution and define $ALG(c) = \bigcup_{k \in \eta(c)} ALG(F^{k,\ell})$. Then for each edge e it holds that*

$$\sum_{i \in ALG(c) \cap T_e} d_i \leq u_e \cdot (1 + 2^{2-q}).$$

Proof: Let \bar{k} be the largest integer in $\eta(c)$ such that $2^{\bar{k}} \leq u_e$. Let $P := \ell + q$. From Proposition 3.3 we conclude that $\sum_{i \in ALG^{k,\ell} \cap T_e} d_i \leq 2 \cdot 2^{k+\ell}$ for each $k < \bar{k}$. This implies that

$$\begin{aligned} \sum_{i \in ALG(c) \cap T_e} d_i &= \sum_{i \in ALG^{\bar{k},\ell} \cap T_e} d_i + \sum_{i=1}^{\infty} \sum_{j \in ALG^{\bar{k}-i,P,\ell} \cap T_e} d_j \\ &\leq u_e + \sum_{i=1}^{\infty} 2 \cdot 2^{\bar{k}-i \cdot P + \ell} \\ &= u_e + 2^{1+\bar{k}+\ell} \sum_{i=1}^{\infty} (2^P)^{-i} \\ &\leq u_e + 2^{1+\ell} \cdot u_e \cdot \frac{1}{2^P - 1} \\ &= u_e \left(1 + \frac{2^{1+\ell}}{2^{\ell+q} - 1} \right) \\ &\leq u_e \cdot (1 + 2^{2-q}). \end{aligned}$$

\square

The following lemma establishes our framework (see Lemma 3.2) in the setting of resource augmentation.

Lemma 3.17 *Let \mathcal{I} be a UFPP instance. Let $\ell \in \mathbb{N}$ and $q \in \mathbb{N}$ be constants. If polynomial time $(\alpha, 0)$ -approximation algorithms exist for each set $F^{k,\ell}$, then there is a polynomial time algorithm that computes a set of tasks whose profit is at least $\alpha^{-1} \cdot \frac{\ell}{\ell+q} \cdot w(OPT(\mathcal{I}))$, and which is feasible if the capacity of every edge e is modified to $u_e(1 + 2^{2-q})$.*

Proof: The algorithm works as follows. For each $c \in \{0, \dots, \ell + q - 1\}$ we define $ALG(c) := \bigcup_{k \in \eta(c)} ALG(F^{k,\ell})$. We output the set $ALG(c^*)$ with maximum profit among all sets $ALG(c)$. Lemma 3.5 implies that $w(ALG(c^*)) \geq \frac{\ell}{\ell+q} \cdot \frac{1}{\alpha} \cdot w(OPT(\mathcal{I}))$. Lemma 3.16 yields that $ALG(c^*)$ is feasible if the capacity of each edge e is changed to $u_e(1 + 2^{2-q})$. \square

The previous lemmas can be combined to obtain the main result of this section:

Theorem 3.18 *Let $\epsilon > 0$, $\beta > 0$. Given an instance \mathcal{I} of the UFPP problem. There is a polynomial time algorithm that computes a set of tasks whose profit is at least $(2 + \epsilon)^{-1} \cdot w(\text{OPT}(\mathcal{I}))$, and which is feasible if the capacity of every edge e is modified to $u_e(1 + \beta)$.*

Proof: We choose $\ell \in \mathbb{N}$, $q \in \mathbb{N}$, and $\epsilon' > 0$ such that $2^{2-q} \leq \beta$ and $\frac{\ell+q}{\ell} \cdot (2 + \epsilon') \leq 2 + \epsilon$, which is always possible (similar as in the proof of Theorem 3.14). The claim follows by combining the framework of Lemma 3.17 with the $(2 + \epsilon', 0)$ -approximation algorithms from Proposition 3.15. \square

4 Large Tasks

In this section we provide a polynomial time $2k$ -approximation algorithm for instances consisting of only $1/k$ -large tasks. In our main algorithm (in Section 5), this is used as a 4-approximation algorithm for the set of $1/2$ -large tasks. The main idea is to restrict to UFPP solutions of a certain form: we will represent tasks by rectangles drawn in the plane, and compute an independent set of rectangles of maximum weight (profit).

By (x_1, y_1, x_2, y_2) we will denote the axis-parallel rectangle in the plane \mathbb{R}^2 with upper left point (x_1, y_1) and lower right point (x_2, y_2) . We will call two rectangles (x_1, y_1, x_2, y_2) and (x'_1, y'_1, x'_2, y'_2) *compatible* if they do not overlap (i.e. do not share an internal point). More precisely, they are compatible if at least one of the following holds: $x_2 \leq x'_1$, $x'_2 \leq x_1$, $y_1 \leq y'_2$, or $y'_1 \leq y_2$.

Definition 4.1 (associated rectangle) *With a task i we associate the rectangle $(s_i, b(i), t_i, \ell(i))$.*

Note that s_i and t_i are integers since the path vertices are labeled $0, \dots, m$, and that $b(i)$ and $\ell(i) := b(i) - d_i$ are non-negative integers as well. Tasks are called *compatible* if their associated rectangles are compatible, and a task set F is called an *independent task set (ITS)* if all tasks are pairwise compatible.




Figure 4.1: A UFPP instance consisting of a path of length eight (with capacities 2,4,4,6,6,5,3,3) and four tasks, represented using rectangles drawn as high as possible under the capacity profile.

The following geometric interpretation motivates this choice of rectangles, see Figure 4.1. For every edge $e = \{x, x + 1\} \in E(P)$, we draw a horizontal line segment between (x, u_e) and $(x + 1, u_e)$. We add vertical line segments to complete this into a closed curve from $(0, 0)$ to $(m, 0)$, called the *capacity profile* (represented as a bold grey curve in Figure 4.1 and later figures). For every task i , the associated rectangle is now a rectangle of height d_i , drawn as high as possible under the capacity profile. The tasks in Figure 4.1 have profits 3, 2, 2, and 2, respectively. Therefore, an optimal UFPP solution consists of tasks 1, 3, and 4, with a total

profit of 7. However, since the rectangles for the tasks 1 and 4 are incompatible, this is not an ITS. The optimal ITS consists of tasks 2, 3 and 4, with total profit 6. This ITS is a UFPP solution as well. This is in fact always true, as shown in the next proposition.

Proposition 4.2 *Let $F \subseteq T$ be an ITS. Then F is a feasible UFPP solution.*

Proof: Consider an edge $e \in E(P)$, and all tasks T_e that use e . Consider two tasks $i, j \in F \cap T_e$. They are compatible, but $s_i < e < t_j$ and $s_j < e < t_i$, so either $\ell(i) \geq b(j)$ or $\ell(j) \geq b(i)$ must hold. It follows that the tasks in $F \cap T_e$ can be numbered i_1, \dots, i_p such that $b(i_1) \leq \ell(i_2)$, $b(i_2) \leq \ell(i_3)$, etc. Hence,

$$\begin{aligned} \sum_{k=1}^p d_{i_k} &= \sum_{k=1}^p (b(i_k) - \ell(i_k)) \\ &= b(i_p) - \ell(i_1) + \sum_{k=1}^{p-1} \underbrace{(b(i_k) - \ell(i_{k+1}))}_{\leq 0} \\ &\leq u_e. \end{aligned}$$

□

In Section 4.1, we will first show that if all tasks are $\frac{1}{k}$ -large, then the profit of an optimal ITS is at most a factor $2k$ worse than the profit of an optimal UFPP solution. We will prove this by showing that any UFPP solution can be partitioned into $2k$ ITSs.

After that we will give a dynamic programming algorithm for finding a maximum profit ITS. The core concept is that of a *corner*, which corresponds to a region of the plane that contains a subset of the rectangles. In Section 4.2 corners are defined, and a recursion formula is given for computing the profit of an optimal ITS that fits in such a corner. The dynamic programming algorithm that is based on this recursion is subsequently given in Section 4.3.

4.1 Partitioning UFPP Solutions into ITSs

In this section we prove that if all tasks are $1/k$ -large, the value of an optimal ITS is at most a factor $2k$ worse than the value of an optimal UFPP solution, by showing that any UFPP solution can be partitioned into $2k$ ITSs. This result is based on the following property.

Proposition 4.3 *Let F be a feasible UFPP solution, and let $k \geq 2$ be an integer, such that every task in F is $1/k$ -large. Let e be a bottleneck edge for $i \in F$. Then there are at most $k - 1$ tasks $j \in F \setminus \{i\}$ that are incompatible with i and use e .*

Proof: Let the set $F' \subseteq F \setminus \{i\}$ consist of all tasks that are incompatible with i and that use e . Suppose $|F'| \geq k$. Consider $j \in F'$. Since j and i are incompatible, but both use e and e is a bottleneck edge for i , it follows that $b(j) > \ell(i) = b(i) - d_i$. Therefore, $d_i + \sum_{j \in F'} d_j \geq d_i + \frac{1}{k} \sum_{j \in F'} b(j) > d_i + \frac{1}{k} |F'| (b(i) - d_i) \geq b(i) = u_e$, contradicting that F is a feasible solution. □

A partition $\{F_1, \dots, F_\ell\}$ of a task set F into ℓ ITSs will be encoded by an ℓ -coloring α of F . This is a function $\alpha : F \rightarrow \{1, \dots, \ell\}$ such that if $\alpha(i) = \alpha(j)$, then i and j are compatible.

