Bachelor Thesis

Approximation Algorithms for Solving the Quality-of-Service Multicast Tree Problem

Marie Diana Sieper

Date of Submission: August 3, 2018
Advisors: Priv.-Doz. Dr. habil. Joachim Spoerhase
Prof. Dr. Alexander Wolff

Julius-Maximilians-Universität Würzburg
Lehrstuhl für Informatik I
Algorithmen, Komplexität und wissensbasierte Systeme

Abstract

The Quality-of-Service Multicast Tree problem represents a generalization of the Steiner tree problem. The input of this (NP-hard) Steiner tree problem is a weighted, undirected graph $G = (V, E)$ and a set of terminals $R \subseteq V$. The goal is to find a tree $T = (V', E')$ with minimal weight connecting all terminals.

The Quality-of-Service Multicast Tree (QoSMT) problem also expects a weighted, undirected graph $G = (V, E)$ with a set of terminals $R \subseteq V$, but now every terminal possesses a non-negative rate, and a source $s \in V$ that every terminal needs to be connected to is given. As in the Steiner tree problem, the solution is a tree $T = (V', E')$ connecting all terminals and the source s with minimal costs, but now the costs of every edge $e \in E'$ are its own weight multiplied with the highest rate of any node connected to the source s over a path containing e. Many names and formulations exist for the Quality-of-Service Multicast Tree (QoSMT) problem such as Multi-Level Steiner tree problem, Grade-of-Service Steiner tree problem or Multi-Tier Tree problem.

We first analyse the best currently known algorithm by Karpinski et al. [KMOZ05] solving the QoSMT problem, which provides a constant approximation guarantee. It relies on approximation algorithms for the Steiner tree problem with the special feature of them being β-convex α-approximation algorithms. We show that $\beta \geq 1$ always holds for every Steiner tree approximation algorithm, and that every α-approximation algorithm for the Steiner tree problem is a α-convex $(\alpha + \varepsilon)$-approximation algorithm for every $\varepsilon > 0$.

We analyse the algorithm of Byrka et al. which provides the lowest known approximation guarantee of $\ln(4) + \varepsilon$ for all $\varepsilon > 0$ for the Steiner tree problem with regard to the characteristic mentioned above, and combine it with the Karpinski-algorithms, obtaining new approximation guarantees of at least 3.769 for the QoSMT problem and at least 1.849 for the special case of the two-rate QoSMT problem, in which the terminals have exactly two different rates. We think and will also argue that for this Byrka-algorithm the value of β can not be significantly smaller than the value of α.

Furthermore we present two different component-based LPs for the two-rate QoSMT problem in order to use them to transfer the component-based approach of Byrka et al. [BGRS13] for it. The first one works with homogenous components (the costs of all edges belonging to one component are all multiplied with the same rate), and we show that an integer solution to this LP does provide a 2-approximation for the two-rate QoSMT problem, but there are graphs for which the computed solution is not optimal. The second LP, which is not restricted to homogenous components, computes an exact solution to the two-rate QoSMT problem in the integral case. We concentrate on the question whether QoSMTs can be approximated arbitrarily close with trees consisting of components each with at most k terminals. We present a theorem which says that

this question can be decided on a small subset of graphs. Furthermore we analyse the approach of Borchers and Du [BD97], who did prove a similar result for the Steiner tree problem, with regard to the question whether it can be adopted for QoSMTs, and show the occuring difficulties.

Zusammenfassung

Das Quality-of-Service Multicast Tree Problem stellt eine Verallgemeinerung des Steinerbaumproblems dar. Bei diesem NP-schweren Problem sind ein gewichteter, ungerichteter Graph $G = (V, E)$ und eine Terminalmenge $R \subseteq V$ mit dem Ziel gegeben, einen Baum $T = (V', E') \subseteq G$ zu finden, der alle Terminals verbindet und dabei minimale Kosten verursacht.

Für das Quality-of-Service Multicast Tree (QoSMT) Problem existieren viele Namen und Formulierungen, wie beispielsweise Multi-Level Steiner Tree Problem, Grade-of-Service Steiner Tree Problem oder Multi-Tier Tree Problem. Dabei ist ebenfalls ein ungerichteter, gewichteter Graph $G = (V, E)$, sowie einen Terminalmenge $R \subseteq V$ gegeben. Jedoch besitzt jetzt auch noch jedes Terminal einen Level, und es ist eine Quelle s gegeben, an die alle Terminale angeschlossen werden müssen. Gesucht ist nun wieder ein Baum $T = (V', E')$ mit minimalen Kosten, der alle Terminals und s miteinander verbindet. Jedoch berechnen sich nun die Kosten jeder einzelnen Kante $e \in E'$ aus ihrem Eigengewicht multipliziert mit dem höchsten Level eines Knotens, der durch e mit der Quelle s verbunden wird.

In dieser Arbeit untersuchen wir daher zunächst den besten zur Zeit bekannten Algorithmus für das QoSMT Problem von Karpinski et al. [KMOZ05], der eine konstante Approximationsschranke von 3,802 besitzt. Dieser benötigt Approximationsalgorithmen für das Steinerbaumproblem, die zusätzlich sogenannte β-konvexe α-Approximationsalgorithmen sind. Wir zeigen, dass $\beta \geq 1$ für jeden das Steinerbaumproblem lösenden Approximationsalgorithmus gelten muss, und dass jeder α-Approximationsalgorithmus für das Steinerbaumproblem auch ein α-konvexer $(\alpha+\varepsilon)$-Approximationsalgorithmus für jedes $\varepsilon > 0$ ist.

Wir untersuchen den Algorithmus von Byrka et al. [BGRS13], der die niedrigste bisher bekannte Approximationsschranke von $\ln(4) + \varepsilon$ für alle $\varepsilon > 0$ für das Steinerbaumproblem aufweist, auf diese oben genannte Eigenschaft und setzen ihn in den Karpinski-Algorithmus ein. Dabei erhalten wir neue Approximationsgüten von 3,769 für das allgemeine QoSMT Problem und 1,849 für den Spezialfall des zwei-Raten-QoSMT Problems, in dem jedes Terminal genau eine von zwei verschiedene Raten r_1 und r_2 besitzt. Wir zeigen außerdem, dass für den Byrka et-al.-Algorithmus β und α immer nahezu identisch sind.

Außerdem präsentieren wir zwei verschiedene komponentenbasierende LPs für das Zwei-Raten-QoSMT Problem, die dem Zweck dienen sollen, den komponentenbasierten

Ansatz von Byrka et al. [BGRS13] auf das Zwei-Raten-QoSMT Problem zu übertragen. Eines davon arbeitet mit homogenen Komponenten (die Kosten aller Kanten einer Komponente werden mit derselben Rate multipliziert), und wir zeigen, dass es sich bei ganzzahligen Ergebnissen dieses LPs um eine 2-Approximation, im Allgemeinen jedoch nicht um eine exakte Lösung für das Zwei-Raten-QoSMT Problem handelt. Das zweite LP, welches nicht nur auf homogene Komponenten beschränkt ist, liefert im ganzzahligen Fall einen exakten Wert. Wir beschäftigen uns mit der Fragestellung, ob wir QoSMTs mit Bäumen approximieren können, die aus Komponenten mit je höchstens k Terminals bestehen. Wir präsentieren einen Satz, der eine starke Einschränkung der Graphenklasse bietet, auf der wir diese Frage beantworten müssen. Weiterhin überprüfen wir den Ansatz von Borchers und Du [BD97], die ein ähnliches Resultat für das Steinerbaumproblem bewiesen haben, auf seine Eignung für diese Fragestellung, und zeigen Schwierigkeiten dabei auf.

Contents

1 Introduction

Motivation The Quality-of-Service Multicast Tree (QoSMT) problem appears in the design of low-cost network structures connecting a transmitter with its receivers. Usually, these receivers have a great variety in Quality-of-Service requests (e.g. video streaming with different resolutions), and sending data with higher quality is more expensive than sending data with lower quality. Therefore it would cost far too much to send the highest quality to every receiver. It is therefore advantageous to construct a network respecting this variety in Quality-of-Service requests. Apart from this application in network design, multi-scale representations of graphs can be used in applications such as geography or network visualization, where users want to examine complex (street, river, or abstract) networks at different levels of detail. Here, the cost models the stability of the visualization. [AAS+18].

Related work The QoSMT problem is based on the Steiner tree problem (definition 2.1), which is one of the most fundamental problems in Computer Science and Operations Research [BGRS13]. Solving it is NP-hard as was shown by Karp [Kar72]. In fact, Chlebík and Chlebíková [CC08] have proven that it is already NP-hard to approximate it with an approximation factor $\leq 96/95$.

A simple 2-approximation algorithm is the terminal-spanning-tree algorithm [GP68]. It is also a good example for an algorithm computing a so-called *k-restricted* tree (definition 2.3) for $k = 2$. Borchers and Du [BD97] have shown that k-restricted trees can be used to arbitrarily approximate Steiner trees (we will cite this result in theorem 2.4).

Several further algorithms working with k-restricted trees habe been published, culminating with an $1+\ln(3)/2+\varepsilon < 1.55$-approximation algorithm by Robins and Zelikovsky [RZ05], $\varepsilon > 0$ being arbitrary small.

