# Faster $(1+\epsilon)$-Approximation for Unsplittable Flow on a Path via Resource Augmentation and Back 

Fabrizio Grandoni $\square$<br>IDSIA, USI-SUPSI, Lugano, Switzerland<br>Tobias Mömke $\square$ ©<br>Department of Computer Science, Universität Augsburg, Germany<br>Andreas Wiese $\square$<br>Department of Industrial Engineering, Universidad de Chile, Santiago, Chile


#### Abstract

Unsplittable flow on a path (UFP) is an important and well-studied problem. We are given a path with capacities on its edges, and a set of tasks where for each task we are given a demand, a subpath, and a weight. The goal is to select the set of tasks of maximum total weight whose total demands do not exceed the capacity on any edge. UFP admits an $(1+\epsilon)$-approximation with a running time of $n^{O_{\epsilon}(\operatorname{poly}(\log n))}$, i.e., a QPTAS [Bansal et al., STOC 2006; Batra et al., SODA 2015] and it is considered an important open problem to construct a PTAS. To this end, in a series of papers polynomial time approximation algorithms have been developed, which culminated in a $(5 / 3+\epsilon)$-approximation [Grandoni et al., STOC 2018] and very recently an approximation ratio of $\left(1+\frac{1}{e+1}+\varepsilon\right)<1.269$ [Grandoni et al., 2020]. In this paper, we address the search for a PTAS from a different angle: we present a faster $(1+\epsilon)$-approximation with a running time of only $n^{O_{\epsilon}(\log \log n)}$. We first give such a result in the relaxed setting of resource augmentation and then transform it to an algorithm without resource augmentation. For this, we present a framework which transforms algorithms for (a slight generalization of) UFP under resource augmentation in a black-box manner into algorithms for UFP without resource augmentation, with only negligible loss.


2012 ACM Subject Classification Theory of computation $\rightarrow$ Dynamic programming; Theory of computation $\rightarrow$ Packing and covering problems; Theory of computation $\rightarrow$ Rounding techniques

Keywords and phrases Approximation Algorithms, Unsplittable Flow, Dynamic Programming
Digital Object Identifier 10.4230/LIPIcs.ESA.2021.49
Funding Fabrizio Grandoni: Partially supported by the SNSF Excellence Grant 200020B_182865/1. Tobias Mömke: Partially supported by the DFG Grant 439522729 (Heisenberg-Grant), the ERC Advanced Investigators Grant 695614 (POWVER), and by DFG Grant 389792660 as part of TRR 248 (CPEC).
Andreas Wiese: Partially supported by FONDECYT Regular grants 1170223 and 1200173.

## 1 Introduction

The unsplittable flow on a path problem (UFP) is a natural packing problem which has received a lot of attention, e.g., $[3,17,15,6,7,4,8,9]$. We are given a path $G=(V, E)$ with a capacity $u(e) \in \mathbb{N}_{0}$ for each edge $e \in E$ and a set of tasks $T$. For each task $i \in T$ we are given a sub-path $P(i)$ of $E$, a demand $d(i) \in \mathbb{N}_{0}$, and a weight (or profit) $w(i) \in \mathbb{N}_{0}$. For any set of tasks $T^{\prime} \subseteq T$, we define $d\left(T^{\prime}\right):=\sum_{i \in T^{\prime}} d(i)$ and $w\left(T^{\prime}\right):=\sum_{i \in T^{\prime}} w(i)$ and for each $e \in E$ we define $T_{e} \subseteq T$ to be the set of all tasks $i \in T$ using $e$, i.e., such that $e \in P(i)$. Similarly we sometimes say that $i$ is contained, contains, or intersects a subpath $P$ if $P(i)$ does. The goal is to select a set of tasks $T^{\prime} \subseteq T$ of maximum total weight $w\left(T^{\prime}\right)$ such that $T^{\prime}$ obeys the edge capacities, i.e., $d\left(T^{\prime} \cap T_{e}\right) \leq u(e)$ for each edge $e \in E$. We denote by $n$ the size of the input, and hence in particular $|T| \leq n$ and $|E| \leq n$.

© Fabrizio Grandoni, Tobias Mömke, and Andreas Wiese;
licensed under Creative Commons License CC-BY 4.0

UFP is a generalization of KnAPSACK (i.e., if $|E|=1$ ) and it is motivated by various applications. For example, the path $G$ can represent a network with a chain of communication links in which we seek to select the most profitable set of possible transmissions that obey the given edge capacities. Also, the edges in $E$ can correspond to discrete time slots, each task models a job that we might want to execute, and the edge capacities model the available amount of a resource shared by the jobs like energy or machines. Also, there is a connection between UFP and general caching, i.e., where pages can have different (possible non-unit) sizes and different costs for being evicted [12].

There is a $(1+\epsilon)$-approximation algorithm known for UFP with running time of $n^{O_{\epsilon}(\log n)}$ [7], i.e., a QPTAS (improving an earlier QPTAS in [5]), and it has been a long-standing open problem to construct a PTAS. Towards this goal, polynomial time approximation algorithms for UFP have been developed and the best known approximation ratio has been gradually improved from $O(\log n)[6]$ to $7+\epsilon[8], 2+\epsilon[3]$, and $5 / 3+\epsilon[17]$, while the currently best known ratio ratio is $1+\frac{1}{1+e}+\epsilon<1.269$ [15].

### 1.1 Our Contribution

In this paper, we contribute to the search for a PTAS for UFP from a different angle: we improve the running time of the known QPTAS for UFP [7] and present a $(1+\epsilon)$ approximation with a running time of only $n^{O_{\epsilon}(\log \log n)}$. Hence, our result is in the same spirit as similar improvements for the Maximum Independent Set of Rectangles problem in [13] and precedence constrained scheduling for unit-size jobs in [18].

We first present our result for the case of resource augmentation, i.e., when we allow ourselves to increase the edge capacities by a factor of $1+\delta$ for some constant $\delta>0$ while the compared optimal solution OPT does not have this privilege. Let $u_{\min }$ and $u_{\max }$ denote the minimum and maximum edge capacities, respectively. In that setting, it is easy to show that we can assume that the edge capacities are in a constant range, namely $u_{\max }=O_{\varepsilon, \delta}\left(u_{\min }\right)$. We classify each task $i \in T$ to be small or large, depending on whether $i$ uses a relatively small or a relatively large fraction of the available capacity on the edges of $P(i)$. Since we have a constant range of edge capacities, one can show easily that each edge $e$ can be used by only a constant number of large tasks in $O P T$.

The high level strategy in the known QPTASs [5, 7] is to take the middle edge $e^{*}$, guess the large tasks using $e^{*}$, split the small tasks using $e^{*}$ into $O(\log n)$ or $(\log n)^{O(1)}$ groups, and guess an under-estimating capacity profile for each group with $O_{\epsilon}(1)$ uniform steps. In more detail, in the latter step one guesses for each edge $e \in E$ approximately how much capacity from $u(e)$ is used in $O P T$ by the tasks crossing $e^{*}$ in each group. For each group one argues that one loses only a factor of $1+\epsilon$ in the profit by underestimating the true capacity. Then one recurses on the subpaths of $E$ on the left of $e^{*}$ and on the right of $e^{*}$ which yields a recursion depth of $O(\log n)$.

Instead, when we consider the small tasks using $e^{*}$, we use only one over-estimating profile with non-uniform step size for all tasks together with only $O_{\varepsilon, \delta}(\log \log n)$ steps. Each step is a power of $1+\delta$ in $\left[\frac{\delta}{\log n} u_{\min }, u_{\max }\right]$. This yields $O_{\varepsilon, \delta}(\log \log n)$ steps since the profile is monotone on the left and on the right of $e^{*}$ and $u_{\max }=O_{\varepsilon, \delta}\left(u_{\min }\right)$. Also, our justification for the error is very different. On some edges $e$ the error is at most $\frac{\delta}{\log n} u_{\text {min }}$ which accumulates to at most $O(\delta) u_{\min } \leq O(\delta) u(e)$ during the recursion. On the other edges $e$ the error is at most a $\delta$-fraction of the total demand used on $e$ by tasks crossing $e^{*}$; over the recursion this can add up to at most $\delta u(e)$ since in $O P T$ edge $e$ is used by tasks with a total demand of at most $u(e)$.

When we recurse we employ another novelty: we consider the tasks crossing at least one of two specially defined edges $e_{1}^{*}, e_{2}^{*}$, and recurse in three subpaths induced by them: $e_{1}^{*}$ is like before the middle edge of current subpath. The other edge $e_{2}^{*}$ is chosen such that half of previously guessed steps are on the left and the other half on the right (like in the QPTAS in [7]). This ensures that at the same time the recursion depth is $O(\log n)$ and each recursive call is described by a subpath of $E$ and only $O_{\varepsilon, \delta}(\log \log n)$ steps from previously guessed profiles. Thus, we can embed this recursion into a dynamic program (DP) with DP-table of size $n^{O_{\varepsilon, \delta}}(\log \log n)$.

Then, we present a new framework using which one can transform algorithms for UFP under resource augmentation (like ours) in a black-box manner into algorithms for UFP without resource augmentation, while losing only a factor of $1+\epsilon$ in the approximation ratio. To this end, we define a slight generalization of UFP that we denote by Bonus-UFP. The key difference to UFP is that in addition to profit from normal tasks, one receives a bonus for subpaths which do not completely contain the path of any selected task. Also, its instances are required to have a simpler structure than general UFP instances, yielding similar properties as obtained via resource augmentation. For example, on each edge one is allowed to select only a constant number of large tasks (independently of the actual edge capacities!), so they can be easily guessed in time $n^{O_{\epsilon}(1)}$ for each edge. Also, the capacity allocated for small tasks in $O P T$ is intuitively in a constant range (similarly as above) and the notion of resource augmentation is defined such that this capacity for the small tasks is increased by a factor $1+\delta$.