(So $\alpha(i) = j$ means that $i \in F_j$.) An edge e is called a *separator* for a task set F if it is a bottleneck edge for some task $i \in F$, such that all tasks in F that use e are incompatible with i . So by Proposition 4.3, if e is a separator edge, then there are at most $k - 1$ tasks in F that use e and are incompatible with i , when F is a UFPP solution consisting of $1/k$ -large tasks. A coloring α of F is called *nice* if for every edge e and task $i \in F$ that has e as its bottleneck edge, all tasks that use e and are incompatible with i are colored differently. The main idea behind these notions and our construction of a coloring is as follows: We will identify a separator edge e , and consider the set F_0 of tasks i with $s_i < e$, and F_1 of tasks i with $t_i > e$. (Note that $F_0 \cup F_1 = F$ and $F_0 \cap F_1 \neq \emptyset$.) Unless $F_0 = F$ or $F_1 = F$, we may use induction to conclude that both admit a nice $2k$ -coloring. Then, since e is a separator edge and the colorings are nice, all tasks in $F_0 \cap F_1$ are colored differently in both colorings. Therefore these can be combined into a single nice $2k$ -coloring for F . However, since it may occur that $F_0 = F$ or $F_1 = F$, we need a slightly more sophisticated argument, which we will now present in detail.

Lemma 4.4 *Let F be a feasible UFPP solution, and let $k \geq 2$ be an integer, such that every task in F is $1/k$ -large. Then there exists a nice $2k$ -coloring for F .*

Proof: We prove the lemma by induction over $|F|$. The statement is trivially true when $|F| \leq 2k$. Now suppose that $|F| > 2k$, and assume that the above statement holds for every such set F' with $|F'| < |F|$. The proof is illustrated in Figure 4.2. Let e_B be a bottleneck edge for some task in F , with minimum capacity among all such edges. Let $L \subseteq F$ be the set of tasks that use e_B . By Proposition 4.3 and the choice of e_B , $|L| \leq k$.

Let $E_S = \{e_1, \dots, e_p\}$ be the set of edges e for which there is a task i with e as bottleneck edge, such that $i \in L$ or i is incompatible with some $j \in L$. (So $e_B \in E_S$.) The edges in E_S are numbered such that if $x < y$, then $e_x < e_y$.

We first argue that all edges in E_S are separators for F . For $e_B \in E_S$, the statement is clear since e_B is a bottleneck edge with minimum capacity. Now consider an edge $e \in E_S \setminus \{e_B\}$, and let $i \in F$ be a task with e as bottleneck, that is incompatible with some task $i' \in L$. Then $\ell(i) < b(i')$ holds. Now suppose there is a task i'' using e that is compatible with i . Since both i and i'' use e and e is a bottleneck edge of i , from their compatibility it follows that $b(i'') \leq \ell(i) < b(i') = e_B$. But this contradicts that e_B is a bottleneck edge for F with minimum capacity.

Now, for $1 \leq j \leq p - 1$, define $F_j \subseteq F$ to be all tasks i with $s_i < e_{j+1}$ and $t_i > e_j$. Similarly, define $F_0 \subseteq F$ to be all tasks i with $s_i < e_1$, and define $F_p \subseteq F$ to be all tasks i with $t_i > e_p$. Observe that $F_0 \cup \dots \cup F_p = F$.

Case 1: For every j , $|F_j| < |F|$.

In this case, we may use induction to conclude that for every F_j there exists a nice $2k$ -coloring α_j . We can combine these into a nice $2k$ -coloring of F : all tasks in $C = F_0 \cap F_1$ use the edge e_1 . Since e_1 is a separator, there exists a task $i \in C$ with e_1 as bottleneck such that all tasks using e_1 are incompatible with i . So, in both the nice $2k$ -coloring α_0 of F_0 and the nice $2k$ -coloring α_1 of F_1 , all tasks in C are colored differently. Therefore, we can permute the colors of α_1 such that the tasks in C receive the same colors in α_0 and α_1 . At this point, the colorings can be combined into a $2k$ -coloring α' of $F_0 \cup F_1$, which is again nice.

Next, we can combine the nice $2k$ -coloring α_2 of F_2 with the nice $2k$ -coloring α' of $F_0 \cup F_1$ in a similar way, and continue like this with F_3, \dots, F_p , until a nice $2k$ -coloring of $F = F_0 \cup \dots \cup F_p$ is obtained. This proves the desired statement in this case.

Case 2: There exists a j such that $F_j = F$.

For such a choice of j , define $F'_j = F_j \setminus L$. Since L is non-empty, $|F'_j| < |F|$ holds, so we may again use induction to conclude that F'_j admits a nice $2k$ -coloring α . We will extend this to a nice $2k$ -coloring of $F_j = F$. For ease of exposition we will assume that $1 \leq j < p$ (the cases $j = 0$ and $j = p$ are similar but easier). So both e_j and e_{j+1} are edges in E_S . We assume also w.l.o.g. that $e_B \geq e_{j+1}$.

We observe that every task $i \in F'_j$ that is incompatible with some task in L uses either e_j or e_{j+1} : by definition of E_S , such a task i must use some edge $e_x \in E_S$. If $x > j + 1$ then by definition of F_j , i also uses e_{j+1} . If $x < j$ then similarly i also uses e_j .

Since $e_B \geq e_{j+1}$ and $F_j = F$, all tasks in L use e_{j+1} . Then, since e_{j+1} is a separator, there are at most $k - |L|$ tasks in F'_j that use e_{j+1} (Proposition 4.3); let F_R denote these tasks. In addition, there are at most k tasks in F'_j that use e_j ; denote these by F_L . Since $|F_L| + |F_R| \leq 2k - |L|$ and there are $2k$ available colors, we can choose $|L|$ different colors that are not used for tasks in $F_L \cup F_R$, in the coloring α . We extend the nice $2k$ -coloring α for F'_j to $F_j = F$ by using these $|L|$ colors for the tasks in L . Since we observed that all tasks in F'_j that are incompatible with some task in L are included in $F_L \cup F_R$, it follows that this yields again a nice $2k$ -coloring for F . Since we now constructed a nice $2k$ -coloring for F in all cases, this concludes the proof of Lemma 4.4. \square




Figure 4.2: An illustration of the proof of Lemma 4.4 for the case $k = 2$ (using logarithmic scale on the y -axis, all tasks are $\frac{1}{2}$ -large). The gray task is the single task in L , and the bold tasks are associated with the indicated separator edges e_1, \dots, e_5 .

We will now show that this bound is tight for every $k \geq 2$: consider a path P of length 5, with capacities $2k^2, 2k^2 + 2k, 2(2k^2 + 2k), 2k^2 + 2k$ and $2k^2$, in order along the path. (So $V(P) = \{0, \dots, 5\}$.) Introduce $k - 1$ tasks with $s_i = 0, t_i = 3$ and $d_i = 2k + 1$, $k - 1$ tasks with $s_i = 2, t_i = 5$ and $d_i = 2k + 1$, one task with $s_i = 1, t_i = 3$ and $d_i = 2k + 3$, and one task with $s_i = 2, t_i = 4$ and $d_i = 2k + 3$. All tasks can be verified to be $\frac{1}{k}$ -large. In addition, they all satisfy $b(i) \geq 2k^2 > \ell(i)$ and $s_i \leq 2 < 3 \leq t_i$, so all are pairwise incompatible, and therefore at least $2k$ colors are needed. Finally, they constitute a feasible UFPP solution: for the first and fifth edge this is easy to see. For the second and fourth edge, the sum of demands using the edge is $(k - 1)(2k + 1) + (2k + 3) = 2k^2 + k + 2 \leq 2k^2 + 2k = u_{\{1,2\}} = u_{\{3,4\}}$, since $k \geq 2$. All tasks use the third edge, so the demand sum is $2 \cdot ((k - 1)(2k + 1) + (2k + 3)) = 2 \cdot (2k^2 + k + 2) = 4k^2 + 2k + 4 \leq 4k^2 + 4k = u_{\{2,3\}}$, since $k \geq 2$. The above is summarized in the following proposition.

Proposition 4.5 *For any $k \in \mathbb{N}$ with $k \geq 2$ there is an instance of UFPP with the properties that*

- *the instance consists of $2k$ tasks which are all $\frac{1}{k}$ -large,*
- *the set of all tasks of the instance yields a feasible UFPP solution, and*
- *the optimal ITS contains only one task.*

4.2 Computing Optimal ITSs in Corners

To bound the complexity of our algorithm, it is useful to assume that all edge capacities are different. In the next lemma we first show that by making small perturbations to the capacities, this property can easily be guaranteed, without changing the set of feasible solutions.

Lemma 4.6 *Let \mathcal{I} be a UFPP instance with task set $T = \{1, \dots, n\}$. In linear time, the capacities and demands can be modified such that for the resulting instance \mathcal{I}' :*

- *All edge capacities are distinct, and*
- *a set of tasks $F \subseteq T$ is feasible for \mathcal{I} if and only if F is feasible for \mathcal{I}' .*

Proof: Recall that the m path edges are labeled $\{i, i + 1\}$, for $i \in \{0, \dots, m - 1\}$. For every edge $e = \{i, i + 1\}$, we change the capacity to $u'_e = mu_e + i$. For every task $j \in T$, we change the demand to $d'_j = md_j$. This gives the instance \mathcal{I}' , in which all capacities are distinct.