For many years, this was the best known approximation algorithm for the Steiner tree problem until Byrka et al. [BGRS13] published an LP-based (consider Definition 3.7 for the LP) algorithm with an approximation guarantee of $\ln(4) + \varepsilon$ for every $\varepsilon > 0$. This algorithm will be one of the main subjects in this work, as we analyse it with regard to its use as a blackbox algorithm in the Karpinski et-al.-algorithm.

The QoSMT problem (definition 2.2) generalizes the Steiner tree problem. Several names, formulations and similiar problems exist for the Quality-of-Service Multicast Tree (QoSMT) problem, such as Multi-Level Steiner tree problem [AAS+18], Grade-of-Service Steiner tree problem [XLD01] or Multi-Tier Steiner tree problem [Mir96].

Charikar et al. [CNS04] were the first to present an approximation algorithm with a constant approximation factor for an unbounded number of rates. Their approximation factor of $(e \cdot \alpha) < 4.212$ for α being an approximation guarantee for the Steiner tree problem and e being Euler's number was later improved by Karpinski et al. [KMOZ05].

They presented an 3.802-approximation-algorithm for the Quality-of-Service Multicast Tree problem for an unbounded number of rates and an 1.960-approximation-algorithm for the special case that there are only two rates, further called two-rate QoSMT problem.

An algorithm without a constant approximation factor, but with very small approximation guarantees for small numbers of rates was published by Ahmed et al. [AAS+18]. Using the algorithm of Byrka et al., they provide an approximation guarantee < 1.848 for two levels in the Multi-Level Steiner tree problem.

Our contribution First, we study β-convex α-approximation algorithms (definition 2.5) which are a special group of Steiner tree approximation algorithms. These are needed as subroutines in the algorithms of Karpinski et al. [KMOZ05]. We prove that for every β-convex α-approximation algorithm, $\beta \geq 1$ always holds; and that every α-approximation algorithm for the Steiner tree problem is an α-convex $\alpha + \varepsilon$-approximation algorithm for every $\varepsilon > 0$.

Using these theorems on the algorithm of Byrka et al. [BGRS13], we improve the approximation guarantees to at least 3.769 for the QoSMT problem and at least 1.849 for the two-rate QoSMT problem. As these new results are just obtained for the highest possible β, a natural question would be whether this value could be proven to be smaller, providing another improvement on the bounds given above. We provide an argument for this being not possible if we want our α to be as good as it can be.

As this approach will not bring any better results, we concentrate on component-based LPs for the two-rate QoSMT problem in order to use the iterative randomized-rounding technique introduced by Byrka et al. [BGRS13]. The first one (definition 4.2) works with homogenous components (the costs of all edges belonging to one component are all multiplied with the same rate), and we show that an integer solution to this LP does provide a 2-approximation for the two-rate QoSMT problem, but there are graphs for which the computed solution is not optimal. The second LP (definition 4.7), which is not restricted to homogenous components, computes an exact solution to the two-rate QoSMT problem in the integral case. We concentrate on the question whether every QoSMT can be approximated arbitrarily close with a tree consisting of components each with at most k terminals. We present a theorem which says that this question can be reduced to a small subset of graphs. Furthermore we consider the approach of Borchers and Du [BD97] who have proven a similar result for the Steiner tree problem with regard to the question, whether it can be adopted for our question, and show the occuring difficulties.

Organization The rest of this work is organized as follows. First, we will introduce some basic definitions in Section 2. We will then take a look at the results of Karpinski et al. [KMOZ05] for solving the Quality-of-Service Multicast Tree problem in Section 3.1, and after that analyse how we can use the algorithm of Byrka et al. [BGRS13] in order to improve the results presented before in Section 3.2. In Section 4 we present our LPs and focus on the question whether the Quality-of-Service Multicast tree problem can be approximated with components each containing only a given limited number of terminals, trying to solve the QoSMT problem with the approach of Byrka et al. [BGRS13]. Finally, we will conclude and give a perspective about what could be studied in future works based on the results presented in this work.

2 Definitions

We first introduce some basic definitions.

Definition 2.1 (Steiner tree problem [BGRS13]). Given an undirected graph $G = (V, E)$, an edge cost function $c\colon E \to \mathbb{Q}_0^+$ and a subset of *terminals* $R \subseteq V$, the *Steiner tree problem* asks for a tree S with minimal costs *spanning all terminals*, this means $c(S) := \sum_{e \in S} c(e)$ is minimal over all trees in G spanning R. This tree S is called a *Steiner tree*, the non-terminal nodes $v \in V \setminus R$ are called *Steiner nodes*. The weight $c(S)$ is called *opt*.

Solving the Steiner tree problem is NP-hard as was shown in 1972 by Karp [Kar72]. In fact, it is already NP-hard to approximate it with an approximation factor $\leq 96/95$ [CC08]. Note that in order to obtain a minimal Steiner tree S, it is possible that S contains several Steiner nodes. It is also reasonable to replace any $G = (V, E)$ we are searching a Steiner tree for with its *metric closure* $G' = (V', E')$. We obtain G' by chosing $V' = V$, completing G and replacing our weight function c with $c'\colon E' \to \mathbb{Q}_0^+$,

$$c'(e = \{v, u\}) := \min\{c(f) \mid f \text{ is } u\text{-}v\text{-path }\}$$

It is obvious that the weights of an optimal Steiner tree S in G and an optimal Steiner tree S' in G' are equal. For any tree T' in G', a tree T in G with $c(T) \leq c'(T')$ can be found by simply taking all the paths which gave their weight to those $e' \in T'$ and removing as many as needed to get a tree. Furthermore, if there is an edge in G, the corresponding edge in G' can not have a greater weight, so if there is a Steiner tree in G, it surely is a terminal connecting tree with at most equal weight in G'. Therefore, without loss of generality we can consider metric graphs only.

The Steiner tree problem can be generalized to the Quality-of-Service Multicast Tree problem.

Definition 2.2 (Quality-of-Service Multicast Tree problem - QoSMT problem [KMOZ05]). Let $G = (V, E)$ be a graph, $l\colon E \to \mathbb{Q}_0^+$ be a function which represents the weight of every edge and $r\colon V \to \mathbb{Q}_0^+$ be a function which assigns a *rate* to each node. Let $\{r_0 = 0, r_1, r_2, \ldots, r_N\}$ be the finite range of r and define S_i to be the set of all nodes with rate r_i. Furthermore we presume that $r_i < r_j$ for $i < j$. The *Quality-of-Service Multicast Tree problem* asks for a minimum-cost subtree $S \subseteq G$, spanning a given source node s and the nodes in S_i for all $i \geq 1$, all referred to as *terminals*. Let r_e be called the *rate of edge e in S*, which is defined as the maximum rate of any node in the component of $S \setminus \{e\}$ which does not contain the source s. The cost of an edge e in S is $c(e) = l(e) \cdot r_e$. The tree S is called a *Quality-of-Service Multicast Tree* and, as in the Steiner tree case, non-terminal nodes are referred to as *Steiner nodes*.

Note that the rate of the source s is not determined. But, without loss of generality, it is valid to and we will do treat s as a node with the highest rate r_N, as it has to be connected to all r_N-nodes with this rate. Therefore, we will treat s as a terminal also.

A special case of the QoSMT problem is the *two-rate QoSMT problem*. In this case, the nodes of G have only two rates, so the range of the rate function $r\colon V \to \mathbb{Q}_0^+$ can be denoted by $\{r_0 = 0, r_1, r_2\}$.

Using the same argument as before, we can always work with the metric closure of G. By studying a graph with just one rate $r_1 = 1$ and treating the source s as an r_1 node, it is obvious that the QoSMT problem is a generalization of the Steiner tree problem and hence also NP-hard. Therefore, assuming $P \neq NP$, no algorithm time-bounded by a polynomial exists which calculates a minimal QoSMT for any given graph G.

Karpinski et al [KMOZ05] published an approximation algorithm for the QoSMT problem with a constant approximation guarantee, regardless of the number of rates. In their algorithm, they use Steiner tree approximation algorithms with a special feature. In order to understand this feature, we need a few more definitions.

Definition 2.3 (k-restricted [KMOZ05])**.** A given Steiner tree S is called *full* if every terminal $v \in R$ is a leaf. It follows that every Steiner node contained by S is an internal node. Given any Steiner tree S', it can be decomposed into a set of full components by breaking S' up at its internal terminal nodes.

A given Steiner tree is called *k-restricted* if every component obtained by this decomposition has at most k leafs (and, equivalent to that, at most k terminals). Since we are working with complete graphs, for every graph $G = (V, E)$ with optimal Steiner tree S and every $k \in \mathbb{N}_{\geq 2}$, an optimal k-restricted Steiner tree S' can be found, possibly using some edges and Steiner nodes in more than one component. Its weight $c(S')$ is denoted by opt_k, the weight $c(S)$ of the optimal tree is denoted by *opt*.

We then denote the *k-Steiner ratio* as ρ_k, which is defined as follows:

$$\rho_k := \max_{G=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\}$$

where the maximum is taken over all instances of the Steiner tree problem.

By considering the metric closure of a star graph with one Steiner node, $k+1$ terminals and an edge e with weight 1 between the Steiner node and every terminal, we easily observe that ρ_k is greater than 1 for every $k \in \mathbb{N}_{\geq 2}$. But as was shown by Borchers and Du [BD97], ρ_k gets arbitrarily close to 1 for sufficiently large k.