Our framework directly transforms any algorithm for Bonus-UFP under resource augmentation to an algorithm for UFP without resource augmentation, while increasing the approximation ratio only by a factor $1+\epsilon$. The transformation uses the slack-lemma [16] which was employed in previous algorithms for UFP [16, 15, 17] in order to gain free capacity on the edges within quite complicated dynamic programs. With our framework one does not need to apply the slack-lemma and construct this technical machinery "by hand" anymore, but it is sufficient to design an algorithm for Bonus-UFP under resource augmentation and then the framework does the transformation automatically.

- Theorem 1 (informal). Given an $\alpha$-approximation algorithm for Bonus-UFP under resource augmentation with a running time of $T(n)$, we can construct a $(1+\epsilon) \alpha$-approximation algorithm for UFP (without resource augmentation) with a running time of $T(n) n^{O_{\epsilon}(1)}$.

We hope that this facilitates future research on UFP for (eventually) finding a PTAS. In particular, we believe that if one constructs an algorithms for UFP under resource augmentation then it is very likely that it can be adjusted to an algorithm for Bonus-UFP under resource augmentation. For example, we demonstrate that this can be easily done with our algorithm for UFP under resource augmentation above, which yields the following theorem.

- Theorem 2. For any $\epsilon>0$ there is a $(1+\epsilon)$-approximation algorithm for UFP with a running time of $n^{O_{\epsilon}(\log \log n)}$.


### 1.2 Other related work

There are some special cases of UFP for which PTASs are known, for example when there are $O(1)$ edges such that each input task uses one of them [16], each input task can be selected an unbounded number of times [16], or when the profit of each task is proportional to its demand [7]. Also, there is an FPT- $(1+\epsilon)$-algorithm known for the unweighted case of UFP where the fixed parameter is the cardinality of the optimal solution [19].

Variations of UFP have been studied like bagUFP where the input tasks are partitioned into bags, and we are allowed to select at most one task from each bag [14]. Also, since the natural LP-formulation of UFP suffers from an integrality gap of $\Omega(n)$ [9], stronger LP-formulations have been investigated [10, 2]. Furthermore, unsplittable flow has been studied on trees, where the best known results are a $O\left(\log ^{2} n\right)$-approximation [10]. This was generalized to a $O(k \cdot \log n)$-approximation [1] for submodular objectives where $k$ denotes the pathwidth of the given tree (bounded by $O(\log n)$ ).

## 2 Faster approximation scheme for UFP with resource augmentation

We first provide a $(1+\epsilon)$-approximation algorithm for UFP with a running time of $n^{O_{\epsilon, \delta}(\log \log n)}$ for the simplified setting of $\delta$-resource augmentation where we permit the algorithm to compute a solution $S$ where for each edge $e \in E$ we have that $\sum_{i \in S \cap T_{e}} d(i) \leq(1+\delta) u(e)$ and we require the optimal solution OPT to respect the original capacities $u(\cdot)$, i.e., $\sum_{i \in \mathrm{OPT} \cap T_{e}} d(i) \leq u(e)$ for each $e \in E$. For a UFP instance with a path $G=(V, E)$ and edge capacities $u(\cdot)$, we define $u_{\min }:=\min _{e \in E} u(e)$ and $u_{\max }:=\max _{e \in E} u(e)$. First, we use the resource augmentation to reduce the general case to the setting of a constant range of (polynomially bounded) edge capacities where $u_{\max } \leq u_{\min } / \delta^{(2 / \epsilon)}$.

- Lemma 3. Assume that for any constants $\epsilon>0, \delta>0$, there is an $\alpha$-approximation algorithm under $\delta$-resource augmentation for UFP with a running time of $T(n)$ for instances such that $u_{\min }=n / \delta$ and $u_{\max } \leq u_{\min } / \delta^{2 / \epsilon}$. Then there is an $(\alpha+\epsilon)$-approximation algorithm for UFP under $4 \delta$-resource augmentation with a running time of $T(n)\left(n \log u_{\max }\right)^{O_{\epsilon, \delta}(1)}$.

Due to Lemma 3, in the following we assume that we are given an instance of UFP such that $n / \delta \leq u(e) \leq n /(\eta \delta)$ for each $e \in E$ and some $\epsilon, \delta>0$, with $\eta:=\delta^{2 / \epsilon}$, and that we are in the setting of $\delta$-resource augmentation. By contracting edges suitably we can assume w.l.o.g. that each edge $e \in E$ is the first or the last edge of the path $P(i)$ of some input task $i \in T$.

### 2.1 Recursive Algorithm

We describe our algorithm first as a recursive algorithm and then embed it into a dynamic program which will have a running time of $n^{O_{\epsilon, \delta}(\log \log n)}$.

Let $e^{*} \in E$ be an edge in the middle of $E$, i.e., such that at most $\lfloor|E| / 2\rfloor$ edges are on the left of $e^{*}$ and at most $\lfloor|E| / 2\rfloor$ edges are on the right of $e^{*}$. We would like to guess the tasks in $\mathrm{OPT} \cap T_{e^{*}}$. Since there are too many possibilities for this, we first guess approximately how much capacity the tasks in OPT $\cap T_{e^{*}}$ use on each edge $e \in E$. For any set of tasks $S \subseteq T$, we define $u_{S}(e):=\sum_{i \in S \cap T_{e}} d(i)$ for each $e \in E$. Observe that $u_{T_{e *} \text { ПOPT }}$ is non-decreasing on the left of $e^{*}$ and non-increasing on the right of $e^{*}$, since all tasks in $T_{e^{*}} \cap$ OPT use $e^{*}$. We will guess an over-estimating profile $\mathrm{mp}_{T_{e^{*}} \text { ПОРT }}: E \rightarrow \mathbb{N}$ of $u_{T_{e^{*}} \cap \mathrm{OPT}}(e)$ defined below, with the properties that

- for each edge $e$ it holds that $\operatorname{mp}_{T_{e^{*} \text { ПOPT }}}(e) \geq u_{T_{e^{*}} \text { ПОРT }}(e)$,
- $\mathrm{mp}_{T_{e}{ }^{*} \text { ПOPT }}$ is a step-function with only $O_{\epsilon, \delta}(\log \log n)$ steps (i.e., $E$ can be partitioned into $O_{\epsilon, \delta}(\log \log n)$ subpaths and $\mathrm{mp}_{T_{e^{*}} \text { ПOPT }}$ is constant on each subpath), and
- $\operatorname{mp}_{T_{e^{*}} \text { ПОРT }}(e)$ and $u_{T_{e^{*}} \text { ПOPT }}(e)$ do not differ too much on each edge $e \in E$.

Formally, let $D:=\log n$; our overall algorithm will intuitively be a recursion with $D$ levels. For any set of tasks $S \subseteq T$ we define the minimal profile $\mathrm{mp}_{S}: E \rightarrow \mathbb{N}$ of $S$ as follows. Let $\theta:=\left\lfloor u_{\text {min }} \cdot \delta / D\right\rfloor$. If for an edge $e \in E$ it holds that $d\left(S \cap T_{e}\right)=0$ then we define


Figure 1 Overestimating profile $\operatorname{mp}\left(u_{T_{e^{*}} \text { ПОРт }}\right)$ (the red line). The edge $e^{\prime}$ halves the number of steps and $e^{\prime \prime}$ the length of the path on the left-hand side of $e^{*}$.
$\operatorname{mp}_{S}(e):=0$. Otherwise, we define $\operatorname{mp}_{S}(e):=\left\lfloor\theta(1+\delta)^{k_{e}}\right\rfloor$ where $k_{e}$ is the smallest integers such that $\theta(1+\delta)^{k_{e}} \geq d\left(S \cap T_{e}\right)$ (see Figure 1). It turns out that $\mathrm{mp}_{S}(\cdot)$ is a step-function with only $O_{\delta, \epsilon}(\log \log n)$ steps, assuming that there is an edge $\bar{e}$ that is used by all tasks in $S$ (like the edge $e^{*}$ above).

Lemma 4. There is a value $\Gamma=O_{\delta, \epsilon}(\log D)=O_{\delta, \epsilon}(\log \log n)$ such that for every edge $\bar{e} \in E$ and every set $S \subseteq T_{\bar{e}}$ with $d\left(S \cap T_{e}\right) \leq u(e)$ for each $e \in E$, we have that $\mathrm{mp}_{S}(\cdot)$ is a step-function with at most $\Gamma$ steps.

Proof. Recall that by Lemma $3 u_{\min }=n / \delta$ and $u_{\max } \leq n /(\delta \eta)$. We therefore search for an upper bound on the smallest value $k$ such that $\left\lfloor\theta(1+\delta)^{k}\right\rfloor \geq n /(\delta \eta)$. We have that

$$
\begin{aligned}
k & \in O\left(\log _{1+\delta}\left(\frac{n}{\delta \eta \theta}\right)\right)=O\left(\log _{1+\delta}\left(\frac{n D}{\eta \delta^{2} u_{\min }}\right)\right) \\
& =O\left(\log _{1+\delta}\left(\frac{D}{\eta \delta^{2}}\right)\right)=O_{\delta, \epsilon}(\log D)=O_{\delta, \epsilon}(\log \log n)
\end{aligned}
$$

The second claim follows due to monotonicity of the profile formed by the tasks $\mathrm{OPT}_{S} \cap T_{e}$.
For any set $S$, the profile $\operatorname{mp}_{S}(\cdot)$ overestimates the true demand of the tasks $S$ on each edge $e$. Therefore, we define the error of this estimation by $\operatorname{err}\left(\mathrm{mp}_{S}, e\right):=\mathrm{mp}_{S}(e)-u_{S}(e)$ for each edge $e \in E$. Intuitively, in our algorithm we will guess profiles $\mathrm{mp}_{S}$ in each of the $D=\log n$ recursion levels, such that at the end each edge $e$ is in the support of at most $2 D$ of these profiles. A key insight is the following lemma which implies that for each edge $e$, the sum of the errors of these $2 D$ profiles is bounded by the extra space of $\delta u(e)$ due to the resource augmentation. Intuitively, if $\mathrm{mp}_{S}(e)=\theta$ on some edge $e$, then the error is at most $\theta=u_{\text {min }} \cdot \delta / D$ and for $2 D$ profiles the errors of this type can accumulate to at most $2 D \cdot \theta \leq 2 \delta u_{\min } \leq 2 \delta u(e)$. On the other hand, if $\operatorname{mp}_{S}(e)>\theta$ then the error is at most $\delta \cdot d\left(S \cap T_{e}\right)$, and this can accumulate to at most $\delta \cdot u(e)$ for all profiles together if these profiles correspond to disjoint sets of tasks $S_{1}, S_{2}, \ldots, S_{D^{\prime}}$ that use at most $u(e)$ units of capacity on $e$ altogether, for any $D^{\prime}$.

