

Placing Regenerators in Optical Networks to Satisfy Multiple Sets of Requests^{*} ^{**}

George B. Mertzios¹, Ignasi Sau², Mordechai Shalom³, Shmuel Zaks²

¹ Department of Computer Science, RWTH Aachen University, Aachen, Germany.

`mertzios@cs.rwth-aachen.de`

² Department of Computer Science, Technion, Haifa, Israel.

`{ignasi,zaks}@cs.technion.ac.il`

³ TelHai Academic College, Upper Galilee, 12210, Israel.

`cmshalom@telhai.ac.il`

Abstract. The placement of regenerators in optical networks has become an active area of research during the last years. Given a set of lightpaths in a network G and a positive integer d , regenerators must be placed in such a way that in any lightpath there are no more than d hops without meeting a regenerator. While most of the research has focused on heuristics and simulations, the first theoretical study of the problem has been recently provided in [10], where the considered cost function is the number of *locations* in the network hosting regenerators. Nevertheless, in many situations a more accurate estimation of the real cost of the network is given by the *total* number of regenerators placed at the nodes, and this is the cost function we consider. Furthermore, in our model we assume that we are given a finite set of p possible traffic patterns (each given by a set of lightpaths), and our objective is to place the minimum number of regenerators at the nodes so that each of the traffic patterns is satisfied. While this problem can be easily solved when $d = 1$ or $p = 1$, we prove that for any fixed $d, p \geq 2$ it does not admit a PTAS, even if G has maximum degree at most 3 and the lightpaths have length $\mathcal{O}(d)$. We complement this hardness result with a constant-factor approximation algorithm with ratio $\ln(d \cdot p)$. We then study the case where G is a path, proving that the problem is NP-hard for any $d, p \geq 2$, even if there are two edges of the path such that any lightpath uses at least one of them. Interestingly, we show that the problem is polynomial-time solvable in paths when all the lightpaths share the first edge of the path, as well as when the number of lightpaths sharing an edge is bounded. Finally, we generalize our model in two natural directions, which allows us to capture the model of [10] as a particular case, and we settle some questions that were left open in [10].

Keywords: optical networks, regenerators, overprovisioning, approximation algorithms, hardness of approximation.

^{*} This research was partially supported by the Israel Science Foundation (grant No. 1249/08) and by the British Council (grant No. UKTELHAI09).

^{**} Appeared in *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP)*, Bordeaux, France, July 2010, Volume 2, pp. 333–344. **Best paper award of Track C.**

1 Introduction

1.1 Background

In modern optical networks, high-speed signals are sent through optical fibers using WDM (Wavelength Division Multiplexing) technology. Networks with each fiber typically carrying around 80 wavelengths are operational, whereas networks with a few hundreds wavelengths per fiber are already experimental. As the energy of the signal decreases with the traveled distance, optical amplifiers are required every some fixed distance (a typical value being around 100 km). However, optical amplifiers introduce noise into the signal, so after a certain number of amplifications, the optical signal needs to be regenerated in order to keep the SNR (Signal-to-Noise Ratio) above a specified threshold. In current technology, the signal is regenerated as follows. An ROADM (Reconfigurable Optical Add-Drop Multiplexer) has the capability of inserting/extracting a given number of wavelengths (typically, around 4) to/from the optical fiber. Then, for each extracted wavelength, an optical regenerator is needed to regenerate the signal carried by that wavelength. That is, at a given optical node, one needs as many regenerators as wavelengths one wants to regenerate. See Fig. 1 for a simplified illustration of the aforementioned devices in the case when the network is a path and the fiber carries 3 wavelengths.

The problem of placing regenerators in optical networks has attracted the attention of several recent research works [5, 8, 9, 13, 18, 19, 22, 23]. Mostly, these articles propose heuristics and run simulations in order to reduce the number of regenerators, but no theoretical analysis is presented. Recently, the first theoretical study of the problem has been done by Flammini *et al.* in [10]. In the next paragraph we discuss how our model differs from the one studied in [10].

Nowadays the cost of a regenerator is considerably higher than the cost of an ROADM (as an example, \$160K vs \$50K). Moreover, the regenerator cost is per wavelength, as opposed to ROADM cost that is paid once per several wavelengths. Therefore the *total* number of regenerators seems to be the right cost to minimize. Another possible criterion is to minimize the number of *locations* (that is, the number of nodes) in which optical regenerators are placed. This measure is the one assumed in [10], which makes sense when the dominant




Fig. 1. A simplified optical network: amplifiers introduce noise into the signal, which needs to be regenerated after at most $d = 3$ hops. When the signal is regenerated through an ROADM, a different regenerator is needed for each wavelength.

cost is given by the set-up of new optical nodes, or when the equipment to be placed at each node is the same for all nodes. Nevertheless, the total number of regenerators seem to be a better estimate of the real cost of the network, and therefore we consider this cost in this article.

It is worth mentioning here that when all the connection requests are known a priori, minimizing the number of regenerators is an easy task. Indeed, suppose that the maximum number of hops a lightpath can make without meeting a regenerator is an integer d (in the example of Fig. 1, we have $d = 3$). Then, for each lightpath ℓ , we need to place one regenerator every d consecutive vertices in ℓ , to get an optimal solution.