Theorem 2.4 (Borchers and Du [BD97])**.** *For a given $k \in \mathbb{N}_{\geq 2}$, let r and s be the non-negative integers saitsfying $k = 2^r + s$ and $s < 2^r$. Then*

$$\rho_k = \frac{(r+1) \cdot 2^r + s}{r \cdot 2^r + s} \leq 1 + \frac{1}{\lfloor \log_2 k \rfloor}.$$ [1]

Knowing this, we can define β-convex α-approximation algorithms, a subgroup of approximation algorithms that solve the Steiner tree problem. We will need them as blackbox algorithms that we can plug into the algorithm of Karpinski et al. [KMOZ05].

Definition 2.5 (β-convex α-approximation algorithm [KMOZ05])**.** A Steiner tree heuristic A is called a *β-convex α-approximation algorithm* if there exist $m \in \mathbb{N}_{\geq 2}$ and $\lambda_i \in \mathbb{R}_{\geq 0}$ for $i = 2, \ldots, m$, so that

$$\beta = \sum_{i=2}^{m} \lambda_i \text{ and } \alpha = \sum_{i=2}^{m} (\lambda_i \cdot \rho_i),$$

is fulfilled as well as the weight $l(A)$ of the tree computed by A is upper bounded by

$$l(A) \leq \sum_{i=2}^{m} (\lambda_i \cdot \text{opt}_i).$$

Note that the last equation guarantees that A is also an α-approximation algorithm for the Steiner tree problem.

The following result provides the lower bound 1 for β for every Steiner tree approximation algorithm.

Theorem 2.6. *Let A be a β-convex α-approximation algorithm, then $\beta \geq 1$.*

Proof. We presume $\beta < 1$. Since A is an α-approximation algorithm, we know that $\alpha > 1$. There exist an $m \in \mathbb{N}_{\geq 2}$ and $\lambda_i, i = 2, \ldots, m$, so that

$$\beta = \sum_{i=2}^{m} \lambda_i \text{ and } l(A) \leq \sum_{i=2}^{m} (\lambda_i \cdot \text{opt}_i).$$

Now we can pick any $n \in \mathbb{N}$ and a metric complete graph G with n nodes in which all nodes are terminals (obviously such a graph exists) and look at an optimal Steiner tree S on it. As every component in S consists of exactly one edge and the adjacent two nodes, it is obvious that $opt_2 = opt_3 = \cdots = opt$. So the length $l(A)$ of the computed tree for G' is upper bounded by

$$\sum_{i=2}^{m} (\lambda_i \cdot \text{opt}_i) = \sum_{i=2}^{m} (\lambda_i \cdot \text{opt}) = \text{opt} \cdot \sum_{i=2}^{m} \lambda_i = \text{opt} \cdot \beta < \text{opt}$$

This is impossible. Therefore $\beta \geq 1$. □

This analysis is sharp. To see this, consider the Minimal Terminal Spanning Tree algorithm which computes a terminal spanning tree T for a Graph G (as stated before, since we are studying metric and complete graphs, this does always exist). T is an optimal 2-restricted Steiner tree for G, because it is the optimal tree where no component contains more than two nodes and every component contains only one edge (the metric structure

[1]Note that Borchers and Du worked with a different definition of ρ_k. Our ρ_k is their $1/\rho_k$.

guarantees us that this is sufficient). Also we know that $\rho_2 = 2$ by inserting $k = 2$ into theorem 2.4. Therefore, the MST heuristic is a 1-convex 2-approximation algorithm.

Note that, given any α-approximation algorithm, it is not clear if such a β exists. However, this changes if we add a small amount to α.

Theorem 2.7. *Given any α-approximation algorithm A and any $\varepsilon > 0$, the following is true: A is a β-convex $(\alpha + \varepsilon)$-approximation algorithm for $\beta = \alpha$.*

Proof. Let A be an α-approximation algorithm which solves the Steiner tree problem and $\varepsilon > 0$. Choose an integer k, so that

$$\rho_k \leq \left(1 + \frac{\varepsilon}{\alpha}\right).$$

Set $m = k$, $\lambda_k = \alpha$ and $\lambda_i = 0$ for all $i = 2, \ldots, m-1$. Then

$$l(A) \leq \alpha \cdot \text{opt} \leq \alpha \cdot \text{opt}_k \leq \lambda_k \cdot \text{opt}_k = \sum_{l=2}^{m} \lambda_l \cdot \text{opt}_l,$$

and

$$\sum_{l=2}^{m} \lambda_l = \alpha,$$

$$\sum_{l=2}^{m} \lambda_l \cdot \rho_l = \lambda_k \cdot \rho_k \leq \alpha \cdot \left(1 + \frac{\varepsilon}{\alpha}\right) = \alpha + \varepsilon.$$

Therefore A is an α-convex $(\alpha+\varepsilon)$-approximation algorithm for the Steiner tree problem. □

3 Combining the algorithms of Byrka et al. and of Karpinski et al.

In this chapter we first take a look at the algorithms given by Karpinski et al. [KMOZ05] for solving the QoSMT problem for both two rates and an unlimited number of rates. We then analyse how we can use the algorithm of Byrka et al. [BGRS13] to improve the results obtained with the aforementioned algorithms.

3.1 The algorithms of Karpinski et al.

The two-rate QoSMT problem

Theorem 3.1 ([KMOZ05])**.** *Let A_1 be an α_1-approximation algorithm and A_2 be a β-convex α_2-approximation algorithm both solving the Steiner tree problem. Then there exists an algorithm* Karpinski$_1$ *(consider algorithm 1) which computes a QoSMT spanning all terminals with an approximation ratio of*

$$\max\left\{\alpha_2, \max_{r\in\mathbb{Q}_{>0}} \alpha_1 \cdot \frac{\alpha_1 - \alpha_2 r + \beta r}{\alpha_1 - \alpha_2 r + \beta r^2}\right\}^{1}$$

for any graph $G = (V, E)$, $l\colon E \to \mathbb{Q}_0^+$, two non-zero rates $r_1 < r_2$ and terminal sets $S_i \subseteq V$ of rate r_i for $i \in \{1, 2\}$.

Note here that r describes the value r_1/r_2. Not knowing which rates our graph G has, we have to assume the worst case. However, if we know which rates our input graph G has, we can refine the approximation guarantee for the specific problem.

According to Karpinski et al. [KMOZ05], the best approximation guarantee they could achieve for Karpinski$_1$ was $1.960 + \epsilon$.

The QoSMT problem for an unlimited number of rates

Theorem 3.2 ([KMOZ05])**.** *Let A be a β-convex α-approximation algorithm. Then there exists an algorithm* Karpinski$_2$ *(consider algorithm 2, which can be derandomized keeping its characteristics) which computes a QoSMT spanning all terminals with an approximation ratio of*

$$\min_{a\in\mathbb{R}_{>1}} \left((\alpha - \beta) \cdot \frac{a-1}{ln(a)} + \beta \cdot \frac{a}{ln(a)}\right)^{2}$$

[1] The exact formula from Karpinski et al. is different from this due to the fact that there is a factor of *opt* written in the second term over which the maximum is taken. By considering the last few lines in the corresponding proof, it is obvious that this *opt* was wrongfully duplicated.

Algorithm 1: Karpinski$_1$($G=(V,E)$, $l\colon E\to\mathbb{Q}_0^+$, source $s\in V$, non-zero rates $r_1<r_2$, terminal sets $S_i\subseteq V$ of rate r_i, $i\in\{1,2\}$)

$ST_1 = A_1(G, T = (s\cup S_1\cup S_2))$
$T_2 = A_1(G, T = (s\cup S_2)$
G' = contract T_2 into s
$T_1 = A_2(G', T = (s\cup S_1'))$
$ST_2 = T_1\cup T_2$
return minimum cost tree among ST_1 and ST_2

for any graph $G=(V,E)$, $l\colon E\to\mathbb{Q}_0^+$, with an unbounded number of rates.

Algorithm 2: Karpinski$_2$($G=(V,E)$, $l\colon E\to\mathbb{Q}_0^+$, source $s\in V$, $n\in\mathbb{N}$ non-zero rates r, positive number a

r = random($[0,1]$)
for $i=1,\dots,n$ **do**
    round r_i up to the next number in the set $\{a^y, a^{y+1}, a^{y+2},\dots\}$
form new sets S_i' of the same new rate
$T=\emptyset$
for $r_i'\neq 0$ in decreasing order **do**
    $T_i = A(G, R = S_i'\cup s)$
    $T = T\cup T_i$
    contract T_i into s
return T

According to Karpinski et al. [KMOZ05], the best approximation guarantee they could achieve for Karpinski$_2$ was ≤ 3.802.

3.2 An analysis of the Byrka et al. algorithm

We will now insert the algorithm of Byrka et al. [BGRS13] into Karpinski$_1$ and Karpinski$_2$.