- Lemma 5. Let $S_{1}, S_{2}, \ldots, S_{D^{\prime}} \subseteq T$ be disjoint sets of tasks with $D^{\prime} \in O(D)$. If $\sum_{j} d\left(S_{j} \cap\right.$ $\left.T_{e}\right) \leq u(e)$ for an edge $e$, then $\sum_{j=1}^{D^{\prime}} \operatorname{err}\left(\operatorname{mp}_{S_{j} \cap T_{e}}, e\right) \leq O(\delta) \cdot u(e)$.

Proof. For each index $j$ and edge $e$ such that $\mathrm{mp}_{S_{j} \cap T_{e}}(e)=\left\lfloor\theta(1+\delta)^{k_{j}}\right\rfloor$ for some $k_{j}$, the error $\operatorname{err}\left(\mathrm{mp}_{S_{j} \cap T_{e}}, e\right)$ is bounded from above by $\lfloor\theta\rfloor$ if $k_{j}=0$ and by $\left\lfloor\theta \cdot(1+\delta)^{k_{j}}\right\rfloor-\left\lceil\theta \cdot(1+\delta)^{k_{j}-1}\right\rceil \leq$ $\theta(1+\delta)^{k_{j}-1} \cdot \delta \leq \delta d\left(S_{j}\right)$ if $k_{j}>0$. (Observe that the rounding is valid because the demands are integers.) Thus $\operatorname{err}\left(\mathrm{mp}_{S_{j} \cap T_{e}}, e\right) \leq \max \left\{\theta, \delta d\left(S_{j}\right)\right\}$. Summing up over all errors we obtain

$$
\sum_{j=1}^{D^{\prime}} \operatorname{err}\left(\mathrm{mp}_{S_{j} \cap T_{e}}, e\right) \leq D^{\prime} \cdot \theta+\sum_{j} \delta d\left(S_{j} \cap T_{e}\right) \leq \frac{D^{\prime} u_{\min } \delta}{D}+\delta u(e) \in O(\delta) u(e)
$$

As mentioned above, we guess $\mathrm{mp}_{T_{e^{*} \cap \mathrm{OPT}}}$, which can be done in time $n^{O_{\delta, \epsilon}(\log \log n)}$ due to Lemma 4. Then, we compute the essentially most profitable set of tasks $T^{\prime} \subseteq T_{e^{*}}$ that fits into $\mathrm{mp}_{T_{e^{*} \cap \text { OPT }}}$, i.e., such that $d\left(T^{\prime} \cap T_{e}\right) \leq \operatorname{mp}_{T_{e^{*} \cap \text { OPT }}}(e)$ on each edge $e \in E$. This can be done using a PTAS in [16] for rooted UFP instances in which all input tasks share a common edge (like the edge $e^{*}$ in our case).

- Theorem 6 ([16]). There is a PTAS for instances of UFP in which there is an edge that is used by every input task (rooted UFP). The same holds if there exist $O(1)$ edges such that each task uses at least one of them.

We recurse on the subpaths on the left and on the right of $e^{*}$. Let us consider the left subpath, i.e., let $E_{L}$ denote the path from the left-most edge of $E$ up to $e^{*}$, not including $e^{*}$. (The recursion to the right is analogous.) We subdivide $E_{L}$ into three parts determined by two edges $e^{\prime}, e^{\prime \prime}$ in $E_{L}$ (or two parts if these edges coincide). We choose the first edge $e^{\prime} \in E_{L}$ on the left of $e^{*}$ such that the number of steps in the profile $\bar{u}:=\mathrm{mp}_{\mathrm{OPT} \cap T_{e^{*}}}$ within $E_{L}$ is halved (a similar trick was used in [7]). Formally, let $\Gamma=O_{\delta, \epsilon}(\log \log n)$ such that $\mathrm{mp}_{\mathrm{OPT} \cap T_{e^{*}}}$ is a step-function with $\Gamma$ steps. We choose $e^{\prime}$ such that $\mathrm{mp}_{\mathrm{OPT} \cap T_{e^{*}}}$ restricted to the subpath of $E_{L}$ on the left of $e^{\prime}$ is a step-function with at most $\lceil\Gamma / 2\rceil$ steps, and the same is true for $\mathrm{mp}_{\mathrm{OPT} \cap T_{e^{*}}}$ restricted to the subpath of $E_{L}$ on the right of $e^{\prime}$. The second edge $e^{\prime \prime}$ lies in the middle of $E_{L}$ such that on the left of $e^{\prime \prime}$ there are at most $\left\lfloor\left|E_{L}\right| / 2\right\rfloor$ edges of $E_{L}$ and the same is true on the right of $e^{\prime \prime}$. We may assume without loss of generality that $e^{\prime}$ lies on the left of $e^{\prime \prime}$.

We guess the profiles $\bar{u}^{\prime}:=\mathrm{mp}_{\left(\mathrm{OPT} \cap T_{e^{\prime}}\right) \backslash T_{e *}}$ and $\bar{u}^{\prime \prime}:=\mathrm{mp}_{\left(\mathrm{OPT} \cap T_{e^{\prime \prime}}\right) \backslash\left(T_{e^{*}} \cup T_{e^{\prime}}\right)}$ and apply Theorem 6 to compute essentially the most profitable subset of tasks that fit into $\bar{u}^{\prime}$ (and $\bar{u}^{\prime \prime}$ ), among the input tasks $i \in T_{e^{\prime}}$ with $P(i) \subseteq E_{L}$ (among the input tasks $i \in T_{e^{\prime \prime}} \backslash T_{e^{\prime}}$ with $\left.P(i) \subseteq E_{L}\right)$. Then we recurse on each of the up to three components of $E_{L} \backslash\left\{e^{\prime}, e^{\prime \prime}\right\}$, where the parameter is the respective subpath and the profile $\bar{u}+\bar{u}^{\prime}+\bar{u}^{\prime \prime}$ restricted to that subpath. Due to the choice of $e^{\prime \prime}$, the depth of our recursion is bounded by $D=\log n$. In particular, at the end each edge $e$ is in the support of at most $O(D)$ guessed profiles. By Lemma 5 , the total error of these profiles is at most $O(\delta) u(e)$ which is compensated by the resource augmentation. Due to the choice of $e^{\prime}$, one can show that in each recursive call the profile of the parameter is a step function with only $4 \Gamma=O_{\delta, \epsilon}(\log \log n)$ steps: each recursive call "inherits" only half of the steps that were given to its parent subproblem, and additionally up to $2 \Gamma$ new steps due to the guessed profiles in the parent subproblem. This allows us to embed this recursion into a dynamic program (DP) with only $n^{O_{\delta, \epsilon}(\log \log n)}$ DP-cells as follows, and hence we obtain an overall running time of $n^{O_{\delta, \epsilon}(\log \log n)}$.

### 2.2 Dynamic program

We define the mentioned DP formally. In our DP table, we have a cell $(P, \hat{u})$ for each subpath $P \subseteq E$ and each function $\hat{u}: P \rightarrow \mathbb{N}_{0}$ such that $0 \leq \hat{u}(f) \leq n /(\delta \eta)$ for each edge $f \in P$ and $\hat{u}$ is a step-function with at most $4 \Gamma$ steps, i.e., one can partition $P$ into $4 \Gamma$ subpaths such that $\hat{u}$ is constant on each one of them. The intuitive meaning of $\hat{u}$ is that for each edge $e \in P$,
$\hat{u}(e)$ units of capacity have already been assigned to some tasks. The goal is to compute a set of tasks with maximum total weight and with $P(i) \subseteq P$ for each selected task $i$, such that on each edge $e \in P$ the selected tasks use a total capacity of at most $u(e)-\hat{u}(e)$.

- Lemma 7. There are at most $n^{O_{\delta, \epsilon}(\log \log n)}$ many different DP cells.

Proof. There are $O\left(n^{2}\right)$ possible choices for $P$. To determine the number of possible profiles $\hat{u}$, recall that $\hat{u}$ has at most $4 \Gamma=O_{\delta, \epsilon}(\log \log n)$ many steps. There are $\binom{O(n)}{O_{\delta, \epsilon}(\log \log n)} \in$ $n^{O_{\delta, \epsilon}(\log \log n)}$ possibilities to partition $P$ into at most $O_{\delta, \epsilon}(\log \log n)$ subpaths. For each subpath, there are at most $n /(\delta \eta)$ possibilities for the value of $\hat{u}$ on this subpath, i.e., $(n /(\delta \eta))^{O_{\delta}(\log \log n)}$ possible combinations overall. Recall that $\eta=\delta^{2 / \varepsilon}$. Multiplying these two numbers gives the total number of profiles.