Suppose $F \subseteq T$ is feasible for \mathcal{I}' . Then for every edge $e = \{i, i + 1\}$, it holds that $\sum_{i \in T_e \cap F} d_i = \frac{1}{m} \sum_{i \in T_e \cap F} d'_i \leq \lfloor \frac{1}{m} u'_e \rfloor = \lfloor u_e + \frac{i}{m} \rfloor = u_e$. Hence F is also feasible for the original instance \mathcal{I} . Clearly, every task set F that is feasible for \mathcal{I} remains feasible for \mathcal{I}' . \square

The central concept of our dynamic program is that of a *corner* (x, y, z) . First we give an informal, geometric explanation of this notion, using the representation of the problem explained before. Subsequently, we will give a formal definition.

A corner (x, y, z) corresponds to a certain region under the capacity profile, see Figure 4.3. In our dynamic program, we will compute for every such corner the profit of an optimal top-drawn solution that fits entirely within this region. This will be done using previously computed values for ‘smaller’ corners. A corner (x', y', z') is *smaller* than the corner (x, y, z) if its region is a strict subset of the region associated with (x, y, z) (we will make this more precise in the proof of Lemma 4.16).

A corner is determined by an integer x -coordinate $0 \leq x \leq m$ (a path vertex), and two integer y -coordinates $y \geq 0$ and $z \geq 0$. If the capacity of the edge $\{x - 1, x\} \in E(P)$ is more than y , then draw a horizontal line segment from the point (x, y) to the left, to the first point that lies on the capacity profile curve. This point will be called (w_L, y) . Otherwise, if the capacity is at most y , then let $w_L = x$. Similarly, if the capacity of $\{x, x + 1\} \in E(P)$ is more than z , then draw a horizontal line segment from (x, z) to the right, to the first point that lies on the capacity profile. This point will be called (w_R, z) . Otherwise, let $w_R = x$. Connect these line segments into a single curve from (w_L, y) to (w_R, z) by adding a vertical line segment from (x, y) to (x, z) . This curve and the capacity profile together now enclose a bounded region, which is shown as a shaded region in Figure 4.3. (In the special case that




Figure 4.3: The corner $C(x, y, z)$ contains all tasks for which the associated rectangles are fully contained in the shaded region.

$u_{\{x, x+1\}} > z \geq u_{\{x-1, x\}} > y$ or $u_{\{x-1, x\}} > y \geq u_{\{x, x+1\}} > z$, the corner actually corresponds to two disjoint regions.) This is the region shown that we associate with the corner (x, y, z) ; we say that a task *fits into* the corner (x, y, z) when its rectangle, drawn as explained earlier, lies fully in (the closure of) this region. Now we give formal definitions.

Definition 4.7 A triple (x, y, z) of integers is called a corner if $0 \leq x \leq m$, $y \geq 0$ and $z \geq 0$. For a corner (x, y, z) , we denote by $w_L(x, y)$ or simply w_L the lowest numbered path vertex such that for all i with $w_L \leq i < x$, it holds that $u_{\{i, i+1\}} > y$. Similarly, $w_R(x, z)$ or simply w_R is defined to be the largest numbered path vertex such that for all i with $w_R \geq i > x$, it holds that $u_{\{i-1, i\}} > z$. Hence, $w_L \leq x$ and $w_R \geq x$.

For each corner (x, y, z) we define a set $C(x, y, z)$ which—intuitively speaking—consists of all tasks for which the associated rectangles are contained in the region associated with this corner.

Definition 4.8 For a corner (x, y, z) , we denote by $C(x, y, z)$ the set of all tasks $i \in T$ for which at least one of the following holds:

- $w_L(x, y) \leq s_i$, $t_i \leq w_R(x, z)$ and $\ell(i) \geq \max\{y, z\}$,
- $w_L(x, y) \leq s_i$, $t_i \leq x$ and $\ell(i) \geq y$, or
- $x \leq s_i$, $t_i \leq w_R(x, z)$ and $\ell(i) \geq z$.

Due to this definition, we say that *task i fits into the corner (x, y, z)* or *corner (x, y, z) contains i* if $i \in C(x, y, z)$. Hence, $C(x, y, z)$ is the set of all tasks that fit into the corner. Note that they are possibly incompatible.

For a given UFPP instance and corner (x, y, z) , we denote by $P(x, y, z)$ the maximum value of $w(F)$ over all ITSs F with $F \subseteq C(x, y, z)$. An ITS F with $F \subseteq C(x, y, z)$ and $w(F) = P(x, y, z)$ is said to *determine* $P(x, y, z)$. Let $u_{\max} := \max_{e \in E(P)} u_e$. Observe that all tasks fit into the corner $(m, 0, u_{\max})$, so computing $P(m, 0, u_{\max})$ yields the desired value:

Proposition 4.9 $P(m, 0, u_{\max})$ equals the profit of an optimal ITS, where m is the path length.

We will now show how $P(x, y, z)$ can be computed in various cases. Proposition 4.10 collects the easy cases (which are easily understood using the above geometric interpretation), and Lemma 4.12 covers the complex case which gives the main recursion formula to compute an optimal ITS. In Proposition 4.10, statement (i) and (ii) show how to rewrite a corner in ‘standard form’ without changing the corresponding region. Statement (iii) deals with the case where the corner corresponds to two disconnected regions.

Proposition 4.10

- (i) If $y = z$, then $C(x, y, z) = C(w_R(x, z), y, u_{\max})$, and hence $P(x, y, z) = P(w_R(x, z), y, u_{\max})$.
- (ii) If $x = 0$ or $y \geq u_{\{x-1, x\}}$, then $C(x, y, z) = C(x, u_{\max}, z)$, and hence $P(x, y, z) = P(x, u_{\max}, z)$.
- (iii) If $y < z$, $x \geq 1$, $y < u_{\{x-1, x\}}$ and $u_{\{x-1, x\}} \leq z < u_{\{x, x+1\}}$, then $C(x, y, z) = C(x, y, u_{\max}) \cup C(x, u_{\max}, z)$ and hence $P(x, y, z) = P(x, y, u_{\max}) + P(x, u_{\max}, z)$.

Proof: The statements (i) and (ii) follow immediately from the definitions. We will now prove statement (iii). Consider an ITS F that determines $P(x, y, z)$. Since $u_{\{x-1, x\}} \leq z$, every task $i \in F$ with $s_i \leq x - 1$ and $t_i \geq x$ has $b(i) \leq z$. Hence F contains no tasks i with $s_i \leq x - 1$ and $t_i \geq x + 1$, and thus can be partitioned into two ITSs that fit into the corners (x, y, u_{\max}) and (x, u_{\max}, z) respectively. It follows that $P(x, y, z) \leq P(x, y, u_{\max}) + P(x, u_{\max}, z)$.

Now let F_1 and F_2 be ITSs that determine $P(x, y, u_{\max})$ and $P(x, u_{\max}, z)$. These are disjoint and both fit in the corner (x, y, z) , so $P(x, y, z) \geq w(F_1) + w(F_2) = P(x, y, u_{\max}) + P(x, u_{\max}, z)$. \square

Because of Proposition 4.10 (i), we only need to consider corners (x, y, z) where $y < z$ or $y > z$ holds. Since these cases are symmetric, *we will only consider corners (x, y, z) for which $y < z$ holds, in the following definitions and lemmas.* One can easily deduce the corresponding statements for the case $y > z$. For the sake of brevity and readability, we will however not explicitly treat this case. By Proposition 4.10 (ii), we may furthermore assume that $x \geq 1$ and $y < u_{\{x-1, x\}}$. Informally, Proposition 4.10 (iii) shows that we only need to consider corners for which the associated region is connected. These observations show that we may restrict our attention to special types of corners, which we call *proper*.

Definition 4.11 *A corner (x, y, z) is called proper if $y < z$, $x \geq 1$, $y < u_{\{x-1, x\}}$ and either $z < u_{\{x-1, x\}}$ or $z \geq u_{\{x, x+1\}}$ holds.*

Now we will consider the more complex case where the given corner (x, y, z) is proper. The main idea is as follows. Consider an ITS F that determines $P(x, y, z)$. Either F also fits into the smaller corner $(x - 1, y, z)$, or there exists a task j with $\ell(j) < z$ and $t_j = x$. In the latter case, we show that F can be partitioned into two task sets that fit into smaller corners, and one single task i . Given a task i with $t_i \leq x$, we will consider the two smaller corners $(s_i, y, b(i))$ and $(x, b(i), z)$. These are illustrated in Figure 4.4. For the indicated task i and the depicted top-drawn set F , we have the desired property that $F \setminus \{i\} \subset C(s_i, y, b(i)) \cup C(x, b(i), z)$. We will show that there always exists a task i for which this holds; such a task is called *special*. This explains the following recursive formula for proper corners.




Figure 4.4: On the left, a proper corner (x, y, z) and an ITS $F \subseteq C(x, y, z)$ is shown. When choosing a special task $i \in F$, it holds that $F \setminus \{i\} \subseteq C(s_i, y, b(i)) \cup C(x, b(i), z)$.

Lemma 4.12 Consider a UFPP instance in which all edge capacities are distinct. Let (x, y, z) be a proper corner. Then

$$P(x, y, z) = \max \left\{ P(x-1, y, z), \max_{i \in C(x, y, z), t_i \leq x} \left\{ w_i + P(s_i, y, b(i)) + P(x, b(i), z) \right\} \right\}.$$

Before we can prove Lemma 4.12, we will define which properties are required for a special task, and prove that such a task always exists. The essential property of a special task i is that the rectangle $(s_i, b(i), x, y)$ is compatible with any associated rectangle of a task in $F \setminus \{i\}$, and does not overlap with the capacity profile. In this case we will say that there are no tasks in F that lie to the right of i or lie below i , which is defined as follows.