A first result The algorithm of Byrka et al., further called B, can have an approximation guarantee of $ln(4)+\varepsilon$, where we can choose $\varepsilon>0$ at will. Without knowing anything else, we can say the following:

Theorem 3.3. *Set $\varepsilon>0$. Then there exists an ($ln(4)+\varepsilon$)-convex ($ln(4)+\varepsilon$)-approximation algorithm solving the Steiner tree problem.*

[2] Karpinski et al. took the minimum not over all $a>1$ but over all $a>0$. As this could (and would in our case) result into negative approximtion guarantees, we think that this was just forgotten to mention.

Proof. We know that B can have an approximation guarantee of $\ln(4)+\varepsilon/2$. It follows by theorem 2.7, that B can be an $(\ln(4)+\varepsilon/2)$-convex $(\ln(4)+\varepsilon)$-approximation algorithm. We can choose a higher convexity as we wish. The claim follows. □

We can now insert B into Karpinski$_1$ and Karpinski$_2$.

Theorem 3.4. *Set $\varepsilon > 0$. Then there exists an $((\ln(4)\cdot 4)/3+\varepsilon)$ - approximation algorithm which solves the QoSMT for every graph $G=(V,E)$ with edge weights $l\colon E \to \mathbb{Q}_0^+$ and with a two-rate function $r\colon V \to \mathbb{Q}_0^+$.*

Proof. We combine the theorems 3.1 and 3.3. By setting $A_1 = A_2 = B$, we know that there exists an algorithm A which solves the QoSMT problem and has an approximation guarantee of

$$\max\left\{\ln(4)+\frac{\varepsilon\cdot 3}{4}, \max_{r\in\mathbb{Q}_{>0}}\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot\frac{\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)-\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot r+\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot r}{\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)-\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot r+\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot r^2}\right\}$$

$$\begin{aligned}
&=\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot\max\left\{1,\max_{r\in\mathbb{Q}_{>0}}\frac{1}{1-r+r^2}\right\}\\
&=\left(\ln(4)+\frac{\varepsilon\cdot 3}{4}\right)\cdot\frac{4}{3}\\
&=\frac{\ln(4)\cdot 4}{3}+\varepsilon.
\end{aligned}$$

□

The factor $((\ln(4)\cdot 4)/3)$ is smaller than 1.849, which is clearly smaller than $1.960+\varepsilon$ which was the best approximation guarantee obtained by Karpinski et al. [KMOZ05].

We will now prove a similar result for an unlimited number of rates.

Theorem 3.5. *Let $\varepsilon > 0$. Then there exists an $(\ln(4)\cdot e+\varepsilon)$ - approximation algorithm which solves the QoSMT for every graph $G=(V,E)$ with edge weights $l\colon E\to\mathbb{Q}_0^+$ and with an arbitrary number of rates $r\colon V\to\mathbb{Q}_0^+$.*

Proof. The proof is quite similar to the proof of theorem 3.4. We combine the theorems 3.2 and 3.3. It follows that there exists an algorithm A which solves the QoSMT problem and has an approximation guarantee of

$$\begin{aligned}
&\min_{a\in\mathbb{R}_{>0}}\left(\left(\left(\ln(4)+\frac{\varepsilon}{e}\right)-\left(\ln(4)+\frac{\varepsilon}{e}\right)\right)\cdot\frac{a-1}{\ln(a)}+\left(\ln(4)+\frac{\varepsilon}{e}\right)\cdot\frac{a}{\ln(a)}\right)\\
&=\min_{a\in\mathbb{R}_{>0}}\left(\left(\ln(4)+\frac{\varepsilon}{e}\right)\cdot\frac{a}{\ln(a)}\right)\\
&=(\ln(4)+\frac{\varepsilon}{e})\cdot\min_{a\in\mathbb{R}_{>0}}(\frac{a}{\ln(a)}).
\end{aligned}$$

The minimum is at $a=e$, e being Eulers number, and we obtain the following:

$$\left(\ln(4) + \frac{\varepsilon}{e}\right) \cdot \min_{a \in \mathbb{R}_{>0}} \left(\frac{a}{\ln(a)}\right)$$
$$= \ln(4) \cdot e + \varepsilon.$$

□

To compare this result to the one obtained by Karpinski et al. [KMOZ05] it is notable that $ln(4) \cdot e < 3.769$. We can see that this is smaller, but there is not a big difference compared to 3.802. To obtain a better bound on the approximation guarantee of Karpinski_1 (algorithm 1) and Karpinski_2 (algorithm 2), we tried to analyse the algorithm of Byrka et al. [BGRS13] in order to get a better β.

Can β be improved? Unfortunately, a refined analysis of B brought no better bound on β. In the following we give an argument (yet no proof) why this is the case.

We first have to know how B works. Therefore, we have to take a look at the LP used to solve the Steiner tree problem it is based on.

Definition 3.6 (directed component cut relaxation - DCR)**.** Let $G = (V, E)$ with edge weights $c\colon E \to \mathbb{Q}_0^+$ be an instance of the Steiner tree problem, $R \subseteq V$ be the subset of terminals and $r \in R$ an arbitrary root node.

We then define a *directed component* C as follows: Consider a subset $R' \subseteq R$ and pick a *sink node* $r' \in R'$ which is denoted as $\text{sink}(C)$, then the component is the *minimal Steiner tree* between R' where *all edges are directed to* r'.

We denote the *set of all components* obtained this way by C_n and say that a component C *crosses* a set $U \subseteq R \backslash \{s\}$ if at least one terminal in C is inside U and $\text{sink}(C)$ is outside. The *set of components crossing* U is denoted by $\delta^+_{C_n}(U)$. The cost of a component C is denoted by $c(C)$.

The LP relaxation is then:

$$min \sum_{C \in C_n} c(C) \cdot x_C$$

with the following constraints:

$$\sum_{C \in \delta^+_{C_n}(U)} x_c \geq 1 \quad \forall U \subseteq R \setminus \{r\}, U \neq \varnothing$$

$$x_C \geq 0 \quad \forall C \in C_n.$$

As the cardinality of C_n is exponential, we have far too many constraints to solve our LP in polynomial time. Therefore, an approximative LP is needed. We achieve this by allowing only a polynomial account of components to be used.

Definition 3.7 (k-DCR)**.** For any $k \in \mathbb{N}_{\geq 2}$, $C_k \subset C_n$ describes the set of components *with at most k terminals*, and for a given cut $U \subseteq R \setminus \{r\}$ the set of componets in C_k crossing U is denoted by $\delta^+_{C_k}(U)$. The LP relaxation is then:

$$\min \sum_{C \in C_k} c(C) \cdot x_C$$

with the following constraints:

$$\sum_{C\in\delta^+_{C_k}(U)} x_c \geq 1 \quad \forall U \subseteq R\setminus\{r\}, U \neq \varnothing$$

$$x_C \geq 0 \quad \forall C \in C_k.$$

Note that this LP is equivalent to DCR with the additional constraint that $x_C = 0$ for every component C with more than k terminals.

The following result will make sure that k-DCR provides a $(1+\varepsilon)$-approximation solution to DCR if we choose our k big enough, and that we can solve k-DCR in polynomial time for every given k.

Theorem 3.8 ([BGRS13])**.** *DCR is a relaxation of the Steiner tree problem, and k-DCR is a relaxation of the k-restricted Steiner tree problem. In the integral case, an optimal solution to the (k-restricted) Steiner tree problem can be gathered from an ILP solution to (k-)DCR by taking every component C where $x_C = 1$ is true and putting them together. It follows from theorem 2.4 that, in this integral case, for the cost $c(S)$ of an optimal Steiner tree and the cost $c(S')$ of an optimal k-restricted Steiner tree obtained this way, $c(S') \leq \rho_k \cdot c(S')$.*

This holds for fractional solutions. This means that for every G being an instance of the Steiner tree problem, for the cost opt_f of an optimal fractional solution to DCR and the cost $opt_{f,k}$ of an optimal fractional solution to k-DCR, the following holds:

$$opt_{f,k} \leq \rho_k \cdot opt_f.$$

Also, for every instance G of the Steiner tree problem and any given $k \in \mathbb{N}_{\geq 2}$, a fractional solution to k-DCR can be found by an algorithm which is time bounded by a polynomial.

The algorithm of Byrka et al. works iteratively and contracts one component per iteration, so the set of components changes. By C^t we denote the set of all components and by x^t the approximative solution to DCR *at the begin of iteration t*. As we solve DCR only approximately, we denote each solution to it as (x, C) where x denotes the solution vector and C the set used in the approximation. Now we have everything we need to take a look at B (algorithm 3).