Given a DP cell $\mathrm{DP}(P, \hat{u})$. We identify two edges $e^{\prime}, e^{\prime \prime} \in P$ similarly as above, i.e., we select $e^{\prime} \in P$ such that if we restrict $\hat{u}$ to the edges of $P$ on the left of $e^{\prime}$ or on the right of $e^{\prime}$, then $\hat{u}$ is a step-function with at most $2 \Gamma$ steps. Also, we select $e^{\prime \prime}$ such that at most $\lfloor|P| / 2\rfloor$ edges of $P$ lie on the left of $e^{\prime \prime}$, and at most $\lfloor|P| / 2\rfloor$ edges of $P$ lie on the right of $e^{\prime \prime}$. Let $P_{1}, P_{2}, P_{3}$ denote the subpaths of $P$ induced by $e^{\prime}$ and $e^{\prime \prime}$, i.e., the components of $P \backslash\left\{e^{\prime}, e^{\prime \prime}\right\}$.

Let $\mathcal{T}$ be the set of all pairs $\left(\hat{u}^{\prime}, \hat{u}^{\prime \prime}\right)$ such that $\hat{u}^{\prime}: P \rightarrow \mathbb{N}_{0}$ and $\hat{u}^{\prime \prime}: P \rightarrow \mathbb{N}_{0}$ are stepfunctions with at most $\Gamma$ steps each, and on each edge $f \in P$ they use at most $(1+\delta) u(f)$ units of capacity, i.e., $\hat{u}(f)+\hat{u}^{\prime}(f)+\hat{u}^{\prime \prime}(f) \leq(1+\delta) u(f)$. We apply the algorithm due to Theorem 6 to the instance whose input tasks contain all tasks in $i \in T_{e^{\prime}}$ with $P(i) \subseteq P$ and with edge capacities given by $\hat{u}^{\prime}$. Similarly, we apply this algorithm to the instance whose input tasks contain all tasks in $i \in T_{e^{\prime \prime}} \backslash T_{e^{\prime}}$ with $P(i) \subseteq P$ and with edge capacities given by $\hat{u}^{\prime \prime}$. Denote by $\operatorname{alg}\left(\hat{u}^{\prime}\right)$ and $\operatorname{alg}\left(\hat{u}^{\prime \prime}\right)$ the respective solutions. With the pair $\left(\hat{u}^{\prime}, \hat{u}^{\prime \prime}\right)$ we associate the solution given by $\operatorname{alg}\left(\hat{u}^{\prime}\right) \cup \operatorname{alg}\left(\hat{u}^{\prime \prime}\right)$ and the solutions stored in the cells $D P\left(P_{j},\left(\hat{u}+\hat{u}^{\prime}+\hat{u}^{\prime \prime}\right)_{\mid P_{j}}\right)$ for $j \in\{1,2,3\}$. We store in $\operatorname{DP}(P, \hat{u})$ the weight of the most profitable solution for all pairs $\left(\hat{u}^{\prime}, \hat{u}^{\prime \prime}\right) \in \mathcal{T}$, i.e.,

$$
\mathrm{DP}(P, \hat{u}):=\max _{t=\left(\hat{u}^{\prime}, \hat{u}^{\prime \prime}\right) \in \mathcal{T}}\left(\left(\operatorname{alg}\left(\hat{u}^{\prime}\right)\right)+w\left(\operatorname{alg}\left(\hat{u}^{\prime \prime}\right)\right)+\sum_{j=1}^{3} \mathrm{DP}\left(P_{j},\left(\hat{u}+\hat{u}^{\prime}+\hat{u}^{\prime \prime}\right)_{\mid P_{j}}\right)\right)
$$

Observe that $\left(\hat{u}+\hat{u}^{\prime}+\hat{u}^{\prime \prime}\right)_{\mid P_{j}}$ has at most $4 \Gamma$ steps as required for valid DP cells: the profile $\hat{u}$ restricted to $P_{j}$ has at most $2 \Gamma$ steps due to the halving and both $\hat{u}^{\prime}$ and $\hat{u}^{\prime \prime}$ introduce at most $\Gamma$ new steps each.

At the end we output $\operatorname{DP}\left(E, u_{0}\right)$ where $u_{0}$ is the profile with $u_{0}(e)=0$ for all $e \in E$. By standard memoization we can also output the solution associated with $\operatorname{DP}\left(E, u_{0}\right)$.

We need to show that the DP computes a near-optimal solution. It is clear that we output a feasible solution since in the computation for each cell, the set $\mathcal{T}$ contains only pairs ( $\hat{u}^{\prime}, \hat{u}^{\prime \prime}$ ) such that $\hat{u}(f)+\hat{u}^{\prime}(f)+\hat{u}^{\prime \prime}(f) \leq(1+\delta) u(f)$ for each edge $f \in P$. Intuitively, Lemma 5 allows us to argue that in each step we can guess for $\hat{u}^{\prime}$ and $\hat{u}^{\prime \prime}$ the profiles $\mathrm{mp}_{\mathrm{OPT}^{\prime} \cap T_{e^{\prime}}}$ and $\operatorname{mp}_{\left(\mathrm{OPT}^{\prime} \cap T_{e^{\prime \prime}}\right) \backslash T_{e^{\prime}}}$, respectively, where $\mathrm{OPT}^{\prime}$ denotes the tasks $i \in \mathrm{OPT}$ with $P(i) \subseteq \stackrel{e}{P}$, without violating the capacities of the edges.

- Lemma 8. The dynamic program computes a solution ALG which is feasible with $O(\delta)$ resource augmentation and $w(\mathrm{ALG}) \geq(1-\epsilon) w(\mathrm{OPT})$.

Proof. Consider a DP-cell ( $P, \hat{u}$ ) for which the DP defined the edges $e^{\prime}, e^{\prime \prime}$ when calculating its solution. Consider the profiles $\mathrm{mp}_{\mathrm{OPT} \cap\left\{i \in T_{e^{\prime}} \mid P(i) \subseteq P\right\}}$ and $\mathrm{mp}_{\mathrm{OPT} \cap\left\{i \in T_{e^{\prime \prime}} \backslash T_{e^{\prime}} \mid P(i) \subseteq P\right\} \text {. Observe }}$ that all tasks $\left\{i \in \mathrm{OPT} \cap\left(T_{e^{\prime}} \cup T_{e^{\prime \prime}}\right) \mid P(i) \subseteq P\right\}$ fit into the combined profile and all remaining tasks from OPT within $P$ will be considered in subproblems. If the two profiles are a feasible choice, then the DP obtains a profit of at least $(1-\epsilon) w\left(\left\{i \in \mathrm{OPT} \cap\left(T_{e^{\prime}} \cup T_{e^{\prime \prime}}\right) \mid P(i) \subseteq P\right\}\right)$
from the tasks in $T_{e^{\prime}} \cup T_{e^{\prime \prime}}$ when choosing these profiles. Therefore, it suffices to proof the claim that if in each recursive call up to some recursion level $\ell$ the DP chooses exactly these profiles, then in each subsequent recursive call of level $\ell+1$ the corresponding profiles will be a feasible choice again. Inductively, we can conclude then that the DP obtains a profit of at least $(1-\epsilon)$ opt.

To prove the mentioned claim, intuitively, due to the halving of the paths, the recursion depth is bounded from above by $O(\log n)$ and therefore the total error does not exceed the available amount of resource augmentation. Formally, we inductively maintain the invariant that for a DP cell $(P, \hat{u})$, if the profile $\hat{u}$ is composed of $2 k$ profiles (from previous DP cells), then the length of $P$ is bounded from above by $2 n / 2^{k}$. Recall that we assumed that $|E| \leq O(n)$. For each edge, we therefore have that at most $O(\log n)$ profiles overlap and by Lemma 5 , the total error of is at most $O(\delta) u(e)$ for each edge $e$, i.e., $O(\delta)$-resource augmentation is sufficient.

Together with Lemma 3 we conclude that we obtain a $(1+\epsilon)$-approximation algorithm for UFP with $\delta$-resource augmentation.

## 3 Bonus-UFP

In this section we define formally the Bonus-UFP problem (BUFP) and formalize how an algorithm with resource augmentation for Bonus-UFP yields an algorithm for (ordinary) UFP without resource augmentation.

### 3.1 Problem definition

Like in (ordinary) UFP, in the Bonus-UFP problem we are given as input a path $G=(V, E)$ and a set of tasks $T$ where each task $i \in T$ has a demand $d(i) \in \mathbb{N}$, a sub-path $P(i)$ of $E$, and a weight $w(i) \in \mathbb{N}_{0}$. The profit of a computed solution with task set $T^{\prime} \subseteq T$ stems as usual from $w\left(T^{\prime}\right)$ and, additionally, from some bonus profit that we obtain from subpaths $E^{\prime} \subseteq E$ (i.e., $E^{\prime}$ forms the edges of the respective subpath) such that no task $i \in T^{\prime}$ satisfies that $P(i) \subseteq E^{\prime}$ (but it might be that $P(i) \cap E^{\prime} \neq \emptyset$ ). The amount of bonus of such a subpath $E^{\prime}$ depends on the tasks $i \in T^{\prime}$ with $P(i) \cap E^{\prime} \neq \emptyset$. The reader may imagine that we get more bonus from $E^{\prime}$ if fewer tasks from $T^{\prime}$ intersect $E^{\prime}$. In particular, the computed solution consists of $T^{\prime}$ and of the subpaths $E^{\prime}$ from which the respective bonus profit is collected. Additionally, the input tasks are divided into large and small tasks, and on any edge we are allowed to select only $O(1)$ large tasks (in particular, for each edge $e$ the large tasks in $O P T \cap T_{e}$ can be guessed easily in time $\left.n^{O(1)}\right)$. Also, for technical reasons, the selected small tasks are allowed to use at most $2 b$ units of capacity of each edge $e \in E$ for some given value $b$ (see Figure 2).