Definition 4.13 Consider a proper corner (x, y, z) , and a task $i \in C(x, y, z)$ with $t_i \leq x$.

- A task $j \in C(x, y, z)$ lies to the right of i if $t_i \leq s_j < x$ and $\ell(j) < b(i)$.
- A task $j \in C(x, y, z)$ lies below i if $b(j) \leq \ell(i)$ and $t_j > s_i$.

Definition 4.14 Let (x, y, z) be a proper corner and $F \subseteq C(x, y, z)$ be an ITS. A task $i \in F$ with $t_i \leq x$ is called special with respect to F and (x, y, z) if

- for all edges e with $s_i < e < x$, it holds that $u_e \geq b(i)$, and
- there is no task $j \in F \setminus \{i\}$ that lies to the right of i or that lies below i .

We now sketch how to find a special task i in case that $F \not\subseteq C(x-1, y, z)$. We start with a task i such that $t_i = x$ and $\ell(i) < z$ holds. Now, assume there is a task j whose associated rectangle is incompatible with the rectangle $(s_i, b(i), x, y)$. Then, j lies below i (that is, $b(j) \leq \ell(i)$), and no associated rectangle of a task crosses the line segment from $(s_j, b(j))$ to $(x, b(j))$. This follows essentially from the fact that all tasks are compatible in F , and all associated rectangles touch the capacity profile. We continue the same procedure with j in the role of i until we find a special task i . We now prove this formally.

Lemma 4.15 Let (x, y, z) be a proper corner and $F \subseteq C(x, y, z)$ be an ITS. Then $F \subseteq C(x-1, y, z)$ or there exists a special task $i \in F$ with respect to F and (x, y, z) .

Proof: A task $i \in C(x, y, z)$ is called a *candidate* for (x, y, z) if

- $t_i \leq x$,
- $\ell(i) < z$ and
- for all edges $e \in E(P)$ with $s_i < e < x$, $u_e \geq b(i)$ holds.

From $u_{\{x-1, x\}} > y$ it follows that $w_L(x, y) = w_L(x-1, y)$. Obviously, $w_R(x-1, z) \leq w_R(x, z)$. Therefore, since $y < z$ it holds that $C(x-1, y, z) \subseteq C(x, y, z)$. Furthermore, $w_R(x-1, z) < w_R(x, z)$ can only occur if $u_{\{x-1, x\}} \leq z$. But since the corner is proper, this implies that $u_{\{x, x+1\}} \leq z$, so $w_R(x, z) = x$. These observations show that the only case when $F \subseteq C(x-1, y, z)$ does not hold is when there is a task $j \in F$ with $t_j = x$ and $\ell(j) < z$. Hence it is easily seen that j is a candidate task (for (x, y, z)). In addition, since $t_j = x$, no task in F lies to the right of j . So to prove the lemma statement, we may assume that there exists at least one candidate task in F with no tasks to the right of it. Choose $i \in F$ to be such a task with minimum value for $b(i)$. We prove that no task in F lies below i , which will prove the lemma. Suppose to the contrary that a task $j \in F$ lies below i , so $b(j) \leq \ell(i) < b(i)$ and $t_j > s_i$. It is easily checked that j is a candidate as well. By choice of i , it must then hold that there exists a task $k \in F$ that lies to the right of j (otherwise j should have been chosen in the role of i). So $t_j \leq s_k < x$, and $\ell(k) < b(j) \leq \ell(i) < b(i)$. Therefore, $s_k < t_i$; otherwise k would also lie to the right of i . But now we can use the fact that i and k are compatible: Recall that we have that $s_k < t_i$, $s_i < t_j \leq s_k$, and $\ell(k) < b(j) < b(i)$. So the only case in which i and k can be compatible is when $b(k) \leq \ell(i) < b(i)$. This however contradicts the fact that all edges between s_i and x have capacity at least $b(i)$; the edge that determines $b(k)$ lies between s_i and x . We conclude that i is a candidate task that satisfies the additional properties that there are no tasks in F below it, or to the right of it, which therefore is special. \square

Now we can prove Lemma 4.12.

Proof of Lemma 4.12: Let (x, y, z) be a proper corner in a UFPP instance in which all capacities are distinct. We will first show that

$$P(x, y, z) \leq \max \left\{ P(x-1, y, z), \max_{i \in C(x, y, z), t_i \leq x} \left\{ w_i + P(s_i, y, b(i)) + P(x, b(i), z) \right\} \right\}. \quad (4.1)$$

Let F be an ITS that determines $P(x, y, z)$. If $F \subseteq C(x-1, y, z)$ then $P(x, y, z) \leq P(x-1, y, z)$. If not, then by Lemma 4.15, there exists a special task i . We argue that $F \setminus \{i\}$ can be partitioned into two ITSs F_1 and F_2 with $F_1 \subseteq C(s_i, y, b(i))$ and $F_2 \subseteq C(x, b(i), z)$, which will prove that $P(x, y, z) \leq \left\{ w_i + P(s_i, y, b(i)) + P(x, b(i), z) \right\}$, and therefore the Inequality (4.1). First we consider the w_L and w_R vertices for both corners. Let $e_B = \{v, v+1\}$ be a bottleneck edge for i , which is unique since all capacities are distinct. So $u_{e_B} = b(i)$ and $s_i < e_B < t_i$. For the corner $(s_i, y, b(i))$, it is easy to see that $w_L(s_i, y) = w_L(x, y)$ and $w_R(s_i, b(i)) = v$. Now consider the corner $(x, b(i), z)$. Obviously the w_R value is the same as for $C(x, y, z)$. In addition, since i is special, all edges e with $s_i < e_B < e < x$ have $u_e \geq b(i)$. Therefore, since we assumed all edge capacities are distinct, $u_e > b(i)$ holds for all such edges e . It follows that $w_L(x, b(i)) = v + 1$.

Now consider a task $j \in F \setminus \{i\}$. We distinguish five cases.

1. $t_j \leq s_i$: $w_L(x, y) = w_L(s_i, y)$ and $\ell(j) \geq y$ imply that $j \in C(s_i, y, b(i))$.
2. $s_i < t_j < e_B$: Since j is compatible with i but does not lie below i , $\ell(j) \geq b(i)$ holds. Therefore $j \in C(s_i, y, b(i))$.
3. $s_j \geq x$: From $j \in C(x, y, z)$ it easily follows that $j \in C(x, b(i), z)$.
4. $e_B < s_j < x$: Since j is compatible with i but does not lie to the right of i , $\ell(j) \geq b(i)$ must hold. In addition, if $t_j > x$ then $\ell(j) \geq z$. Therefore $j \in C(x, b(i), z)$.
5. $s_j < e_B$ and $e_B < t_j$: We argue that this is not possible. We have that $b(j) \leq u_{e_B} = b(i)$. Because i and j are compatible, it then must be that $b(j) \leq \ell(i)$. But this contradicts that there is no task below i .

Since we considered all possibilities, this concludes the proof of Inequality (4.1).

Next, we argue that

$$P(x, y, z) \geq \max \left\{ P(x-1, y, z), \max_{i \in C(x, y, z), t_i \leq x} \left\{ w_i + P(s_i, y, b(i)) + P(x, b(i), z) \right\} \right\}. \quad (4.2)$$

Recall that since (x, y, z) is a proper corner, $C(x-1, y, z) \subseteq C(x, y, z)$, so $P(x, y, z) \geq P(x-1, y, z)$. Now consider a task $i \in C(x, y, z)$ with $t_i \leq x$. Let F_1 and F_2 be the ITSs that determine $P(s_i, y, b(i))$ and $P(x, b(i), z)$, respectively. We will argue that $F_1 \cap F_2 = \emptyset$ and $F_1 \cup F_2 \subseteq C(x, y, z)$. Since clearly $i \notin F_1$ and $i \notin F_2$, it will follow that $P(x, y, z) \geq w_i + P(s_i, y, b(i)) + P(x, b(i), z)$, proving Inequality (4.2).

Now let e_B a bottleneck edge for i . Then $w_R(s_i, b(i)) < e_B$. On the other hand, since $t_i \leq x$, it holds that $w_L(x, b(i)) > e_B$. It follows that $F_1 \cap F_2 = \emptyset$. Finally, we prove that both sets fit in the corner (x, y, z) . Recall that $w_R(s_i, b(i)) < e_B < t_i \leq x$. Since in addition $b(i) > y$, it follows that $C(s_i, y, b(i)) \subseteq C(x, y, z)$. (Because $w_R(s_i, b(i)) < x$, it is irrelevant whether $b(i) > z$ or not.) It is obvious that $C(x, b(i), z) \subseteq C(x, y, z)$, since $b(i) > y$. This concludes the proof of Inequality (4.2). \square

4.3 A Dynamic Programming Algorithm for Finding Optimal ITSs

In this section we give an $O(n^4)$ dynamic programming algorithm for computing an optimal ITS (Theorem 4.18). The main recursive routine $\mathcal{P}(x, y, z)$ for computing $P(x, y, z)$ is shown in Algorithm 1, which is based on the recursive formulas from the previous section (Proposition 4.10 and Lemma 4.12). At the end of this section, we summarize how this algorithm, together with Lemma 4.4, yields the $2k$ -approximation algorithm for UFPP in the case of $1/k$ -large tasks (Theorem 4.20).