Theorem 3.9 ([BGRS13])**.** *B is a randomized $ln(4)+\varepsilon$ approximation algorithm for the Steiner tree problem. It can be derandomized keeping its approximation guarantee.*

It is not known if the approximation guarantee given in 3.9 is sharp. Therefore, all following statements on a lower bound for β are only assumptions (because a better approximation guarantee would directly result in a better β according to theorem 2.7). But we will now argue that for the best approximation guarantee α (which probably is a $\alpha' + \varepsilon$ factor for any $\varepsilon > 0$), the difference between α and β is nearly zero. The reason for this lies in the approximation of DCR. In order to get the best approximation guarantee

Algorithm 3: Byrka($G = (V, E)$. terminals $R \subset V$, edge weights $c \colon E \to \mathbb{Q}_0^+$, source $r \in R$, $\varepsilon > 0$

for $t = 1, 2, \dots$ **do**
  compute a $(1+\varepsilon)$-approximate solution (x^t, C^t) to DCR w.r.t. G
  sample one component $C^{t\prime}$, each component C has probability $x_C^t / \sum_{C' \in C^t} x_{C'}^t$ to be chosen
  contract C^t into its sink $sink(C^t)$
  if only one terminal remains **then**
    return $\biguplus_{i=1}^{t} C^{i\prime}$

we can, we have to solve k-DCR for a very big k. So we have to assume that components with k terminals are contracted and therefore taken into the resulting tree. We know that, in order to have a β-convex α-approximation algorithm, we need an $m \in \mathbb{N}_{\geq 2}$ and $\lambda_i \in \mathbb{R}_{\geq 0}$ for $i = 2, \dots, m$, so that

$$\beta = \sum_{i=2}^{m} \lambda_i,\ \alpha = \sum_{i=2}^{m} \lambda_i \cdot \rho_i \text{ and the length of the tree } l(B) \leq \sum_{i=2}^{m} \lambda_i \cdot opt_i.$$

As stated before it is most likely that k-components are used, which means that we could not estimate the costs of our solution in terms of opt_l for any $l \in 2, \dots, (k-1)$. This would mean that all λ_l are zero for every $l \in 2, \dots, (k-1)$. It follows that

$$\beta = \sum_{i=k}^{m} \lambda_i \text{ and } \beta \leq \alpha = \sum_{i=k}^{m} \lambda_i \cdot \rho_i \leq \rho_k \cdot \sum_{i=k}^{m} \lambda_i = \rho_k \cdot \beta.$$

As ρ_k is nearly 1 for a big k, this means that β and α are nearly the same.

We assume in our statement that we want to have the best possible approximation guarantee first and analyse the corresponding β second. An interesting question (and not further discussed in this work) would be if we could get a better β if we would willingly choose a low k.

4 Component-based LPs for the two-rate QoSMT problem

The intention of this section was to adapt the approach of Byrka et al. for the two-rate QoSMT problem. We therefore present two different component-based LPs solving it, the first one works only with components which contain edges of only one rate, and we will prove that it is therefore only approximative. The second one is exact, but the corresponding fractional linear program cannot prohibit a forbidden combination of components. We will also give a reasonable assumption that a QoSMT can be approximated with k-restricted trees, but we can not prove it.

4.1 An approach to the two-rate QoSMT problem with homogenous components

We now introduce an LP which computes a tree with minimal weight for every instance G of the QoSMT problem due to the condition that every component contains only edges of one rate. Therefore, we have to modify our definition of a component.

Definition 4.1 (directed homogenous two-rate component)**.** Let $G = (V, E)$, $l\colon E \to \mathbb{Q}_0^+$ and $r\colon V \to \mathbb{Q}_0^+$ be an entity of the two-rate QoSMT problem and $R \subseteq V$ be the subset of terminals. Define a *directed homogenous two-rate component* C as follows: Consider a subset $R' \subseteq R$ and note whether it is treated as a r_1 or a r_2 component (denoted by r_C which is 0 in case of r_1 and 1 in case of r_2). Pick a sink node $r \in R'$ which is denoted as $\text{sink}(C)$. Then the component is the minimal QoSMT between R' where every terminal v is seen as given in r_C, and all edges are directed to r. Note that this means that C contains only edges of the rate given in r_C.

Similar to the Steiner tree case, we denote the *set of all components* obtained this way by C_n^* and say that a component C *crosses a set* $U \subseteq R \setminus \{s\}$ if at least one terminal of C is inside U and $\text{sink}(C)$ is outside. The set of components with $r_C = r_2$ is denoted as $C_{n,2}^*$, the set of components crossing a cut U is denoted by $\delta^+_{C_n^*}(U)$. The cost of a component C is denoted by $c(C)$, in which the rate factor is already considered.

We can now modify DCR:

Definition 4.2 (homogenous two-rate DCR)**.** Let $G = (V, E)$ with edge weights $l\colon E \to \mathbb{Q}_0^+$, rates $r\colon V \to \mathbb{Q}_0^+$ and a given source $s \in V$ be an entity of the two-rate QoSMT problem and $R \subseteq V$ be the subset of terminals, $S_2 \subset R$ be the terminals with rate r_2. The LP relaxation then is:

$$\min \sum_{C \in C_n^*} c(C) \cdot x_C$$

with the following constraints:

$$\sum_{C \in \delta^+_{C_n^*}(U)} x_c \geq 1 \quad \forall U \subseteq R \setminus \{s\}, U \neq \varnothing$$

$$\sum_{C \in \delta^+_{C_{n,2}^*}(U)} x_c \geq 1 \quad \forall U \subseteq S_2 \setminus \{s\}, U \neq \varnothing$$

$$x_C \geq 0 \quad \forall C \in C'_n.$$

Note that the only real difference between DCR and homogenous two-rate DCR (apart from the fact that they expect different input classes) is the second constraint, which assures that a flow of 1 connects every r_2 terminal with the source s.

This LP can easily be adapted for more than two rates by adding more constraints similar to the second for each additional rate. We now have to show that homogenous two-rate DCR fulfills what we claimed before.

Theorem 4.3. *For every instance of the two-rate QoSMT problem, the integral homogenous two-rate DCR computes a QoSMT T which is minimal due to the condition that every component in T contains only edges of one rate (we will refer to this subproblem as homogenous QoSMT problem).*

Proof. Let $G = (V, E)$ with edge weights $l \colon E \to \mathbb{Q}_0^+$, rates $r \colon V \to \mathbb{Q}_0^+$ and a given source $s \in V$ be an entity of the two-rate QoSMT problem and $R \subseteq V$ be the subset of terminals, $S_2 \subset R$ be the terminals with rate r_2.

We will show first that every homogenous QoSMT S for G can be converted into an integral solution of homogenous two-rate DCR with the same costs. Take every component C in S and set $x_C = 1$ for the corresponding x_C, all other $x_{C'}$ get assigned the value 0, so the costs of both S and the minimized LP-function are equal. We still have to show that the constraints are fulfilled: As all variables have the value 0 or 1, that is trivial for the last constraint. Now consider any $U \subseteq R \setminus \{s\}$, $U \neq \varnothing$. Take a terminal $r \in U$. There exists a path in S from r to s, therefore there has to be an edge e in S which crosses U. We know that e is part of a component C in S, so $C \in \delta^+_{C_n^*}(U)$ and $x_C = 1$, therefore the first constraint is fulfilled. The fulfillment of the second constraint can be shown exactly the same way.

The other way around, if we have a feasible solution for homogenous two-rate DCR, we can obtain a graph T connecting all terminals to the source and connecting all r_2 nodes with an r_2 path to the source with equal costs by building T exactly out of the components (with the same rate) where $x_C = 1$. To show that such a $r - s$ path exists for every terminal $r \in R$, we consider components crossing $\{r\}$. Due to the first constraint, there exists at least one which we will call C. Therefore we have a $r - \mathrm{sink}(C)$-path in T. If $\mathrm{sink}(C)$ is not already s, we can now iterate this process, each time looking at the cross between the already discovered terminals and the rest, and will always find a component contained in T crossing this cut. As G is finite and every cut really contains the one before, this will result in taking a component C' with $s = \mathrm{sink}(C')$ after a finite number of iterations. Therefore, r and s are connected, and such a path exists. We can show that a r-s-path with rate r_2 exists for every $r \in S_2$ exactly the same way by using the second constraint. We know that T, if it is no QoSMT itself, contains a homogenous QoSMT S so that $c(S) < c(T)$.

As homogenous two-rate DCR calculates a minimal solution, the claim follows. □

We will now provide some bounds on how good the approximation of homogenous two-rate DCR is. Therefore we will first introduce some notations.

Definition 4.4 (homogenous QoSMT (k-) ratio)**.** Take any $k \in \mathbb{N}_{\geq 2}$ and denote for every G being an entity of the two-rate QoSMT problem the weight of an optimal homogenous QoSMT S^* by opt_{Hom}. We then denote by

$$\rho^* := \max_{G=(V,E)} \left\{ \frac{\mathrm{opt}_{\mathrm{Hom}}}{\mathrm{opt}} \right\},$$

the *homogenous QoSMT ratio*, which is the worst factor between an optimal QoSMT S and an optimal homogenous QoSMT S^* for all graphs G being entities of the two-rate QoSMT problem.

We similarly denote the *homogenous QoSMT k-ratio* by

$$\rho_k^* := \max_{G=(V,E)} \left\{ \frac{\mathrm{opt}_{\mathrm{Hom},k}}{\mathrm{opt}} \right\}$$

where $opt_{\mathrm{Hom},k}$ denotes the value of the optimal homogenous QoSMT which contains only components with $\leq k$ terminals.

It is clear that $\rho_k^* \geq \rho^*$ for every $k \in \mathbb{N}_{\geq 2}$. To ensure that homogenous two-rate DCR produces an optimal solution to the two-rate QoSMT problem, $\rho^* = 1$ would be necessary. Unfortunately, that is not the case, which is proven together with some other bounds on ρ^* and ρ_k^* in the following theorem.