Formally, in the input we are given a constant $\tau \in \mathbb{N}$ and values $\mu \in(0,1), b>0$ which partition the tasks $T$ into a set of large tasks $T_{L}=\{i \in T: d(i)>\mu b\}$ and small tasks $T_{S}=\{i \in T: d(i) \leq \mu b\}$. For each pair of a subpath $E^{\prime} \subseteq E$ and a set $L_{\text {int }}^{\prime} \subseteq T_{L}$ such that - the path of each task $i \in L_{\text {int }}^{\prime}$ uses the leftmost or the rightmost edge of $E^{\prime}$ and - $\left|L_{\text {int }}^{\prime} \cap T_{e}\right| \leq \tau$ for each edge $e \in E^{\prime}$
we are given a possible bonus $b n\left(E^{\prime}, L_{\text {int }}^{\prime}\right) \geq 0$. As described above, we obtain the bonus $b n\left(E^{\prime}, L_{\text {int }}^{\prime}\right)$ if no selected task (large or small) is contained in $E^{\prime}$ and the tasks in $L_{\text {int }}^{\prime}$ are exactly the selected large tasks that use some edge of $E^{\prime}$ (the amount of bonus does not depend on the selected small tasks $i$ with $\left.P(i) \cap E^{\prime} \neq \emptyset\right)$. A feasible solution to an instance of BUFP consists of a set of tasks $R \subseteq T_{L} \cup T_{S}$ and a collection of node-disjoint subpaths $E_{1}, \ldots E_{q}$ of $G$ such that the following holds:


Figure 2 A Bonus-UFP instance with two bonus intervals $E^{\prime}$ and $E^{\prime \prime}$. The hatched task within $E^{\prime \prime}$ is not allowed. The bonus depends on the intervals and the tasks depicted in dark gray - the large tasks that intersect with an interval.

1. $R$ is a feasible UFP solution, i.e., for every $e \in E$ it holds that $d\left(R \cap T_{e}\right) \leq u(e)$,
2. the small tasks in $R$ use at most a capacity of $2 b$ on each edge $e \in E$, i.e., $d\left(R \cap T_{S} \cap T_{e}\right) \leq$ $2 b$,
3. each edge $e \in E$ is used by at most $\tau$ tasks in $R \cap T_{L}$, i.e., $\left|R \cap T_{L} \cap T_{e}\right| \leq \tau$, and
4. for no $i \in R$ the path $P(i)$ is contained in any $E_{j}$ (but possibly $P(i) \cap E_{j} \neq \emptyset$ ).

The profit of the solution is $w(R)$ plus the bonus for each subpath $E_{j}$, where the latter depends on the large tasks in $R$ that use $E_{j}$, i.e., the total profit is $w(R)+\sum_{j=1}^{q} b n\left(E_{j}, L_{i n t, j}\right)$ where $L_{i n t, j}=\left\{i \in R \cap T_{L} \mid P(i) \cap E_{j} \neq \emptyset\right\}$.

We say that a solution to the above problem is feasible under $\delta$-resource augmentation if we relax the first two conditions to
1.' for every $e \in E$ it holds that $d\left(R \cap T_{e}\right) \leq u(e)+\delta b$,
2.' for every $e \in E$ it holds that $d\left(R \cap T_{S} \cap T_{e}\right) \leq 2 b+\delta b$.

So intuitively this increases the capacity for the small tasks by at least a factor $1+O(\delta)$. Also, note that Property 3 ensures that each edge $e$ is used by at most $\tau=O(1)$ large tasks, and hence we can guess the large tasks for each edge in time $n^{O(\tau)}=n^{O(1)}$.

For non-negative values $\alpha, \beta, \gamma \leq 1$ we say that a tri-criteria $(\alpha, \beta, \gamma)$-approximation algorithm for BUFP with $\delta$-resource augmentation is an algorithm that computes a solution that is feasible under $\delta$-resource augmentation of profit at least $\alpha \mathrm{opt}_{S}+\beta \mathrm{opt}_{L}+\gamma \mathrm{opt}_{B}$, assuming that there exists an (optimal) solution $R^{*}$ with bonuses $\left\{b n\left(E_{j}^{*}, L_{i n t, j}^{*}\right)\right\}_{j}$ such that opt $_{S}=w\left(R^{*} \cap T_{S}\right)$, opt $_{L}=w\left(R^{*} \cap T_{L}\right)$, and opt ${ }_{B}=\sum_{j=1}^{q} b n\left(E_{j}^{*}, L_{i n t, j}^{*}\right)$. Then our main result states that if $\gamma=1$ we can translate such an algorithm to an algorithm for normal UFP whose approximation ratio is $1 / \min \{\alpha, \beta\}$. We will prove the following theorem in Section 3.2.

- Theorem 9 (Black-box reduction). Assume that there is a ( $\alpha, \beta, 1$ )-approximation algorithm for BUFP with $\delta$-resource augmentation with a running time of $T_{\tau, \delta}\left(n, u_{\max }\right)$. Then there is $a \frac{1+\epsilon}{\min \{\alpha, \beta\}}$-approximation algorithm for UFP (without resource augmentation) with a running time of $T_{O_{\epsilon}(1), O_{\epsilon}(1)}\left(n, u_{\max }\right) \cdot\left(n \cdot \log u_{\max }\right)^{O_{\epsilon}(1)}$.


### 3.2 Black-box reduction

In this section we explain the key ideas for the black-box reduction due to Theorem 9 (omitting details due to space constraints). We start with a lemma indicating that there are near-optimal solutions in which each edge has a certain amount of slack. This slack is a
constant fraction of the capacity used by small tasks according to a suitable definition of small and large tasks. To avoid confusion, we refer to large and small tasks as in definition of BUFP as bonus-large and bonus-small, respectively.

In the lemma, we assign a level $\ell(e)$ to each edge $e$. We say that a task $i$ is of level $\ell$ if $P(i)$ includes an edge $e$ with $\ell(e)=\ell$ and no edge $e^{\prime}$ with $\ell\left(e^{\prime}\right)<\ell$; let $T^{(\ell)} \subseteq T$ denote all tasks of level $\ell$. Intuitively, each edge $e$ of level $\ell$ has an amount of slack $\varepsilon^{4} a(\ell)$ for some value $a(\ell)$ corresponding to level $\ell$. We define small and large tasks such that on each edge $e$ of level $\ell$, the small tasks of level $\ell$ use a total capacity of at most $2 b(\ell)$ with $b(\ell)=O_{\epsilon}(a(\ell))$ and each edge $e$ of level $\ell$ is used by at most $\tau=O_{\epsilon}(1)$ large tasks $i$ with $d(i)=\Omega(a(\ell))$. Formally, for an offset $h^{\prime} \in\{0, \ldots, 1 / \epsilon-1\}$ and a value $h=\Theta\left(\frac{1}{\varepsilon} \ln \frac{7}{\varepsilon^{2}}\right)$ defined in the next lemma, we set $a(\ell):=(1+\varepsilon)^{h^{\prime}+\ell h \cdot\left(1+\frac{1}{\varepsilon}\right)-\frac{h}{\varepsilon}}$ and $b(\ell)=a(\ell) \cdot(1+\varepsilon)^{\frac{h}{\varepsilon}}$. Let opt denote the weight of the optimal solution to the given instance.

- Lemma 10 (Slack Lemma). Let $\epsilon>0$. There are two constants $\mu_{1}, \mu_{2} \in\left(0, \varepsilon^{4}\right)$ with $\mu_{1}<\mu_{2} /(1+\varepsilon)^{1 / \varepsilon^{3}}$, values $h=\Theta\left(\frac{1}{\varepsilon} \ln \frac{7}{\varepsilon^{2}}\right)$ and $\tau=O_{\varepsilon}(1)$, an offset $h^{\prime}$ contained in a set of size $O_{\epsilon}(1)$ that can be computed in polynomial time, a near-optimal solution OPT with $w(\mathrm{OPT}) \geq(1-O(\epsilon)) \mathrm{opt}$, and a level $\ell(e)$ for each edge $e \in E$ with the following properties. We define
- $T_{L}:=\left\{i \in T: d(i) \geq \mu_{2} \cdot a(\ell(e))\right.$ for some edge $\left.e \in P(i)\right\} \quad$ (large tasks)
- $T_{S}:=\left\{i \in T: d(i)<\mu_{1} \cdot b(\ell(e))\right.$ for every edge $\left.e \in P(i)\right\} \quad$ (small tasks)
- $O P T_{L}=O P T \cap T_{L}$ and $O P T_{S}^{(\ell)}:=\mathrm{OPT} \cap T_{S} \cap T^{(\ell)}$ for each $\ell$.

Then, for each edge $e$ of level $\ell(e)=\ell$, it holds that:

- $d\left(T_{e} \cap O P T_{L}\right)+d\left(T_{e} \cap O P T_{S}^{(\ell)}\right) \leq u(e)-\varepsilon^{4} a(\ell)$,
- $d\left(T_{e} \cap O P T_{S}^{(\ell)}\right) \leq 2 b(\ell)$,
- $e$ is used by at most $\tau$ tasks $i \in \mathrm{OPT}_{L} \cap T^{(\ell)}$ with $d(i) \geq \mu_{2} \cdot a(\ell)$.

Our goal is to define an algorithm $B B$ that solves UFP given an oracle $O R$ for BUFP with $\delta$-resource augmentation. We will use the notation $G^{B}, T^{B}, T_{S}^{B}, T_{L}^{B}, u^{B}, \tau^{B}$ for the input path, input tasks, edge capacities, and threshold $\tau$ of the constructed instances of BUFP, while we use $G, T, u$ to denote the corresponding values of the given instance of UFP.