First, we show that the recursion terminates, by showing that all recursive calls are made with arguments (x', y', z') that yield a ‘smaller’ corner than the original argument (x, y, z) .

Lemma 4.16 *Algorithm 1 terminates.*

Proof: We show that when given an input (x, y, z) , all recursive calls will be with arguments (x', y', z') such that the corner (x', y', z') is strictly smaller than (x, y, z) , with respect to the following size measure: For a corner (x, y, z) with $y \leq u_{\max}$ and $z \leq u_{\max}$, we define

$$\text{area}(x, y, z) = (x - w_L)(u_{\max} - y) + (w_R - x)(u_{\max} - z).$$

Algorithm 1: $\mathcal{P}(x, y, z)$: a recursive algorithm for computing $P(x, y, z)$

Input: A UFPP instance where u_{\max} is the maximum edge capacity, all edge capacities are distinct, and integers $0 \leq x \leq m$, $0 \leq y \leq u_{\max}$, $0 \leq z \leq u_{\max}$.

Output: $P(x, y, z)$.

```

1 if  $x = 0$  or  $y \geq u_{\{x-1,x\}}$  then  $y := u_{\max}$ ;
2 if  $x = m$  or  $z \geq u_{\{x,x+1\}}$  then  $z := u_{\max}$ ;
3 Compute  $w_L := w_L(x, y)$  and  $w_R := w_R(x, z)$ ;
4 if  $w_L = w_R$  then return  $0$ ;
5 if  $y = z$  then  $z := u_{\max}$ ;  $x := w_R$ ;
6 if  $y < z$  then
7   if  $u_{\{x-1,x\}} \leq z < u_{\{x,x+1\}}$  then return  $(\mathcal{P}(x, y, u_{\max}) + \mathcal{P}(x, u_{\max}, z))$ ;
8   return  $\max \left\{ \mathcal{P}(x-1, y, z), \right.$ 
       $\left. \max_{i \in C(x,y,z), t_i \leq x} \left( w_i + \mathcal{P}(s_i, y, b(i)) + \mathcal{P}(x, b(i), z) \right) \right\}$ ;
9 end
   (In the remaining case,  $y > z$  holds:);
10 if  $u_{\{x,x+1\}} \leq y < u_{\{x-1,x\}}$  then return  $(\mathcal{P}(x, y, u_{\max}) + \mathcal{P}(x, u_{\max}, z))$ ;
11 return  $\max \left\{ \mathcal{P}(x+1, y, z), \right.$ 
       $\left. \max_{i \in C(x,y,z), s_i \geq x} \left( w_i + \mathcal{P}(t_i, b(i), z) + \mathcal{P}(x, y, b(i)) \right) \right\}$ ;

```

Note that for all corners (x, y, z) that may be considered, $\text{area}(x, y, z)$ is a non-negative integer. Hence the lemma statement then follows.

A recursive call is made in Line 7 whenever a corner (x, y, z) is considered that satisfies the following bounds. Firstly $y < z \leq u_{\max}$, and therefore $x \geq 1$ and $y < u_{\{x-1,x\}}$. Hence $w_L < x$. If the if-condition is satisfied then $w_R > x$ holds. This shows that for both calls $\mathcal{P}(x, y, u_{\max})$ and $\mathcal{P}(x, u_{\max}, z)$, the area strictly decreases. Now consider Line 8, which is reached whenever the corner is a proper corner. In this case, when considering the corner $(x-1, y, z)$, the w_L value does not change, and the w_R value can only decrease. Since $y < z$ it then follows that $\text{area}(x-1, y, z) < \text{area}(x, y, z)$. Now consider a task $i \in C(x, y, z)$ with $t_i \leq x$. We have that $w_R(s_i, b(i)) < x$ and $b(i) > y$, so $\text{area}(s_i, y, b(i)) < \text{area}(x, y, z)$. Clearly, $\text{area}(x, b(i), z) < \text{area}(x, y, z)$. So in all recursive calls in Line 8, the area decreases. By symmetry, the same holds for Lines 9 and 10. \square

Lemma 4.17 *When given a corner (x, y, z) , Algorithm 1 returns $P(x, y, z)$.*

Proof: We prove the statement by induction over the total number of recursive calls (which is possible since Lemma 4.16 shows that the algorithm terminates). That is, we will prove that a call $\mathcal{P}(x, y, z)$ returns the value $P(x, y, z)$, assuming that this holds for recursive calls $\mathcal{P}(x', y', z')$ that may be made.

The modification in Line 1 yields an equivalent corner by Proposition 4.10(ii). By symmetry, the same holds for Line 2. If $w_L = w_R$ then $C(x, y, z) = \emptyset$, so the correct answer is returned in Line 4. The modification in Line 5 is correct as well (Proposition 4.10(i)). So we may assume now that $w_L < w_R$, and $y \neq z$.

Consider the case that $y < z$. This implies that $y < u_{\max}$, so $x \geq 1$ and $y < u_{\{x-1,x\}}$. If $u_{\{x-1,x\}} \leq z$ and $z < u_{\{x,x+1\}}$, then it is a non-proper corner, so Proposition 4.10 (iii) shows that Line 7 returns the correct answer. So if Line 8 is reached, the considered corner is proper. By Lemma 4.12, Line 8 returns the correct answer.

The case that $y > z$ is analog; by symmetric arguments, the answers given in Lines 9 and 10 are correct. This shows that Algorithm 1 returns the correct answer in every case. \square

Lemma 4.17 and Proposition 4.9 show that the profit of an optimal ITS can be computed by calling Algorithm 1 with the argument $(m, 0, u_{\max})$. However, to attain the desired complexity of $O(n^4)$, it is necessary to use a dynamic programming table that stores previously computed values, to avoid repeating the same recursive calls. The complexity then follows from the fact that this table needs at most n^3 entries.

Theorem 4.18 *There is an $O(n^4)$ algorithm for computing a maximum profit ITS for a UFPP instance.*

Proof: The algorithm is as follows. First, in time $O(n+m)$, we transform the given instance into an equivalent instance, in which all edge capacities are distinct (Lemma 4.6). As mentioned in Section 2, in time $O(n+m)$ we can transform the instance to an equivalent instance in which all vertices occur as the start or end vertex of some task, so we may furthermore assume that $m < 2n$. This preprocessing does not change n and cannot increase m . On the resulting instance we call Algorithm 1, with arguments $(x, y, z) = (m, 0, u_{\max})$, to obtain the total profit of an optimal ITS (Lemma 4.17, Proposition 4.9). With a straightforward extension, we can compute an optimal ITS itself.

A small modification is necessary to obtain a time complexity of $O(nm^3)$: whenever any value $P(x, y, z)$ is computed during the course of computation, this value is stored in a table. Whenever a recursive call to $\mathcal{P}(x, y, z)$ would be made, the algorithm instead returns the stored value $P(x, y, z)$, if it has been computed before (in a different recursion branch). This, together with Lemma 4.16, ensures that for every corner (x, y, z) the routine $\mathcal{P}(x, y, z)$ is called at most once during the course of computation. (Lemma 4.16 ensures that a given corner is not considered more than once in the *same* recursion branch.)

We call a corner (x, y, z) *relevant* if $y = 0$ or y is equal to the capacity of some edge, and if the same holds for z . Observe that when calling Algorithm 1 with a relevant corner as argument, then every possible recursive call is with a relevant corner as well. So the total number of recursive calls is bounded by the total number of relevant corners. There are at most $m+1$ possibilities for x , and at most $m+1$ for y and z (since these need to correspond to actual capacities), so the total number of recursive calls is bounded by $O(m^3)$. Now consider the complexity of a single call. The computations in Line 3 can be done in time $O(m)$. Ignoring recursive calls, Lines 7 and 9 take constant time. Lines 8 and 10 take time $O(n)$; for every task i , in constant time one can decide whether $i \in C(x, y, z)$ and $t_i \leq x$ (resp. $s_i \geq x$), and evaluate the given expressions. The remaining lines clearly take constant time, so it follows that a single call of Algorithm 1 takes time $O(n) + O(m) \subseteq O(n)$. Hence, the total complexity becomes $O(nm^3) \subseteq O(n^4)$. \square

Our results can be formulated as follows in terms of the Maximum Weight Independent Set of Rectangles (MWISR) problem:

Theorem 4.19 *MWISR can be solved in time $O(n^4)$ for instances with n axis-parallel rectangles, in which each rectangle contains a point (x, y) such that no point (x, y') with $y' > y$ is part of a rectangle.*

We summarize the results of Section 4 in the following theorem.

Theorem 4.20 *For every integer $k \geq 2$, there is an $O(n^4)$ time $2k$ -approximation algorithm for UFPP restricted to instances where every task is $\frac{1}{k}$ -large.*

Proof: In time $O(n^4)$ we compute an optimal ITS F for the instance (Theorem 4.18). This is a feasible UFPP solution (Proposition 4.2). Let F^* be an optimal UFPP solution. Then F^* can be partitioned into at most $2k$ ITSs (Lemma 4.4), so $w(F) \geq \frac{1}{2k}w(F^*)$. \square

5 Main Approximation Algorithms

We now apply the main results from the previous two sections to obtain our main approximation algorithm.