Theorem 4.5. *Let $k \in \mathbb{N}_{\geq 2}$. For the homogenous two-rate QoSMT ratio ρ^* and for the homogenous two-rate QoSMT k-ratio ρ_k^*, the following holds:*

$$\frac{6}{5} \leq \rho^* \leq 2;$$

$$\frac{6}{5} \leq \rho_k^* \leq \rho_k \cdot 2.$$

Paticularly this implies that ρ^* and every ρ_k^* are finite and therefore exist. Yet the two-rate QoSMT problem can not be solved exactly using homogenous components only.

Proof. We will at first proof $6/5 \leq \rho^*$. Consider the metric closure G' of the graph G given in figure 4.1a with one source, one Steiner node, one r_1 and one r_2 terminal, in which the Steiner node is connected to every other node with weight 1. Set $r_1 = 1$ and $r_2 = 2$. Then the weight $c(S)$ of the optimal QoSMT S for G' (note that S is equivalent to G) is 5. If we only allow homogenous components, every minimal tree we can obtain (an example is given in figure 4.1b) has weight 6, which can be verified due to the small number of nodes by looking at all possible homogenous QoSMTs. Therefore,

$$\rho^* = \max_{G=(V,E)} \left\{ \frac{\text{opt}_{\text{Hom}}}{\text{opt}} \right\} \geq \frac{6}{5}.$$

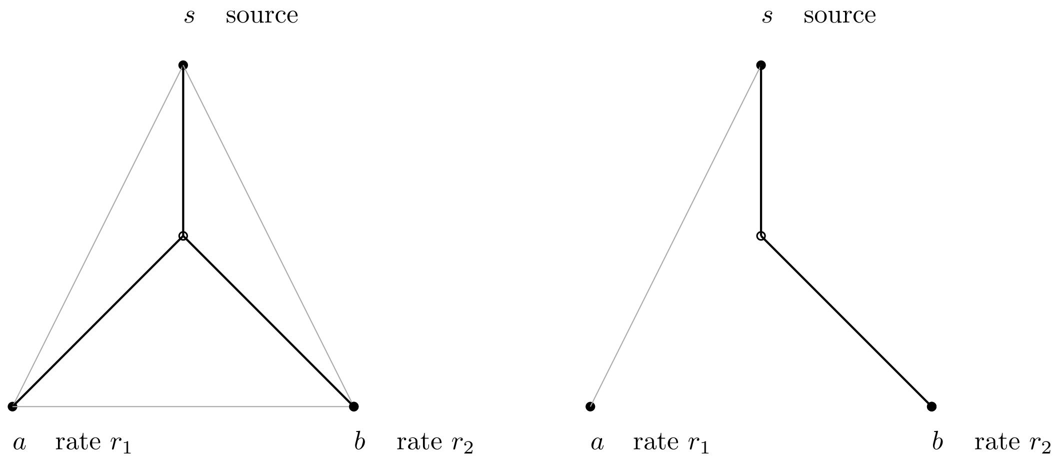

(a) The edges in G are black, all with weight one, the grey edges form together with the black edges the metric closure G' of G.

(b) A minimal homogenous QoSMT on G'. The grey component is treated as r_1 component, the black one as r_2 component.

Fig. 4.1: A counterexample for $\rho^* < 6/5$.

To show that $\rho^* \leq 2$ we consider an arbitrary graph G being an entity of the QoSMT problem with the two rates r_1 and r_2, and any minimal QoSMT S for G. We can describe the weight of S by $c(S) = r_1 \cdot c(S_1) + r_2 \cdot c(S_2)$, S_i each being the induced subgraph of S containing exactly all edges of weight r_i. We can now construct a homogenous graph S^* by first taking every component in S_2 as r_2-component and then adding every component in S as r_1-component. So every terminal in S is connected to the source with at least an r_1 path and every r_2 terminal is connected to the source via an r_2 path, as there exists such a path in S_2. Note that we now may have more edges and components than we might need, but that is not important, because S^* contains a homogenous QoSMT in every case. Then we can say that

$$\rho^* \leq \frac{c(S^*)}{c(S)} = \frac{r_1 \cdot c(S_1) + r_1 \cdot c(S_2) + r_2 \cdot c(S_2)}{r_1 \cdot c(S_1) + r_2 \cdot c(S_2)} \leq \frac{r_1 \cdot c(S_1) + r_2 \cdot c(S_2) + r_2 \cdot c(S_2)}{r_1 \cdot c(S_1) + r_2 \cdot c(S_2)} \leq 2.$$

Let us come to the homogenous k-ratios. Since $\rho_k^* \leq \rho^*$ for every $k \in \mathbb{N}_{\geq 2}$, $6/5 \leq \rho_k^*$ follows immediately.

It is left to show that $\rho_k^* \leq 2 \cdot \rho_k$. We again consider any graph G being an entity of the QoSMT problem with the two rates r_1 and r_2, and any minimal QoSMT S for G. Like above, we find a homogenous QoSMT S^* so that $c(S^*) \leq 2 \cdot c(S)$. If we consider each component C in S^* of its own, we find a k-restricted tree T_C, so that $c(T_C) \leq \rho_k \cdot c(C)$, which connects all terminals in C, as this problem can be seen as a Steiner tree problem. By giving each T_C the rate of the corresponding component C and combining all these trees together into one k-restricted tree T, we obtain an upper bound $c(T)$ on the weight of the optimal k-restricted homogenous QoSMT. We can now argue that

$$\rho_k^* = \max_{G=(V,E)} \left\{ \frac{\text{opt}_{\text{Hom},k}}{\text{opt}} \right\} \leq \frac{c(T)}{c(S)} \leq \frac{\rho_k \cdot c(S^*)}{c(S)} \leq \rho_k \cdot 2.$$

□

We think both homogenous two-rate DCR and the proven bounds for ρ^* and ρ_k^* are of independent interest, especially since the bounds are not sharp and it is possible that smaller upper bounds can be proven. Nevertheless theorem 4.5 shows that homogenous two-rate DCR is useless for our purpose, as we need an exact solution (or at least a $(1+\varepsilon)$-approximation for any $\varepsilon > 0$) for our two-rate QoSMT problem. So, to have the possibility to get a useable component-based LP, we have to allow components with mixed edge rates.

4.2 Approximating QoSMTs with k-restricted trees

We give consideration to the results proven above by introducing a new class of components.

Definition 4.6 (directed two-rate component)**.** Let $G = (V, E)$, $l\colon E \to \mathbb{Q}_0^+$ and $r\colon V \to \mathbb{Q}_0^+$ be an entity of the two-rate QoSMT problem and $R \subseteq V$ be the subset of terminals. Define a *directed two-rate component* C as follows: Consider a subset $R' \subseteq R$, pick a sink node $r \in R'$ which is denoted as $\text{sink}(C)$, and note for every $v \in R'$ whether it is treated as an r_1 or an r_2 node (denoted by $r_{C,v}$ which is 0 in case of r_1 and 1 in case of r_2), observing the following rules: The rate of all r_2 nodes has to be r_2 in every component, and for a component C, $v_{C,\text{sink}(C)} \geq v_{C,v'}$ has to be true for every terminal $v' \in C$. Then the component is the minimal QoSMT between R' where every terminal v' is seen as given in $r_{C,v'}$, and all edges are directed to r.

We denote the *set of all components* obtained this way by C'_n and say that a component C *crosses* a set $U \subseteq R \setminus \{s\}$ if at least one terminal of C is inside U and $\text{sink}(C)$ is not. The *set of components crossing* U is denoted by $\delta^+_{C'_n}(U)$. The cost of a component C is denoted by $c(C)$, where the assignments of the terminals in C are considered.

We say that two components C and C' *violate* each other if there exists a terminal $v \in R$, so that $v = \text{sink}(C)$, $v \in C' \setminus \{\text{sink}(C')\}$ and $v_{C,v} > v_{C',v}$ (in other words, C flows into C', but they treat the rate of v differently and therefore C' cannot support the flow coming from C). We denote the *set of pairs* $\{C, C'\}$ *violating each other* by K'_n.

We can now introduce another LP for the two-rate QoSMT problem.

Definition 4.7 (two-rate DCR)**.** Let $G = (V, E)$ with edge weights $l \colon E \to \mathbb{Q}_0^+$, rates $r \colon V \to \mathbb{Q}_0^+$ and a given source $s \in V$ be an entity of the two-rate QoSMT problem, $R \subseteq V$ be the subset of terminals and $S_2 \subset R$ be the terminals with rate r_2.

The LP relaxation then is:

$$\min \sum_{C \in C'_n} c(C) \cdot x_C$$

with the following constraints:

$$\sum_{C \in \delta^+_{C'_n}(U)} x_c \geq 1 \quad \forall U \subseteq R \setminus \{s\}, U \neq \varnothing$$

$$x_C \geq 0 \quad \forall C \in C'_n$$

$$x_C + x'_C \leq 1 \quad \forall \{C, C'\} \in K'_n.$$

Note that in the integral case, the last constraint ensures that no conflicts emerge in how a node should be seen.

As before homogenous two-rate DCR, we can generalize two-rate DCR for any number of rates, using the same definition of components violating each other. We will now show that two-rate DCR really is a relaxation of the two-rate QoSMT problem.