Intuitively, $B B$ solves a subproblem of BUFP for each maximally long path $G^{B}$ in which each edge is of level at least $\ell$ for some given value $\ell \in\left\{0, \ldots, \ell_{\max }\right\}$ (where $\ell_{\text {max }}$ denotes the maximum level due to Lemma 10), and the bonus subpaths $E_{j}$ selected in an optimal solution will correspond to maximal subpaths of $G^{B}$ consisting of edges of level at least $\ell+1$. With this interpretation, the amount of received bonus on such a subpath $E_{j}$ is calculated via subproblems of BUFP defined on this subpath $E_{j}$, level $\ell+1$, and any possible set of $O_{\epsilon}(1)$ selected large tasks of level $\ell$ that use at least one edge of $E_{j}$ but which are not contained in $E_{j}$. The parameters of the calls to $O R$ for this a subproblem (i.e., path $G^{B}$ and level $\ell$ ) will be intuitively as follows. We set $b=b(\ell), \mu=\mu_{1}$, and $a=a(\ell)$. We place in $T^{B}$ the tasks in $T$ contained in $G^{B}$ whose demand is at most $\mu_{1} b(\ell)$ or at least $\mu_{2} a(\ell)$. This yields that $T_{L}^{B}=\left\{i \in T^{B} \mid d(i) \geq \mu_{2} a(\ell)\right\}$ and $T_{S}^{B}=\left\{i \in T^{B} \mid d(i) \leq \mu_{1} b(\ell)\right\}$ according to the definition of BUFP. We remark that either $\mu_{2} a(\ell)>\mu_{1} b(\ell)$ or $\mu_{1} b(\ell)<1$ (in which case $T_{L}^{B}=\emptyset$ ); hence $T_{L}^{B}$ and $T_{S}^{B}$ are distinct. This way we will enforce that $T_{L}^{B} \subseteq T_{L}$ and $T_{S}^{B} \subseteq T_{S}$ for the sets $T_{L}, T_{S}$ due to Lemma 10 . We will set $\tau^{B}$ to be the respective value $\tau$ due to Lemma 10.

Note that in this recursive call there are $O_{\epsilon}(1)$ (previously selected) large tasks that use some but possibly not all edges of $E_{j}$. However, in the definition of Bonus-UFP, we have a global upper bound of $\tau^{B}$ for the number of allowed large tasks using an edge. To this end, we guess some further large tasks whose paths are contained in $E_{j}$, profiles for some of the
small tasks, and a partition of $E_{j}$ into $O_{\epsilon}(1)$ subpaths, such that we split this problem into subproblems in which each edge can be used by the same number $\tau^{B} \leq \tau$ of large tasks, so that we can call $O R$ on the resulting subproblem. We denote by opt ${ }_{S}$ and $\mathrm{opt}_{L}$ the profit due to small and large tasks, resp., in the solution OPT from Lemma 10.

- Theorem 11. Given a $T_{\tau, \delta}\left(n, u_{\max }\right)$ time $(\alpha, \beta, 1)$-approximation algorithm $O R$ for BUFP with $\delta$-resource augmentation, for every given constant $\delta>0$. Then, for every constant $\varepsilon>0$, there exists a $T_{O_{\epsilon}(1), O_{\epsilon}(1)}\left(n, u_{\max }\right) \cdot\left(n \log u_{\max }\right)^{O_{\epsilon}(1)}$ time algorithm $B B$ for UFP that computes a solution of profit at least $(1-\epsilon) \alpha \mathrm{opt}_{S}+(1-\epsilon) \beta \mathrm{opt}_{L}$.

Theorem 9 then follows from Lemma 10 and Theorem 11, where we set $\alpha=1-\varepsilon$ and $\beta=1$.
The reader might wonder why we proved the above theorem for general $\alpha$ and $\beta$. The reason is that it turns out that tasks $i \in O P T_{S}=O P T \cap T_{S}$ from Lemma 10 satisfy $d(i) \leq \epsilon^{2} u(e)$ for each $e \in P(i)$. Therefore, we can use LP-rounding techniques as in [11] to compute an alternative UFP solution with profit at least $(1-O(\epsilon)) \operatorname{opt}_{S}$. Hence any algorithm for BUFP as in Theorem 11 implies a $\left(\frac{1+\beta-\alpha}{\beta}+O(\varepsilon)\right)$-approximation for UFP. Recent work showed how to obtain ( $\alpha, \beta$ )-approximation algorithms for UFP in the above sense with $\beta=1$ and $\alpha$ equal to $\frac{1}{3}[17]$. This result becomes substantially easier to prove with resource augmentation in the BUFP setting. So we hope that Theorem 11 as stated can be a handy tool for future work along the same line.

## 4 Algorithm for Bonus-UFP under resource augmentation

We describe now how to adjust our algorithm for ordinary UFP with resource augmentation from Section 2 to an algorithm for Bonus-UFP with resource augmentation. Intuitively, there are two extra issues than one needs to address. First, the edges $e^{\prime}$ and $e^{\prime \prime}$ that partition the considered subpath $P$ might happen to fall within the interval of some bonus. This has to be taken into account in the definition of the subproblems which are solved recursively. Second, we need to keep track of the large tasks that use each edge so that the threshold $\tau$ is not exceeded. Therefore it is convenient to guess them explicitly, and use the over-estimated profiles for the small tasks only.

- Theorem 12. There is a ( $1-\epsilon, 1,1$ )-approximation algorithm for Bonus-UFP with $\delta$ resource augmentation with a running time of $n^{O_{\varepsilon, \delta}(\log \log n)}$.

Theorems 9 and 12 together yield Theorem 2. It remains a challenging open problem to construct a PTAS for (Bonus-)UFP. One key bottleneck in our approach is that already in the first iteration we need to guess a profile with up to $\Omega(\log \log n)$ many steps, and it is not clear how to do this in polynomial time.

In the remaining section, we prove Theorem 12. Again, we start with normalizing the instance. Since we have an additive resource augmentation of $\delta b$, we can use half of it in order to ensure $u_{\text {min }} \geq \delta b / 2$. Thus $2 b \leq 4 u_{\min } / \delta$, i.e., $u_{\mathrm{OPT}_{S}}(e)$ together with half of the resource augmentation has a constant capacity profile. Furthermore, analogous to Lemma 3, we can assume that each task $i \in T$ has an integer demand $d(i) \in \mathbb{N}$ and $b \in O_{\delta}(n)$. Unlike before, we do not lose a profit of $\epsilon w\left(\mathrm{OPT}_{L}\right)$.

- Lemma 13. For arbitrary $\delta>0$, suppose there is an ( $\alpha, \beta, \gamma$ )-approximation algorithm for BUFP instances $I$ with $(1+\delta)$ resource augmentation such that $u_{\min }=\delta b, b \in O_{\delta}(n)$, and $d(i) \in \mathbb{N}$ for all $i \in T$. Then there is an $((1-\epsilon) \alpha, \beta, \gamma)$-approximation algorithm for BUFP with $(1+4 \delta)$ resource augmentation.

Due to Lemma 13, in the following we assume that we are given an instance of BUFP with $(1+\delta)$ resource augmentation such that $u_{\text {min }}=\delta b, 2 b=n / \eta$ for some $\eta \in O_{\delta}(1)$, and $d(i) \in \mathbb{N}$ for all $i \in T$.

We adapt the strategy used for UFP with resource augmentation and we use the same notation. Again we overestimate profiles based on the value $\theta:=\left\lfloor u_{\text {min }} \cdot \delta / D\right\rfloor$, now with $u_{\min }=\delta b$ and $D=\log n$. Lemma 4 is still valid for BUFP, if we restrict the set $S$ to small tasks.

- Lemma 14. There is a value $\Gamma=O_{\delta, \epsilon}(\log D)=O_{\delta, \epsilon}(\log \log n)$ such that for every edge $\bar{e} \in E$ and every set $S \subseteq T_{S} \cap T_{\bar{e}}$ with $d\left(S \cap T_{e}\right) \leq 2 b$ for each $e \in E$, we have that $m p_{S}(\cdot)$ is a step-function with at most $\Gamma$ steps.

Proof. Since the demand $d(S)$ is bounded from above by $2 b$, it is sufficient to consider the profile that for each edge $e \in E$ has the capacity $\min \{2 b, u(e)\}$. With the modification, we can apply Lemma 4.

Furthermore, Lemma 5 does not change for BUFP, if we restrict it to small tasks, i.e., $S_{j} \subseteq T_{S}$ for all $j$.

For ease of notation, we assume that for each edge $f$ and each set $L \subseteq T_{L} \cap T_{e}$ with $d(L) \leq u(f)$, the BUFP instance to be solved has is a (dummy) bonus interval ( $\{f\}, L$ ). If the bonus interval is not part of the original instance, it has zero bonus, i.e., $\operatorname{bn}(\{f\}, L)=0$. By duplicating edges, we can assume without loss of generality that there is no task $i$ with $P(i)=f$.

### 4.1 Dynamic program for BUFP

In our DP table, we have a cell $(P, \hat{u}, L)$ for each subpath $P \subseteq E$, each function $\hat{u}: P \rightarrow \mathbb{N}_{0}$, and each set of large tasks $L \subseteq T_{L}$ with the following properties. We require $0 \leq \hat{u}(f) \leq n / \eta$ for each edge $f \in P$ and $\hat{u}$ is a step-function with at most $8 \Gamma$ steps, i.e., one can partition $P$ into $8 \Gamma$ subpaths such that $\hat{u}$ is constant on each of them. A set $L$ is valid, if each task in $L$ uses the leftmost or the rightmost edge of $P$. Note that $|L| \leq 2 \tau$ since there can only be at most $\tau$ large tasks crossing each of the two boundaries of $P$.