Theorem 5.1 *For every $\epsilon > 0$ there is a $(7 + \epsilon)$ -approximation algorithm for the UFPP problem.*

Proof: We partition the tasks into $\frac{1}{2}$ -small and $\frac{1}{2}$ -large tasks. For the $\frac{1}{2}$ -small tasks, we have a $(3 + \epsilon)$ -approximation algorithm (Theorem 3.14), and for the $\frac{1}{2}$ -large tasks, a 4-approximation algorithm (Theorem 4.20). Returning the best solution of the two then yields a $(7 + \epsilon)$ -approximation algorithm (Fact 2.1). \square

The running time of the above algorithm is dominated by the dynamic program which is needed for handling the ‘medium size’ tasks, see Proposition 3.11. Unfortunately, the exponent of the running time is quite large. However, in the following theorem we show that already a moderate running time of $O(n^4 \log n)$ suffices to obtain a constant factor approximation, by partitioning the tasks differently in order to avoid this expensive dynamic program. Hence, at the cost of a higher approximation factor, this algorithm has a much better running time.

Theorem 5.2 *There exists a 25.12-approximation algorithm for UFPP with running time $O(n^4 \log n)$.*

Proof: We partition the tasks into $\frac{1}{9}$ -small and $\frac{1}{9}$ -large tasks. Using Theorem 4.20, we obtain an 18-approximation algorithm for the $\frac{1}{9}$ -large tasks, that runs in time $O(n^4)$. For the $\frac{1}{9}$ -small tasks, we apply our framework with constants $\ell = 3$, $q = 5$. According to Lemma 3.2, we then choose $\beta := 2^{1-q} = 1/16$. We first verify that these choices satisfy the conditions of Lemma 3.8:

- $\delta \leq (1 - \beta)/2^\ell$ holds, since $(1 - \beta)/2^\ell = 15/128 > 15/135 = 1/9 = \delta$.
- $\delta/(1 - \beta) \leq (3 - \sqrt{5})/2$ holds, since $\delta/(1 - \beta) = 16/135 \approx 0.12 < 0.38 \approx (3 - \sqrt{5})/2$.




Figure 6.1: An example of a graph G with coloring α and the resulting instance $\text{UFPP}(G)$. (The demand of a task is indicated by the height of its bar.)

Hence we may apply Lemma 3.8 to conclude that for every integer k , in time $O(n^3 \log n)$, an $(f(\delta')/(1 - \beta), \beta)$ -approximative solution can be computed for the $\frac{1}{9}$ -small tasks in $F^{k, \ell}$, where $\delta' = \delta/(1 - \beta) = 16/135$, and

$$f(\delta') = \frac{1 + \sqrt{\delta'}}{1 - \sqrt{\delta'} - \delta'} \approx \frac{1.3443}{0.5372} < 2.503.$$

Therefore this is an (α, β) -approximative solution, with $\alpha = 2.503/(1 - \beta) < 2.67$. Our framework (Lemma 3.2) then gives an approximation algorithm for the instance consisting of all $\frac{1}{9}$ -small tasks, with complexity $O(n^4 \log n)$ and approximation ratio $\frac{\ell+q}{\ell} \cdot \alpha < 7.12$.

We apply the 18-approximation algorithm for the $\frac{1}{9}$ -large tasks and the 7.12-approximation algorithm for the $\frac{1}{9}$ -small tasks. Returning the best solution of these two gives a 25.12-approximation algorithm for the entire instance (Fact 2.1). \square

Chakrabarti et al. [17] showed that any α -approximation algorithm for UFPP implies a $(1 + \alpha + \epsilon)$ -approximation algorithm for UFP on cycles (also called ring networks). Note that on a cycle, for each task i two possible paths can be chosen from s_i to t_i . (Even though in [17] only UFP-NBA was considered, we observe that their argument applies to UFP as well.) This yields the following corollary.

Corollary 5.3 *For every $\epsilon > 0$ there is a $(8 + \epsilon)$ -approximation algorithm for the Unsplittable Flow Problem on Cycles.*

6 Strong NP-Hardness

In this section we prove that UFPP is strongly NP-hard for instances with demands in $\{1, 2, 3\}$, using a reduction from Maximum Independent Set in Cubic Graphs. Let G, k be an independent set instance, where G is a connected graph of maximum degree 3 with $G \neq K_4$. The question is whether G contains an independent set of size at least k . This problem is NP-hard [28] (even for *cubic* graphs, in which all degrees are exactly 3). Let $V(G) = \{v_1, \dots, v_n\}$, and $E(G) = \{e_1, \dots, e_m\}$. By Brooks' Theorem (see e.g. [33]), G is 3-colorable since G contains no K_4 . Such a coloring can be found in polynomial time [33], so for our polynomial transformation we may assume that a proper 3-coloring $\alpha : V(G) \rightarrow \{1, 2, 3\}$ is given.

We now construct an equivalent instance of UFPP as follows, see Figure 6.1. The path P has $2n + 2m + 1$ vertices, labeled $0, \dots, 2n + 2m$. For the proofs below it will be useful to distinguish different types of edges: an edge $\{i - 1, i\} \in E(P)$ (with $1 \leq i \leq 2n + 2m$) is a *left edge* if $i \leq 2m$ and a *right edge* otherwise. It is an *odd edge* if i is odd, and an *even edge* otherwise. Even left edges $\{2k - 1, 2k\}$ (with $1 \leq k \leq m$) will be associated with edges e_k of G , and even right edges $\{2k - 1, 2k\}$ (with $m + 1 \leq k \leq n + m$) will be associated with vertices v_{k-m} of G .

For every vertex $v_i \in V(G)$, introduce the following tasks: First, introduce one *long* task with start vertex 0 and end vertex $2m + 2i - 1$. Next, let $\sigma_1, \dots, \sigma_d$ be the indices of the edges incident with v_i , in increasing order. Introduce $d + 1$ *short* tasks with the following start and end vertex pairs: $(0, 2\sigma_1 - 1), (2\sigma_1, 2\sigma_2 - 1), \dots, (2\sigma_{d-1}, 2\sigma_d - 1), (2\sigma_d, 2m + 2i)$. For all of these tasks for vertex v_i , the demand is equal to $\alpha(v_i)$, and the profit is equal to $\alpha(v_i)n$ times the number of odd edges used by the task. The aforementioned tasks will be called the *high-profit tasks* (for v_i). Finally, for v_i we introduce one *low-profit task* from $2m + 2i - 1$ to $2m + 2i$ with profit 1 and demand $\alpha(v_i)$.

Doing this for all vertices gives the tasks of the UFPP instance. It remains to set the edge capacities of P :

- For odd left edges $e = \{2k - 2, 2k - 1\} \in E(P)$ for integer $1 \leq k \leq m$, set $u_e = \sum_{v \in V(G)} \alpha(v)$.
- For even left edges $e = \{2k - 1, 2k\} \in E(P)$ for integer $1 \leq k \leq m$, set $u_e = (\sum_{v \in V(G)} \alpha(v)) - 1$.
- For right edges $e = \{2k - 2, 2k - 1\} \in E(P)$ and $e = \{2k - 1, 2k\} \in E(P)$ for integer $m + 1 \leq k \leq n + m$, set $u_e = \sum_{i=k-m}^n \alpha(v_i)$.

This gives the instance $\text{UFPP}(G)$ of UFPP, to which we will refer in the remainder of this section. With T we will denote the set of all tasks of $\text{UFPP}(G)$. With an independent set I of G we associate a solution of $\text{UFPP}(G)$ of the following form.

Definition 6.1 *A set $F \subseteq T$ of tasks of $\text{UFPP}(G)$ is in standard form if:*

1. *For every $i \in \{1, \dots, n\}$:*
 - *Either F contains the long high-profit task for v_i and the low-profit task for v_i , but no short high-profit tasks for v_i ,*
 - *or F contains all short high-profit tasks for v_i , but neither the long high-profit task for v_i nor the low-profit task for v_i .*
2. *For every $\{v_i, v_j\} \in E(G)$, F does not contain both the long high-profit task for v_i and the long high-profit task for v_j .*

Given an independent set I of G , we can construct a solution F to $\text{UFPP}(G)$ by selecting the long high-profit task and the low-profit task for v_i if $v_i \in I$, and all short high-profit tasks for v_i otherwise. Since I is an independent set, this set F is in standard form; it satisfies Property 2 in the above definition. We will first show that the resulting task set F is a feasible solution to $\text{UFPP}(G)$, by showing that in fact any task set in standard form is a feasible solution. Secondly, we show that the task set F constructed this way is an optimal

solution for $\text{UFPP}(G)$. To this end, we first argue that the total profit of a solution F in standard form depends only on the number of long high-profit tasks in F . Subsequently, we will show that any optimal solution for $\text{UFPP}(G)$ is in standard form. Together this will show that the instances G and $\text{UFPP}(G)$ are equivalent (when considering them as decision problems with appropriately chosen target objective values).