Theorem 4.8. *For every instance of the two-rate QoSMT problem, the integral two-rate DCR computes a minimal QoSMT T.*

Proof. Let $G = (V, E)$ with edge weights $l \colon E \to \mathbb{Q}_0^+$, rates $r \colon V \to \mathbb{Q}_0^+$ and a given source $s \in V$ be an entity of the two-rate QoSMT problem, $R \subseteq V$ be the subset of terminals and $S_2 \subset R$ be the terminals with rate r_2.

We will show first that a minimal QoSMT S for G can be converted into an integral solution of two-rate DCR with the same costs. Similarly to the proof of theorem 4.3 we take every component C in S and set $x_C = 1$ for the corresponding x_C, all other $x_{C'}$ get assigned the value 0, so the costs of both S and the minimized LP-function are equal. To show the fulfillment of the first constraint in the proof of theorem 4.3 we have not used that we have homogenous components, so we can just repeat the steps to obtain the same result. It is to show that the last constraint is fulfilled. If we take a look at our graph G, it is obvious that no pair of components in it can violate each other. So for every pair $\{C, C'\} \in K'_n$, only one C or C' can be used. This proves the fulfillment of the last constraint.

The other way around, if we have an optimal solution for homogenous two-rate DCR, we can obtain a graph T connecting all terminals to the source and connecting all r_2 nodes with an r_2 path to the source with equal costs by building T exactly out of the components (with the same rate for each node and edge) where $x_C = 1$. To show that an r-s-path exists for every terminal $r \in R$, we can just follow the corresponding steps in the proof of theorem 4.3, as we have not used the fact that our components are homogenous there. It is left to show that every terminal with rate r_2 ($\neq s$) is connected to the source with an r_2 path in T. We therefore take a look at the u-v-path f given by the seqence of components C_i and their sinks between r and s. The first sink s_1 has to be treated as an r_2 node in the component containing r, because of our rules for the treatment of sinks. So r is connected to s_1 with a rate r_2 path. Now s_1 is again part of a component C_2 as a terminal (if it is not already s), and because we know that C_1 and C_2 cannot violate each other, s_1 has to be treated as an r_2 node in C_2. We can now iterate this process over all components f is part of, arguing that every edge $e \in f$ has rate r_2. The claim follows. □

Theorem 4.8 proves that we can now try to repeat all the steps done by Byrka et al. [BGRS13] to use two-rate DCR instead of DCR to obtain an approximation algorithm for the two-rate QoSMT problem. But there are some problems.

At first, we have to note that two-rate DCR is a very weak LP as any fractional solution can use violating components (e.g. by setting $x_C = x_{C'} = 1/2$ for a pair $\{C, C'\} \in K'_n$), which leads to edge rates which are hard to use or not useful at all. So it would be preferable to find another, stronger LP.

The second problem is that given a mixed component, it is not clear how the contraction has to be done, because we have to assign a rate to the new node v (a statement, which is true not only for two-rate DCR, but for every component-based LP exactly solving the QoSMT problem). If we assign the highest rate, we will think that we can connect r_2 terminals to v, even if the connection would have led over an r_1 node (and therefore not be allowed in the former graph). On the other hand, if we assign the lower rate to v, we can not connect any r_2 terminals (even if we could have done that before), which can lead to higher costs which would have been unnecessary. A possible solution to this problem could be not to contract the component C but to assign its variable x_C a fix value of 1 before solving the LP again.

The last problem (which we study in the rest of this section) is that we do not know if we can approximate QoSMTs with k-restricted trees. This would be necessary, since (like in the DCR case) we have far too many components and constraints to solve the LP in polynomial time.

Therefore, to procede, we will introduce a notation describing how good the approximation with k-components is.

Definition 4.9 (QoSMT k- ratio). Let $k \in \mathbb{N}_{\geq 2}$. We then denote the *k-restricted QoSMT ratio* by

$$\rho'_k := \max_{G=(V,E)} \left\{ \frac{\mathrm{opt}_k}{\mathrm{opt}} \right\}$$

where opt_k denotes the value of the optimal QoSMT which contains only components with $\leq k$ terminals.

As we can use all components, the homogenous as well, it follows that $\rho'_k \leq \rho^*_k$ for every $k \in \mathbb{N}_{\geq 2}$. Therefore it follows from theorem 4.5 that all ρ'_k are finite.

As stated before, we do not know if ρ'_k gets as close to 1 as we need. It is therefore necessary to find a proof or a counter example. The following theorem shows that we can restrict our search to a small subgroup of graphs.

Theorem 4.10. *Take any $k \in \mathbb{N}_{\geq 2}$. Then the following equation holds:*

$$\rho_k' = \max_{G^1=(V,E)} \left\{ \frac{opt_k}{opt} \right\} = \max_{G'^2=(V,E)} \left\{ \frac{opt_k}{opt} \right\}$$

where the first maximum is taken over all graphs G which possess a full *optimal QoSMT S and second the maximum is taken over all graphs G' which are the* metric closure *of a tree T whichs leafs are coincident with the terminals. Note that T is also an optimal QoSMT for the corresponding G'.*

Proof. We will prove the following three inequalities:

$$\rho_k' \leq^1 \max_{G^1=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \leq^2 \max_{G'^2=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \leq^3 \rho_k'.$$

We will refer to the various maxima as the originial (being the maximum in definition 4.9), the first and the second one (the ones used in this theorem, depending on their order in the equation). Now take any $k \in \mathbb{N}_{\geq 2}$.

We will prove the first inequality by showing that, given any graph G with optimal QoSMT S, we can construct a k-restricted tree S', so that

$$c(S') \leq \max_{G^1=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \cdot c(S).$$

Take S and deconstruct it into its components C_i. Assign a source s_i to every component C_i by picking the node that connects C_i to the rest of the tree containing the source and treat each node as if it would have the highest rate it is connected with to s in S. Now we have a number of independent graphs all having a full QoSMT. We take an optimal k-restricted T_i' for each C_i. By putting them together (we know that the edge rates still fulfill their requirements because of our treatments) we obtain a tree S', so that

$$c(S') \leq \max_{G^1=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \cdot \sum_i c(C_i) = \max_{G^1=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \cdot c(S).$$

We will prove the second inequality similarly to the first by taking any graph $G = (V, E)$ which has a full Steiner tree S and show that we can construct a k-restricted tree S', so that

$$c(S') \leq \max_{G'^2=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \cdot c(S).$$

We define a new graph $G' = (V', E')$ being the metric closure of S. Note that $V' \subseteq V$, $E' \subseteq E$, S is an optimal QoSMT for G' and every edge $e \in E'$ has the same or less costs in E' as in E. We can now find an optimal k-restricted tree T in G'. This is also a terminal connecting k-restricted tree S' in G, so with that $c(S') \leq c(T)$, and we know that

$$c(S') \leq c(T) \leq \max_{G^2=(V,E)} \left\{ \frac{\text{opt}_k}{\text{opt}} \right\} \cdot c(S).$$

The third inequality follows directly out of the fact that we take a maximum in both cases, and every graph considered in the second maximum is also considered in the original one. $\square$

As the same problem has already been solved for the Steiner tree case by Borchers and Du [BD97], an implied question is:

Can the approach of Borchers and Du [BD97] be used to prove that ρ'_k gets arbitrarily close to 1? To outline an answer for this question, we will first take a look at the structure of the proof of theorem 2.4. To prove the inequality

$$\rho_k \leq \frac{(r+1)\cdot 2^r + s}{r \cdot 2^r + s}$$

they first reduce the problem to the case where they have a given graph G with a full Steiner tree S (similar to our theorem 4.10). All further work is performed on S and all edges used in the constructed components are edges in S. Then they transform S into a full binary tree S' with the same weight in which all leaves (and therefore terminals) are in the last layer or the one before. For r and s being the non-negative integers satisfying $k = 2^r + s$, a number of $r \cdot 2^r + s$ k-restricted trees T_l for $l = 1, \dots, r \cdot 2^r + s$ is created, each of them spanning all terminals. The Steiner tree S' is disjointed into a number of disjoint paths, and it can be shown that the weight of every path is counted only 2^r times in the sum $c(L_1) + c(L_2) + \dots + c(L_{r\cdot 2^r+s})$ where each L_i is a subgraph of T_i, so that the part of T_i disjoint to L_i covers S' exactly once. They draw the conclusion that

$$c(L_1) + c(L_2) + \dots + c(L_{r\cdot 2^r+s}) \leq 2^r \cdot c(S)$$

and therefore there has to be a $d \in \{1, \dots, r \cdot 2^r + s\}$ so that

$$c(L_d) \leq \frac{2^r}{r \cdot 2^r + s} \cdot c(S).$$

Combining the previos results, they conclude

$$c(T_d) = c(S) + c(L_d) \leq (1 + \frac{2^r}{r \cdot 2^r + s}) \cdot c(S)$$

and have therefore found a k-restricted tree T_d so that $c(T_d)/c(S) \leq \rho_k$ fulfills the inequality which had to be proven.

As can be seen this proof relies on the fact that we know the exact edge costs in S' and how often these edges are used in the k-restricted trees. If we now take a look at our corresponding k-restricted QoSMT problem, we do know what the edge costs in the optimal QoSMT are, but if we just construct a big number of k-restricted trees we can not say how the edges are treated in these. This is getting problematic at the point where it can happen that the rate r_1 nodes which could be connected to the source over a low-cost r_1 path are the docking points for components containing nodes with rate r_2. But in order to obtain a bound on ρ'_k which gets arbitrarily close to 1, we can not pre-estimate the costs of these edges being treated as r_2 edges, therefore we think that this approach leads only to the $\rho'_k \leq 2 \cdot \rho_k$ factor proven above.