The intuitive meaning of $\hat{u}$ and $L$ is that for each edge $e \in P, \hat{u}(e)$ units of capacity have already been assigned to some small tasks and the tasks $L$ have already been selected. For a set of tasks $T^{\prime}$, we define $u_{T^{\prime}}: E \rightarrow \mathbb{N}_{0}$ to be the profile with $u_{T}(e)=d\left(T^{\prime} \cap T_{e}\right)$ for each $e \in E$. The goal is to compute a set of tasks with maximum total weight and with $P(i) \subseteq P$ for each selected task $i$, such that on each edge $e \in P$ the selected tasks use a total capacity of at most $u(e)-\hat{u}(e)-u_{L}(e)$.

- Lemma 15. There are at most $n^{O_{\delta, \epsilon, \tau}(\log \log n)}$ many different DP cells.

Proof. There are $O\left(n^{2}\right)$ possible choices for $P$. To determine the number of possible profiles $\hat{u}$, recall that $\hat{u}$ has at most $8 \Gamma=O_{\delta, \epsilon}(\log \log n)$ many steps and $L$ has a size of at most $2 \tau$.

There are $\left(\underset{O_{\delta, \epsilon, \tau}(\log \log n)}{O(n)}\right) \in n^{O_{\delta, \epsilon, \tau}(\log \log n)}$ possibilities to partition $P$ into at most $O_{\delta, \epsilon, \tau}(\log \log n)$ subpaths. For each subpath, there are at most $n / \eta$ possibilities for the value of $\hat{u}$ on this subpath and $\binom{n}{2 \tau} \leq n^{2 \tau}$ possibilities for selecting $L$, i.e., $(n / \eta)^{O_{\delta, \epsilon}(\log \log n)} \cdot n^{2 \tau}$ possible combinations overall. Multiplying these two numbers gives the total number of profiles.

Suppose that we are given a DP cell $\mathrm{DP}(P, \hat{u}, L)$. We identify two edges $e^{\prime}, e^{\prime \prime} \in P$. We select $e^{\prime} \in P$ such that if we restrict $\hat{u}$ to the edges of $P$ on the left of $e^{\prime}$ or on the right of $e^{\prime}$, then $\hat{u}$ is a step-function with at most $4 \Gamma$ steps. Also, we select $e^{\prime \prime}$ such that at most
$\lfloor|P| / 2\rfloor$ edges of $P$ lie on the left of $e^{\prime \prime}$, and at most $\lfloor|P| / 2\rfloor$ edges of $P$ lie on the right of $e^{\prime \prime}$. Let $P_{1}, P_{2}, P_{3}$ denote the subpaths of $P$ induced by $E_{j^{\prime}}$ and $E_{j^{\prime \prime}}$, i.e., the components of $P \backslash\left(E_{j^{\prime}} \cup E_{j^{\prime \prime}}\right)$.

Let $\mathcal{T}$ be the set of all tuples $\left(p_{1}, p_{2}, \hat{u}^{\prime}, \hat{u}^{\prime \prime}\right)$ with $p_{1}:=\left(E_{j^{\prime}}, L_{\mathrm{int}, j^{\prime}}\right)$ and $p_{2}:=\left(E_{j^{\prime \prime}}, L_{\mathrm{int}, j^{\prime \prime}}\right)$ specified as follows. The pair $\left(E_{j^{\prime}}, L_{\mathrm{int}, j^{\prime}}\right)$ is a bonus interval with $e^{\prime} \in E_{j^{\prime}}$ and $\left(E_{j^{\prime \prime}}, L_{\mathrm{int}, j^{\prime \prime}}\right)$ is a bonus interval with $e^{\prime \prime} \in E_{j^{\prime \prime}}$ such that $E_{j^{\prime}} \cap E_{j^{\prime \prime}}=\emptyset$ or $p_{1}=p_{2}$. A task $i$ which overlaps with both $E_{j^{\prime}}$ and $E_{j^{\prime \prime}}$ (i.e., $P(i) \cap E_{j^{\prime}} \neq \emptyset$ and $P(i) \cap E_{j^{\prime \prime}} \neq \emptyset$ ) is either in both $L_{\mathrm{int}, j^{\prime}}$ and $L_{\text {int }, j^{\prime \prime}}$ or in none.

Each of the remaining two entries of the tuple $\hat{u}^{\prime}: P \rightarrow \mathbb{N}_{0}$ and $\hat{u}^{\prime \prime}: P \rightarrow \mathbb{N}_{0}$ is composed of two step-functions with at most $\Gamma$ steps each, and on each edge $f \in P, \hat{u}^{\prime}(f)+\hat{u}^{\prime \prime}(f) \leq(2+\delta) b$. Furthermore, the capacity of all step-functions use at most $u(f)+\delta b$ units of capacity, i.e., $\hat{u}(f)+\hat{u}^{\prime}(f)+\hat{u}^{\prime \prime}(f)+u_{L}(f)+u_{L_{\mathrm{int}, j^{\prime}}}(f)+u_{L_{\mathrm{int}, j^{\prime \prime}}}(f) \leq u(f)+\delta b$. Let $e_{\ell}^{\prime}, e_{r}^{\prime}, e_{\ell}^{\prime \prime}, e_{r}^{\prime \prime}$ be the left-most and right-most edge of $E_{j^{\prime}}$ and $E_{j^{\prime \prime}}$, respectively. We apply the algorithm due to Theorem 6 to the instance whose input tasks contain all tasks in $i \in T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime}}$ with $P(i) \subseteq P$ and $P(i) \nsubseteq E_{j^{\prime}}$, with edge capacities given by $\hat{u}^{\prime}$. Similarly, we apply this algorithm to the instance whose input tasks contain all tasks in $i \in T_{e_{\ell}^{\prime \prime}} \cup T_{e_{r}^{\prime \prime}} \backslash\left(T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime}}\right)$ with $P(i) \subseteq P$ and $P(i) \nsubseteq E_{j^{\prime \prime}}$, with edge capacities given by $\hat{u}^{\prime \prime}$. Denote by $\operatorname{alg}\left(\hat{u}^{\prime}\right)$ and $\operatorname{alg}\left(\hat{u}^{\prime \prime}\right)$ the respective solutions. With the tuple ( $\left.p_{1}, p_{2}, \hat{u}^{\prime}, \hat{u}^{\prime \prime}\right)$ we associate the solution given by $\operatorname{alg}\left(\hat{u}^{\prime}\right) \cup \operatorname{alg}\left(\hat{u}^{\prime \prime}\right)$, the bonuses $\operatorname{bn}\left(p_{1}\right), \operatorname{bn}\left(p_{2}\right)$, the profit from large tasks $w\left(L_{\mathrm{int}, j^{\prime}} \cup L_{\mathrm{int}, j^{\prime \prime}} \backslash L\right)$, and the solutions stored in the cells $D P\left(P_{j},\left(\hat{u}+\hat{u}^{\prime}+\hat{u}^{\prime \prime}\right)_{\mid P_{j}}, L^{\prime}\right)$ for $j \in\{1,2,3\}$ and $L^{\prime}$ the tasks from $L \cup L_{\text {int }, j^{\prime}} \cup L_{\text {int, } j^{\prime \prime}}$ crossing $P_{j}$. We store in $\mathrm{DP}(P, \hat{u}, L)$ the weight of the most profitable solution for all tuples $\left(p_{1}, p_{2}, \hat{u}^{\prime}, \hat{u}^{\prime \prime}\right) \in \mathcal{T}$, i.e.,

$$
\begin{aligned}
\operatorname{DP}(P, \hat{u}, L) & :=\max _{t=\left(p_{1}, p_{2}, \hat{u}_{S}^{\prime}, \hat{u}_{S}^{\prime \prime}\right) \in \mathcal{T}}\left(w\left(\operatorname{alg}\left(\hat{u}^{\prime}\right)\right)+w\left(\operatorname{alg}\left(\hat{u}^{\prime \prime}\right)\right)\right. \\
& \left.+\sum_{p \in\left\{p_{1}, p_{2}\right\}} \operatorname{bn}(p)+w\left(L_{\mathrm{int}, j^{\prime}} \cup L_{\mathrm{int}, j^{\prime \prime}} \backslash L\right)+\sum_{j=1}^{3} \operatorname{DP}\left(P_{j},\left(\hat{u}+\hat{u}^{\prime}+\hat{u}^{\prime \prime}\right)_{\mid P_{j}}\right)\right),
\end{aligned}
$$

where $\hat{u}^{\prime}$ and $\hat{u}^{\prime \prime}$ are derived from $t$ as described before. Observe that if $p_{1}=p_{2}$, we collect only one bonus. We have to ensure that $\hat{u}^{(j)}:=\left(\hat{u}+\hat{u}^{\prime}+\hat{u}^{\prime \prime}\right)_{\mid P_{j}}$ is a valid profile. Due to the halving, restricted to $P_{j}$ the profile $\hat{u}$ has at most $4 \tau$ steps and each of the other two profiles has at most $2 \Gamma$ steps from small tasks, which results in $8 \Gamma$ steps in total. Independently, we always have a total number of at most $4 \tau$ steps from large tasks

At the end we output $\operatorname{DP}\left(E, u_{0}\right)$ where $u_{0}$ is the profile with $u_{0}(e)=0$ for all $e \in E$. By standard memoization we can also output the solution associated with $\operatorname{DP}\left(E, u_{0}\right)$.

We need to show that the DP computes a near-optimal solution. It is clear that we always output a feasible solution since in the computation for each cell, the set $\mathcal{T}$ contains only tuples $\left(\hat{u}_{S}^{\prime}, \hat{u}_{S}^{\prime \prime}, p_{1}, p_{2}\right)$ such that $\hat{u}(f)+\hat{u}^{\prime}(f)+\hat{u}^{\prime \prime}(f)+u_{L \cup L_{\text {int }, j^{\prime}} \cup L_{\text {int }, j^{\prime \prime}}} \leq u(f)+\delta b$ for each edge $f \in P$. Intuitively, Lemma 5 allows us to argue that in each step we can guess for $\hat{u}^{\prime}$ and $\hat{u}^{\prime \prime}$, the pairs $p_{1}$ and $p_{2}$, the profiles $\mathrm{mp}_{\mathrm{OPT}^{\prime} \cap T_{e^{\prime}}}$ and $\mathrm{mp}_{\left(\mathrm{OPT}^{\prime} \cap T_{e^{\prime \prime}}\right) \backslash T_{e^{\prime}}}$, respectively, where $\mathrm{OPT}^{\prime}$ denotes the tasks $i \in \mathrm{OPT}$ with $P(i) \subseteq P$, without violating the capacities of the edges. Let $\mathrm{OPT}_{B}$ be the set of bonus pairs of an optimal solution and $w\left(\mathrm{OPT}_{B}\right)$ the sum of bonuses.