Proposition 6.2 *Let $F \subseteq T$ be in standard form. Then F is a feasible solution for $\text{UFPP}(G)$.*

Proof: Since F is in standard form, for every $v_i \in V(G)$, every edge of P is used by at most one task for v_i . Therefore, for odd left edges $e = \{2k - 2, 2k - 1\} \in E(P)$ (with $1 \leq k \leq m$), the capacity $u_e = \sum_{v \in V(G)} \alpha(v)$ is not exceeded. Similarly, for right edges $e = \{2k - 2, 2k - 1\} \in E(P)$ (with $m + 1 \leq k \leq n + m$), the capacity $\sum_{i=k-m}^n \alpha(v_i)$ is not exceeded. Now consider even left edges $e = \{2k - 1, 2k\} \in E(P)$ with $1 \leq k \leq m$. These correspond to edges $e_k \in E(G)$. Let $e_k = \{v_i, v_j\}$. Since F is in standard form, it cannot contain high-profit tasks for both v_i and v_j . Note that no short high-profit tasks for v_i and v_j use the edge $\{2k - 1, 2k\}$. Hence the total capacity of e used by F is at most

$$\max \left\{ \left(\sum_{w \in V(G)} \alpha(w) \right) - \alpha(v_i), \left(\sum_{w \in V(G)} \alpha(w) \right) - \alpha(v_j) \right\} \leq \left(\sum_{w \in V(G)} \alpha(w) \right) - 1 = u_e.$$

□

Proposition 6.3 *Let $F \subseteq T$ be a solution in standard form for $\text{UFPP}(G)$, which contains x long high-profit tasks. Then $w(F) = x + \sum_{i=1}^n \alpha(v_i)n(m + i)$.*

Proof: Recall that for high-profit tasks for a vertex v_i , the profit equals $\alpha(v_i)n$ times the number of odd edges used by the task. The long-high profit task for v_i (from 0 to $2m + 2i - 1$) uses all odd edges $\{2k - 2, 2k - 1\}$ of P with $1 \leq k \leq m + i$. Hence if selected, it contributes $\alpha(v_i)n(m + i)$ to the total profit. Together, the low-profit tasks for v_i use exactly the same set of odd edges of P . So either way, the high-profit tasks for v_i contribute $\alpha(v_i)n(m + i)$ to the total profit. The single low-profit task for v_i , with a profit of 1, is selected if and only if the long high-profit task is selected. This yields the stated total profit. □

Lemma 6.4 *If G has an independent set I with $|I| = x$, then $\text{UFPP}(G)$ admits a solution of profit $x + \sum_{i=1}^n \alpha(v_i)n(m + i)$.*

Proof: Construct F by selecting the long high-profit task and the low-profit task for v_i if $v_i \in I$, and all short high-profit tasks for v_i otherwise. Since I is an independent set, this set F is in standard form. So by Proposition 6.2, it is a feasible solution to $\text{UFPP}(G)$. Proposition 6.3 now gives the total profit. □

Now we will show that any optimal solution to $\text{UFPP}(G)$ is in standard form, which will yield the converse of Lemma 6.4. We say that a solution F *fully uses the capacity* of an edge $e \in E(P)$ if $\sum_{i \in F \cap T_e} d_i = u_e$. The majority of the total profit of a solution is proportional to the capacity use of all odd edges. Since there is at least one solution that fully uses the capacity of all of these edges (a solution in standard form), the next proposition follows.

Proposition 6.5 *Let F be an optimal solution for UFPP(G). For every odd edge $e \in E(P)$, F fully uses the capacity of e .*

Proof: Recall that for every high-profit task with a demand of d , the profit equals dnx , where x is the number of odd edges used by the task. Hence the total profit of high-profit tasks in a feasible solution F is bounded by n times the sum of capacities of all such edges. This capacity sum is

$$m \sum_{i=1}^n \alpha(v_i) + \sum_{k=1}^n \sum_{i=k}^n \alpha(v_i) = \sum_{i=1}^n \alpha(v_i)(m+i).$$

Suppose that for at least one odd edge the full capacity is not used by F . The low-profit tasks in F can in total only yield a profit of n , so then the total profit is at most

$$n \left(\sum_{i=1}^n \alpha(v_i)(m+i) - 1 \right) + n = n \sum_{i=1}^n \alpha(v_i)(m+i).$$

Since there exists a feasible solution for UFPP(G) with strictly higher profit (Lemma 6.4), this shows that F is not optimal. \square

Now we will apply Proposition 6.5 to prove that every optimal solution for UFPP(G) is in standard form.

Proposition 6.6 *Let F be an optimal solution for UFPP(G). Then for every $i \in \{1, \dots, n\}$, F either contains no short high profit tasks for v_i , or contains all short high profit tasks for v_i .*

Proof: Suppose that the statement is not true. Then we may assume that there exists an even left edge $e = \{x, x+1\}$ in P such that F contains the short high profit task for v_i that ends at x , but not the short high profit task for v_i that starts at $x+1$ (the argument in the case where this is reversed is the same). Let $v_i v_j \in E(G)$ be the edge that e corresponds to, so v_i and v_j are the only vertices of G for which e is not included in the set of corresponding short high profit tasks. Since F fully uses the capacity of both of the odd edges $\{x-1, x\}$ and $\{x+1, x+2\}$ next to e (Proposition 6.5), F must contain tasks starting on x or $x+1$, with a total capacity of exactly $\alpha(v_i)$, to compensate for the fact that the short high profit task for v_i starting at $x+1$ is not included. But the only such task is a short high profit task for v_j , which has a capacity of $\alpha(v_j) \neq \alpha(v_i)$ (since α is a proper coloring of G), a contradiction. \square

Lemma 6.7 *Every optimal solution F for UFPP(G) is in standard form.*

Proof: We first prove that Property 1 of Definition 6.1 holds. Considering the capacity $\alpha(v_n)$ of the rightmost odd edge $e = \{2m+2n-2, 2m+2n-1\}$ shows that F cannot contain both the long high profit task for v_n and the short high profit task for v_n that uses this edge. However, since F fully uses the capacity of e (Proposition 6.5), exactly one of these cases is true. Combined with Proposition 6.6 this shows that F either contains the long high profit task for v_n , or all short high profit tasks. Since F is optimal, the low-profit task for v_n is included if and only if the long high profit task is included. This shows that Property 1

of Definition 6.1 holds for v_n . Then the same argument can be applied using the odd edge $\{2m + 2n - 4, 2m + 2n - 3\}$ to the left of e , which has capacity $\alpha(v_n) + \alpha(v_{n-1})$, to prove this property for v_{n-1} . Continuing this way proves Property 1 of Definition 6.1 for all v_i .

Next, we will prove that Property 2 of Definition 6.1 also holds for the optimal solution F . We will prove that the vertices of G for which F contains a long high-profit task form an independent set in G . More precisely, for any edge $\{v_i, v_j\} \in E(G)$, F does not contain both the long high-profit task for v_i and the long high-profit task for v_j . Let $e_k = \{v_i, v_j\}$ (with $k \in \{1, \dots, m\}$). According to our construction this edge is associated with an even left edge $e = \{2k - 1, 2k\} \in E(P)$ of $\text{UFPP}(G)$. The capacity of e is one unit smaller than the capacity of its adjacent odd edge $e' = \{2k - 2, 2k - 1\}$, and yet the capacity of e' is fully used by F (Proposition 6.5). The only tasks using e' but not e are a short task for v_i , and a short task for v_j . So for at least one of v_i and v_j , F includes this short task. Since we already showed in the first part of the proof that F cannot contain both a long high-profit task and a short high-profit task for the same vertex v_k , it follows that for at least one of v_i and v_j , F does not contain the long task. Hence Property 2 of Definition 6.1 also holds for F , and thus F is in standard form. \square

Combining the above lemma again with a calculation of the profit of a solution in standard form, this shows that any optimal solution to $\text{UFPP}(G)$ corresponds to an independent set of G :

Lemma 6.8 *If $\text{UFPP}(G)$ admits a solution of profit at least $x + \sum_{i=1}^n \alpha(v_i)n(m + i)$, then G has an independent set of size at least x .*

Proof: Consider an optimal solution F , so $w(F) \geq x + \sum_{i=1}^n \alpha(v_i)n(m + i)$. By Lemma 6.7, F is in standard form. Let I be the set of vertices v_i of G for which F contains the long high-profit task. Since F is in standard form, I is an independent set of G . Proposition 6.3 shows that $w(F) = |I| + \sum_{i=1}^n \alpha(v_i)n(m + i)$. It follows that $|I| \geq x$, so I is the desired independent set. \square

Our construction uses only polynomially bounded numbers (profits, demands and capacities), so even if the numbers are encoded in unary, this is a polynomial transformation. For this it is also essential that a 3-coloring α of G can be found in polynomial time [33]. Lemmas 6.4 and 6.8 show that G has an independent set of size at least x if and only if P, F admits a solution of profit at least $x + \sum_{i=1}^n \alpha(v_i)n(m + i)$. Since the problem of finding a maximum independent set is NP-hard when restricted to graphs of maximum degree 3 [28], this proves that UFPP is strongly NP-hard. The problem obviously lies in NP, which yields:

Theorem 6.9 *UFPP is strongly NP-complete when restricted to instances with demands in $\{1, 2, 3\}$.*

In fact, with a small addition we can show that the problem remains NP-complete when all edge capacities are equal. Let u_m be the maximum capacity used in the instance $\text{UFPP}(G)$ constructed above, and let $X = n + \sum_{i=1}^n \alpha(v_i)n(m + i)$ be an upper bound on the profit of any solution (Lemma 6.8). Note that X is bounded by a polynomial in n . For every edge $e = \{k, k + 1\}$ of P with $u_e < u_m$, we can increase the capacity to u_m , and introduce $u_m - u_e$ dummy tasks, from k to $k + 1$ with demand 1 and profit X . For a polynomial transformation

we must ensure that the number of dummy tasks is polynomially bounded. Given an instance with demands in $\{1, 2, 3\}$, a maximum capacity of $3n$ suffices, so w.l.o.g. $u_e \leq 3n$ holds for all capacities. Hence, on each edge at most $O(n)$ dummy tasks are introduced. Since the number of edges is bounded by $O(n)$, at most $O(n^2)$ dummy tasks with profit X and demand 1 need to be added. In an optimal solution, clearly all of the dummy tasks are selected, so this yields an equivalent instance. Therefore the previous theorem can be strengthened to:

Theorem 6.10 *UFPP is strongly NP-complete when restricted to instances with demands in $\{1, 2, 3\}$, where all edges have equal capacities.*

7 Discussion

In this paper, we presented the first constant factor approximation algorithm for the unsplittable flow problem on a path. From a technical point of view, our key contributions are the framework employing the (α, β) -approximation algorithms for the small tasks, and the geometrically inspired dynamic program for the large tasks. In particular, for the latter we establish a connection between UFPP and Maximum Weight Independent Set of Rectangles, two problems which at first glance might seem totally unrelated. It may be interesting to study whether our new techniques for finding maximum non-overlapping sets of rectangles can be applied to other geometric (packing) problems.