We will substantiate this vague argument on shortcutting trees. We obtain these shortcutting trees as follows: Let $G = (V, E)$ be a complete, metric instance of the

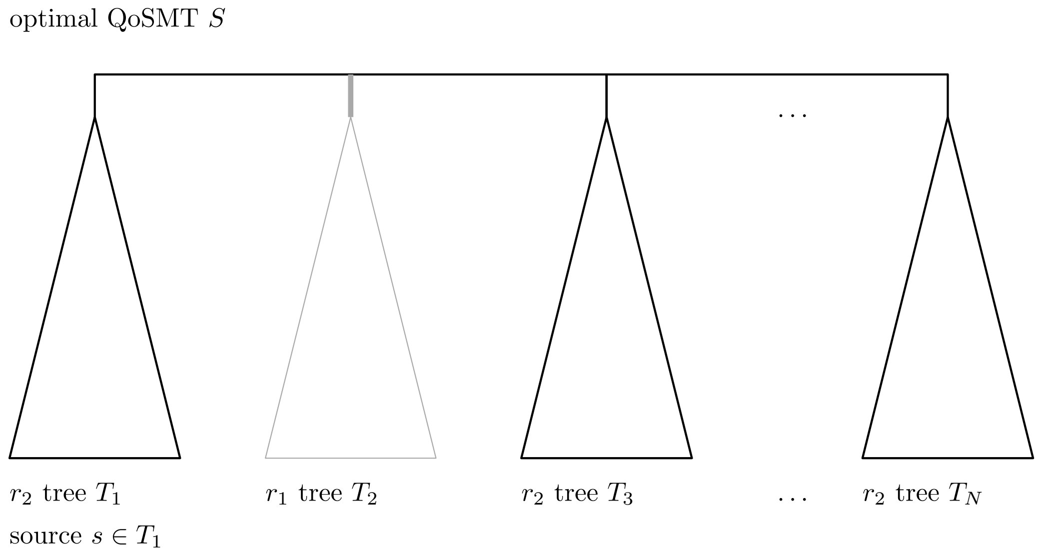

Fig. 4.2: An example for the difficulty in using the shortcutting approach for finding low-cost k-restricted QoSMTs.

Steiner tree problem with an optimal Steiner tree S. We know that for a minimal terminal-spanning tree T', $c(T') \leq 2 \cdot c(S)$ holds. To prove this, take a look at the Steiner tree S. By doubling it, we obtain a cycle S' connecting all terminals, so that $c(S') = 2 \cdot c(S)$. The cycle S' provides an order $(v_1, v_2, \ldots, v_{|T|})$ on the set of terminals according to their order in S' and starting with an arbitrary terminal v_1. (Note that S' may contain several Steiner nodes which are not contained in the sorted terminal set.). We know that S' contains a v_1-v_2-path f. By replacing f with the edge $\{v_1, v_2\}$ we obtain a new cycle with costs at most $c(S')$. By iterating this process, we obtain a cycle S^* including only terminals and with costs $c(S^*) \leq 2 \cdot c(S)$. We know that S^* contains a terminal-spanning tree, so a minimal terminal-spanning tree T' has costs $c(T') \leq c(S^*) \leq 2 \cdot c(S)$ [GP68]. Now in order to create a k-restricted shortcutting tree, we start at an arbitrary terminal $v \in T$. Now we take v and the next $(k-1)$ terminals in the cycle S' as terminal leafs for a component C and make a shortcut from C to the next node v' in S'. This procedure is iterated until all terminals are part of the constructed k-restricted tree. It is possible that we just started at a point where many shortcut edges have a high cost compared to the other ones. We handle this case by repeating this procedure for every possible starting terminal $r \in R$ and taking a result with the lowest costs. We think (but have not proven) that this shortcutting approach would result in an upper bound $b_k \geq \rho_k$ for the k-Steiner ratio, so that b_k is getting arbitrarily close to 1.

To outline why shortcutting can not be used to prove that ρ'_k gets arbitrarily close to 1, we consider the metric closure G of the full tree (which is its own QoSMT S) given in figure 4.2. This tree contains several separated subtrees, each containing k terminals of only one rate. The r_1-subtrees are connected to the rest with an edge with very high costs. Now if we try to use our shortcutting-approach, we always shortcut between an r_1 and an r_2 node. So the high-costing edges are all used with rate r_2. If the rate-ratio r_2/r_1 is very high, the cost-difference between S and the computed tree is also very high. This demonstrates our statement made before.

We have pointed out the difficulty in adapting the approaches for proving bounds on ρ_k to ρ'_k. But note that our example is no counterexample for our assumption that ρ'_k gets arbitrarily close to one, as we could connect the r_1-subtrees to the nearest r_2 node with rate r_1. So it could be possible to work with one of the approaches mentioned above if a good criterion would be found when to separatly handle r_1-nodes and when not to.

5 Conclusion and future work

An important aspect of this work was the question whether the algorithm of Byrka et al. [BGRS13] could be used together with the algorithms of Karpinski et al. [KMOZ05] in order to obtain better approximation guarantees on the QoSMT problem. We have shown that we can use every Steiner tree approximation algorithm as a β-convex α-approximation algorithm, answering the question if we can combine these two algorithms, and obtained the slightly improved bounds of at least 3.769 for the QoSMT problem and at least 1.849 for the the two-rate QoSMT problem. Another question was if this upper bound on β could be improved. We argued why this should not be possible if we want the algorithm to be a β-convex $(\ln(4) + \varepsilon)$-approximation algorithm for $\varepsilon > 0$ as small as possible and if the analysis given by Byrka et al. [BGRS13] is sharp. An interesting question would be if this changes for a willingly higher ε.

In order to use a component-based approach to find approximation algorithms for the (two-rate) QoSMT problem, it would be good if QoSMTs could be arbitrarily approximated with k-restricted trees. We have shown that this is not the case if all components in the k-restricted tree are homogenous. An outstanding question is whether we can approximate QoSMTs with k-restricted trees using components with mixed edge-rates. We have shown that this can be proven or refuted only using graphs which are the metric completion of their own full QoSMT. We further pointed out the difficulties in adapting approches for proving a similar result for ρ_k in Steiner trees. In order to avoid these difficulties, good criteria for deciding whether to treat a specific group of r_1-terminals independently in advance would be needed. Additionally we presented two component-based LPs for the QoSMT problem.

In order to use the randomized rounding approach of Byrka et al. [BGRS13] for the QoSMT problem, another strong component-based LP would be needed. To find such an LP is an open problem and an interesting subject for future work.

Bibliography

[AAS+18] R. Ahmed, P. Angelini, F. D. Sahneh, A. Efrat, D. Glickenstein, M. Gronemann, N. Heinsohn, S. G. Kobourov, R. Spence, J. Watkins, and A. Wolff: Multi-level steiner trees. *SEA*, 17:15:1–15:14, 2018.

[BD97] A. Borchers and D. Du: The k-steiner ratio in graphs. *SIAM Journal on Computing*, 26(3):857–869, 1997.

[BGRS13] Jaroslaw Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità: Steiner tree approximation via iterative randomized rounding. *J. ACM*, 60(1):6:1–6:33, 2013.

[CC08] Miroslav Chlebík and Janka Chlebíková: The Steiner tree problem on graphs: Inapproximability results. *Theoret. Comput. Sci.*, 406(3):207–214, 2008.

[CNS04] Moses Charikar, Joseph (Seffi) Naor, and Baruch Schieber: Resource optimization in QoS multicast routing of real-time multimedia. *IEEE/ACM Trans. Networking*, 12(2):340–348, 2004.

[GP68] Edgar N. Gilbert and Henry O. Pollak: Steiner minimal trees. *SIAM J. Appl. Math.*, 16(1):1–29, 1968.

[Kar72] Richard M. Karp: Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (editors): *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[KMOZ05] Marek Karpinski, Ion I. Mandoiu, Alexander Olshevsky, and Alexander Zelikovsky: Improved approximation algorithms for the quality of service multicast tree problem. *Algorithmica*, 42(2):109–120, 2005.

[Mir96] Prakash Mirchandani: The multi-tier tree problem. *INFORMS J. Comput.*, 8(3):202–218, 1996.

[RZ05] Gabriel Robins and Alexander Zelikovsky: Tighter bounds for graph Steiner tree approximation. *SIAM J. Discrete Math.*, 19(1):122–134, 2005.

[XLD01] Guoliang Xue, Guo-Hui Lin, and Ding-Zhu Du: Grade of service Steiner minimum trees in the Euclidean plane. *Algorithmica*, 31(4):479–500, 2001.

Erklärung

Hiermit versichere ich die vorliegende Abschlussarbeit selbstständig verfasst zu haben, keine anderen als die angegebenen Quellen und Hilfsmittel benutzt zu haben, und die Arbeit bisher oder gleichzeitig keiner anderen Prüfungsbehörde unter Erlangung eines akademischen Grades vorgelegt zu haben.

Würzburg, den August 3, 2018

....
Marie Diana Sieper