- Lemma 16. The dynamic program computes a solution ALG which is feasible with $O(\delta)$ resource augmentation and $w(\mathrm{ALG}) \geq(1-\epsilon) w\left(\mathrm{OPT}_{S}\right)+w\left(\mathrm{OPT}_{L}\right)+w\left(\mathrm{OPT}_{B}\right)$.

Proof. Initially, the DP chooses the middle edge $e$ (i.e., $e=e^{\prime}=e^{\prime \prime}$ ) and a bonus interval containing $e$. One of the choices is the bonus pair $p=\left(E_{j}, L_{\mathrm{int}, j}\right)$ such that $p \in \mathrm{OPT}_{B}$ and $e \in E_{j}$. Since OPT is feasible, it has no tasks contained in $E_{j}$. Let $e_{\ell}$ and $e_{r}$ be the left-most and right-most edge of $E_{j}$. Then one of the options for the profile is to choose $\hat{u}:=\mathrm{mp}_{\mathrm{OPT} \cap T_{e_{\ell}}}+\mathrm{mp}_{\mathrm{OPT} \cap T_{e_{r}}}$.

All tasks OPT $\cap\left(T_{e_{\ell}} \cup T_{e_{r}}\right)$ fit into the profile $\hat{u}$. Since the DP approximates the profit from these tasks, the value of the DP cell without the values from the subproblems is at least $(1-\epsilon) w\left(\mathrm{OPT}_{S} \cap T_{e}\right)+w\left(\mathrm{OPT}_{L} \cap\left(T_{e_{\ell}} \cup T_{e_{r}}\right)\right)+\mathrm{bn}(p)$. The instance is split into the left hand side of $E_{j}$ and the right hand side of $E_{j}$. Let ( $P, \hat{u}, L$ ) be a sub-problem reached (recursively) by the described choice of profile. Each edge $e^{\prime}$ chosen subsequently within $P$ has the property that the DP can choose the bonus pair $p_{1}:=\left(E_{j^{\prime}}, L_{\mathrm{int}, j^{\prime}}\right)$ (where as before we assume that $e_{\ell}^{\prime}$ and $e_{r}^{\prime}$ are the leftmost and rightmost edges of $E_{j^{\prime}}$, respectively) from $\mathrm{OPT}_{B}$ containing $e^{\prime}$ and a feasible choice of $\hat{u}^{\prime}$ is $\operatorname{mp}_{\mathrm{OPT}_{S} \cap\left\{i \in T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime}} \mid P(i) \subseteq P\right\} \text {. Each edge } e^{\prime \prime} \text { chosen subsequently within }}$ $P$ has the property that the DP can choose the bonus pair $p_{2}:=\left(E_{j^{\prime \prime}}, L_{\mathrm{int}, j^{\prime \prime}}\right)$ (where as before we assume that $e_{\ell}^{\prime \prime}$ and $e_{r}^{\prime \prime}$ are the leftmost and rightmost edges of $E_{j^{\prime \prime}}$, respectively) from $\mathrm{OPT}_{B}$ containing $e^{\prime \prime}$ and a feasible choice of $\hat{u}^{\prime}$ is $\left.\operatorname{mp}_{\mathrm{OPT}_{S} \cap\left\{i \in T_{e_{\ell}^{\prime \prime}} \cup T_{e_{r}^{\prime \prime}} \backslash\left(T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime}}\right)\right.} \mid P(i) \subseteq P\right\}$. Observe that all tasks $\left\{i \in \mathrm{OPT}_{S} \cap\left(T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime}} \cup T_{e_{\ell}^{\prime \prime}} \cup T_{e_{r}^{\prime \prime}}\right) \mid P(i) \subset P\right\}$ fit into the combined profiles, we select all tasks specified in $p_{1}$ and $p_{2}$, and all remaining tasks from OPT within $P$ will be available in sub-problems. The DP collects a profit of at least $(1-\epsilon) w(\{i \in$ $\left.\left.\mathrm{OPT}_{S} \cap\left(T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime \prime}} \cup T_{e_{\ell}^{\prime \prime}} \cup T_{e_{r}^{\prime \prime}}\right) \mid P(i) \subseteq P\right\}\right)+w\left(\left\{i \in \mathrm{OPT}_{L} \cap\left(T_{e_{\ell}^{\prime}} \cup T_{e_{r}^{\prime \prime}} \cup T_{e_{\ell}^{\prime \prime}} \cup T_{e_{r}^{\prime \prime}}\right) \backslash L \mid\right.\right.$ $P(i) \subseteq P\})+\mathrm{bn}\left(p_{1}\right)+\mathrm{bn}\left(p_{2}\right)$. We inductively conclude that if all choices as described above are feasible, the DP obtains a profit of at least $(1-\epsilon) w\left(\mathrm{OPT}_{S}\right)+w\left(\mathrm{OPT}_{L}\right)+w\left(\mathrm{OPT}_{B}\right)$.

To show feasibility, we have to argue that none of the described choices exceeds the overall capacity of $u(f)+\delta b$ for an edge $f \in E$. The argument is analogous to the proof of Lemma 8.

Together with Lemma 13 we conclude that we obtain a ( $1-\epsilon, 1,1$ )-approximation algorithm for BUFP with $\delta$-resource augmentation.

## References

1 Anna Adamaszek, Parinya Chalermsook, Alina Ene, and Andreas Wiese. Submodular unsplittable flow on trees. Math. Program., 172(1-2):565-589, 2018. doi:10.1007/ s10107-017-1218-4.
2 Aris Anagnostopoulos, Fabrizio Grandoni, Stefano Leonardi, and Andreas Wiese. Constant integrality gap LP formulations of unsplittable flow on a path. In IPCO, pages 25-36, 2013. doi:10.1007/978-3-642-36694-9_3.
3 Aris Anagnostopoulos, Fabrizio Grandoni, Stefano Leonardi, and Andreas Wiese. A mazing $2+\varepsilon$ approximation for unsplittable flow on a path. ACM Trans. Algorithms, 14(4):55:1-55:23, 2018. doi:10.1145/3242769.

4 Y. Azar and O. Regev. Combinatorial algorithms for the unsplittable flow problem. Algorithmica, 44(1):49-66, 2006. doi:10.1007/s00453-005-1172-z.
5 N. Bansal, A. Chakrabarti, A. Epstein, and B. Schieber. A quasi-PTAS for unsplittable flow on line graphs. In STOC, pages 721-729. ACM, 2006.
6 N. Bansal, Z. Friggstad, R. Khandekar, and R. Salavatipour. A logarithmic approximation for unsplittable flow on line graphs. In $S O D A$, pages 702-709, 2009.
7 Jatin Batra, Naveen Garg, Amit Kumar, Tobias Mömke, and Andreas Wiese. New approximation schemes for unsplittable flow on a path. In SODA, pages 47-58, 2015. doi:10.1137/1.9781611973730.5.

8 Paul Bonsma, Jens Schulz, and Andreas Wiese. A constant-factor approximation algorithm for unsplittable flow on paths. SIAM Journal on Computing, 43:767-799, 2014.
9 A. Chakrabarti, C. Chekuri, A. Gupta, and A. Kumar. Approximation algorithms for the unsplittable flow problem. Algorithmica, 47:53-78, 2007.
10 C. Chekuri, A. Ene, and N. Korula. Unsplittable flow in paths and trees and column-restricted packing integer programs. In APPROX-RANDOM, pages 42-55, 2009.
11 C. Chekuri, M. Mydlarz, and F. Shepherd. Multicommodity demand flow in a tree and packing integer programs. ACM Transactions on Algorithms, 3, 2007.
12 M. Chrobak, G. Woeginger, K. Makino, and H. Xu. Caching is hard, even in the fault model. In $E S A$, pages 195-206, 2010.
13 Julia Chuzhoy and Alina Ene. On approximating maximum independent set of rectangles. In IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS 2016), pages 820-829. IEEE, 2016.
14 Fabrizio Grandoni, Salvatore Ingala, and Sumedha Uniyal. Improved approximation algorithms for unsplittable flow on a path with time windows. In WAOA, pages 13-24, 2015.
15 Fabrizio Grandoni, Tobias Mömke, and Andreas Wiese. Unsplittable flow on a path: The game!, 2020. unpublished.
16 Fabrizio Grandoni, Tobias Mömke, Andreas Wiese, and Hang Zhou. To augment or not to augment: Solving unsplittable flow on a path by creating slack. In SODA, pages 2411-2422, 2017.

17 Fabrizio Grandoni, Tobias Mömke, Andreas Wiese, and Hang Zhou. A (5/3+ 6 )-approximation for unsplittable flow on a path: placing small tasks into boxes. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, pages 607-619, 2018. doi:10.1145/3188745.3188894.
18 Shi Li. Towards PTAS for precedence constrained scheduling via combinatorial algorithms. In SODA, pages 2991-3010. SIAM, 2021.
19 Andreas Wiese. A (1+epsilon)-approximation for unsplittable flow on a path in fixed-parameter running time. In ICALP 2017, volume 80 of LIPIcs, pages 67:1-67:13, 2017. doi:10.4230/ LIPIcs.ICALP. 2017.67.