A challenging open question is whether a PTAS exists for UFPP, or whether the problem is APX-hard. Note that even in the much simpler case of uniform capacities (RAP), this question is open; the best known approximation algorithm has an approximation ratio of $2 + \epsilon$ [14]. It seems that for a PTAS one would have to use an approach that deals with all types of tasks (that is, small and large) in a combined manner, rather than using a best-of-two analysis like in our algorithm, or in other algorithms such as the $(2 + \epsilon)$ -approximation algorithm for UFPP-NBA [20].

We proved strong NP-hardness for UFPP with uniform capacities (RAP), even if the input is restricted to demands in $\{1, 2, 3\}$. Due to this result it would be interesting to know whether the special case of demands in $\{1, 2\}$ is also NP-hard, or admits a polynomial time algorithm. Note that if the demands are uniform, the problem can be solved in polynomial time with a small extension of the algorithm by Arkin and Silverberg [4].

Furthermore, our research raises the question what generalizations of UFPP also admit a constant factor approximation. One might consider generalizations such as those that have been considered for RAP and UFPP-NBA, either in a scheduling context or in a network flow context (see Section 1.2). In particular, an interesting open problem is whether the unsplittable flow problem on trees (without the NBA) also admits a constant factor approximation algorithm.

Acknowledgements

We would like to thank David Williamson for helpful comments on a draft of the paper, and Marek Chrobak for pointing us to [22].

References

- [1] P. K. Agarwal, M. van Kreveld, and S. Suri. Label placement by maximum independent set in rectangles. *Computational Geometry*, 11:209–218, 1998.
- [2] M. Andrews, J. Chuzhoy, S. Khanna, and L. Zhang. Hardness of the undirected edge-disjoint paths problem with congestion. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS '05)*, pages 226–244. IEEE, 2005.
- [3] M. Andrews and L. Zhang. Hardness of the undirected edge-disjoint paths problem. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC '05)*, pages 276–283. ACM, 2005.
- [4] E. M. Arkin and E. B. Silverberg. Scheduling jobs with fixed start and end times. *Discrete Appl. Math.*, 18:1–8, 1987.
- [5] Y. Azar and O. Regev. Combinatorial algorithms for the unsplittable flow problem. *Algorithmica*, 44:49–66, 2006.
- [6] N. Bansal, A. Chakrabarti, A. Epstein, and B. Schieber. A quasi-PTAS for unsplittable flow on line graphs. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 721–729. ACM, 2006.
- [7] N. Bansal, Z. Friggstad, R. Khandekar, and R. Salavatipour. A logarithmic approximation for unsplittable flow on line graphs. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '09)*, pages 702–709. SIAM, 2009.
- [8] A. Bar-Noy, R. Bar-Yehuda, A. Freund, J. Naor, and B. Schieber. A unified approach to approximating resource allocation and scheduling. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC '00)*, pages 735–744. ACM, 2000.
- [9] R. Bar-Yehuda, M. Beder, and Y. Cohen. Approximation algorithms for bandwidth and storage allocation. Technical report, 2005. available at <http://www.cs.technion.ac.il/~reuven/PDF/sap.pdf>.
- [10] R. Bar-Yehuda, M. Beder, Y. Cohen, and D. Rawitz. Resource allocation in bounded degree trees. In *Proceedings of the 14th Annual European Symposium Algorithms (ESA '06)*, volume 4168 of *LNCS*, pages 64–75. Springer, 2006.
- [11] M. Bartlett, A. M. Frisch, Y. Hamadi, I. Miguel, A. Tarim, and C. Unsworth. The temporal knapsack problem and its solution. In *Proceedings of Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR '05)*, volume 3524 of *LNCS*, pages 34–48. Springer, 2005.
- [12] A. Baveja and A. Srinivasan. Approximation algorithms for disjoint paths and related routing and packing problems. *Math. Oper. Res.*, 25:255–280, May 2000.
- [13] V. Bonifaci, A. Marchetti-Spaccamela, and S. Stiller. A constant-approximate feasibility test for multiprocessor real-time scheduling. In *Proceedings of the 16th Annual European Symposium on Algorithms (ESA '08)*, volume 5193 of *LNCS*, pages 210–221. Springer, 2008.

- [14] G. Calinescu, A. Chakrabarti, H. Karloff, and Y. Rabani. Improved approximation algorithms for resource allocation. In *Proceedings of the 9th International Conference on Integer Programming and Combinatorial Optimization (IPCO '02)*, volume 2337 of *LNCS*, pages 401–414. Springer, 2002.
- [15] V. Chakaravarthy, A. Kumar, S. Roy, and Y. Sabharwal. Resource allocation for covering time varying demands. In *Proceedings of the 19th Annual European Symposium on Algorithms (ESA '11)*, volume 6942 of *LNCS*, pages 543–554. Springer, 2011.
- [16] V.T. Chakaravarthy, G.R. Parija, S. Roy, Y. Sabharwal, and A. Kumar. Minimum cost resource allocation for meeting job requirements. In *Proceedings of the 2011 IEEE International Parallel & Distributed Processing Symposium (IPDPS '11)*, pages 14–23. IEEE, 2011.
- [17] A. Chakrabarti, C. Chekuri, A. Gupta, and A. Kumar. Approximation algorithms for the unsplittable flow problem. *Algorithmica*, 47:53–78, 2007.
- [18] P. Chalermsook and J. Chuzhoy. Maximum independent set of rectangles. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '09)*, pages 892–901. SIAM, 2009.
- [19] C. Chekuri, A. Ene, and N. Korula. Unsplittable flow in paths and trees and column-restricted packing integer programs. In *Proceedings of the 12th International Workshop (APPROX '09) and 13th International Workshop (RANDOM '09) on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 5687 of *LNCS*, pages 42–55. Springer, 2009.
- [20] C. Chekuri, M. Mydlarz, and F. Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Trans. on Algorithms*, 3, Article 27, 2007.
- [21] B. Chen, R. Hassin, and M. Tzur. Allocation of bandwidth and storage. *IIE Transactions*, 34:501–507, 2002.
- [22] M. Chrobak, G. Woeginger, K. Makino, and H. Xu. Caching is hard – even in the fault model. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA '10)*, volume 6346 of *LNCS*, pages 195–206. Springer, 2010.
- [23] A. Darmann, U. Pferschy, and J. Schauer. Resource allocation with time intervals. *Theor. Comp. Sc.*, 411:4217–4234, 2010.
- [24] C. W. Duin and E. Van Sluis. On the complexity of adjacent resource scheduling. *J. of Scheduling*, 9:49–62, 2006.
- [25] F. Eisenbrand and T. Rothvoß. A PTAS for static priority real-time scheduling with resource augmentation. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP '08)*, volume 5125 of *LNCS*, pages 246–257. Springer, 2008.
- [26] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '01)*, pages 671–679. SIAM, 2001.

- [27] A. M. Frieze and M. R. B. Clarke. Approximation algorithms for the m -dimensional 0 – 1 knapsack problem: worst-case and probabilistic analyses. *European J. Oper. Res.*, 15:100–109, 1984.
- [28] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theor. Comp. Sc.*, 1:237–267, 1976.
- [29] N. Garg, V. V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18:3–20, 1997.
- [30] S. Khanna, S. Muthukrishnan, and M. Paterson. On approximating rectangle tiling and packing. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '98)*, pages 384–393. SIAM, 1998.
- [31] T. W. Lam and K. K. To. Trade-offs between speed and processor in hard-deadline scheduling. In *Proceedings of the 10th Annual ACM-SIAM symposium on Discrete algorithms (SODA '99)*, pages 623–632. SIAM, 1999.
- [32] S. Leonardi, A. Marchetti-Spaccamela, and A. Vitaletti. Approximation algorithms for bandwidth and storage allocation problems under real time constraints. In *Proceedings of the 20th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS '00)*, volume 1974 of *LNCS*, pages 409–420. Springer, 2000.
- [33] L. Lovász. Three short proofs in graph theory. *J. Combinatorial Theory Ser. B*, 19:269–271, 1975.
- [34] F. Nielsen. Fast stabbing of boxes in high dimensions. *Theor. Comp. Sc.*, 246:53 – 72, 2000.
- [35] C. Phillips, C. Stein, E. Torng, and J. Wein. Optimal time-critical scheduling via resource augmentation. *Algorithmica*, 32:163–200, 2008.
- [36] C. A. Phillips, R. N. Uma, and J. Wein. Off-line admission control for general scheduling problems. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '00)*, pages 879–888. SIAM, 2000.