Unfortunately, when designing a network, it is usually the case that the traffic requests are not known in advance. For instance, the traffic in a given network may change dramatically depending on whether in the foreseeable future an Internet supplier or an email storage server opens or closes a site within the area of the network. In such a situation of uncertain traffic forecast, a common approach in order to minimize capital expenses is to predeploy (or overprovision) resources [12,14,15,17]. That is, the network is designed to satisfy several possible traffic patterns. A similar setting arises in networks in which there are several possible traffic configurations that alternate according to some phenomena, like the weather, the season, an overflow of the capacity of another network, or a breakdown. In that case, the network must be designed so that it can satisfy each of the traffic configurations independently.

In our model, we assume that we are given a finite set of p possible traffic patterns (each given by a set of lightpaths), and our objective is to place the minimum total number of regenerators at the nodes so that each of the traffic patterns is satisfied. That is, the number of regenerators that must be placed at a node of the network is the maximum of the number of regenerators needed by any of the traffic patterns at that node. We aim at minimizing the total number of regenerators placed at the network. We formally define the problem in Section 1.2.

1.2 Definitions

Given an undirected underlying graph $G = (V, E)$ that corresponds to the network topology, a *lightpath* is a simple path in G . That is, we assume that the routing of the requests is given (see [10] for complexity results when the routing of the requests is not given). We also assume that lightpaths sharing an edge use different wavelengths. That is, we deal with optical networks without traffic grooming [2]. The *length* of a lightpath is the number of edges it contains. We consider *symmetric* lightpaths, that is, a lightpath with endpoints u and v consists of a request from u to v and a request from v to u . The *internal vertices* (resp. *edges*) of a lightpath or a path ℓ are the vertices (resp. edges) in ℓ different from the first and the last one. Given an integer d , a lightpath ℓ is *d-satisfied* if there are no d consecutive internal vertices in ℓ without a regenerator. A set of lightpaths is *d-satisfied* if each of its lightpaths is *d-satisfied*. Given p sets of lightpaths L_1, \dots, L_p , with $L_i = \{\ell_{i,j} \mid 1 \leq j \leq x_i\}$, we consider the union of

all lightpaths in the p sets $\cup L_i = \{\ell_{i,j} \mid 1 \leq i \leq p, 1 \leq j \leq x_i\}$. An *assignment* of regenerators is a function $\text{reg} : V \times \cup L_i \rightarrow \{0, 1\}$, where $\text{reg}(v, \ell) = 1$ if and only if a regenerator is used at vertex v by lightpath ℓ .

We study the following problem: given $p \geq 1$ sets of lightpaths, and a distance $d \geq 1$, determine the smallest number of regenerators that d -satisfy each of the p sets. Formally, for two fixed integers $d, p \geq 1$, the optimization problem we study is defined as follows.

(d, p) -TOTAL REGENERATORS $((d, p)$ -TR)

Input: A graph $G = (V, E)$ and p sets of lightpaths $\mathcal{L} = \{L_1, \dots, L_p\}$.

Output: A function $\text{reg} : V \times \cup L_i \rightarrow \{0, 1\}$ such that each lightpath in $\cup L_i$ is d -satisfied.

Objective: Minimize $\sum_{v \in V} \text{reg}(v)$, where
 $\text{reg}(v) = \max_{1 \leq i \leq p} \sum_{\ell \in L_i} \text{reg}(v, \ell)$.

Note that, as mentioned in Section 1.1, in the case $p = 1$ (that is, when there is a single set of requests) the problem is trivially solvable in polynomial time, as the regenerators can be placed for each lightpath independently. The case $d = 1$ is not interesting either, as for each internal vertex $v \in V$ and each $\ell \in \cup L_i$, $\text{reg}(v, \ell) = 1$, so there is only one feasible solution, which is optimal.

1.3 Our contribution

In this article we provide hardness results and approximation algorithms for the (d, p) -TOTAL REGENERATORS problem $((d, p)$ -TR for short). We first prove in Section 3 that for any two fixed integers $d, p \geq 2$, (d, p) -TR does not admit a PTAS unless $P = NP$, even if the underlying graph G has maximum degree at most 3, and the lightpaths have length at most $\lceil \frac{7d}{2} \rceil$. In Section 4 we complement this hardness result with a constant-factor approximation algorithm with ratio $\min\{p, H_{d \cdot p} - 1/2\}$, where $H_n = \sum_{i=1}^n \frac{1}{i}$ is the n -th harmonic number. Section 5 is devoted to the case where the underlying graph is a path. In Section 5.1 we prove that (d, p) -TR is NP-hard in paths for any fixed $d, p \geq 2$, even if there are two edges of the path such that any lightpath uses at least one of them. Interestingly, we show in Section 5.2 that the problem is polynomial-time solvable in paths when all the lightpaths share the first (or the last) edge, as well as when the maximum number of lightpaths sharing an edge is bounded. In Section 6 we generalize the model presented in Section 1.2 in two natural directions. This generalization allows us to capture the model of [10] as a particular case, and to settle some complexity issues that were left open in [10]. (Since we need some further definitions, we defer the precise statement of these results to Section 6.) Finally, in Section 7 we conclude the article and present a number of interesting avenues for further research. We first provide in Section 2 some standard preliminaries. Due to space limitations, almost all proofs are omitted in this extended abstract (except that of Proposition 1); they can be found in [16].

2 Preliminaries

We use standard terminology concerning graphs, complexity, and algorithms; see for instance [7, 11, 21], respectively.

Graphs. All the graphs considered in this article are simple and undirected. Given a graph G we denote by $V(G)$ and $E(G)$ the sets of vertices and edges of G , respectively. If H is a subgraph of G , we denote it by $H \subseteq G$. Given a graph G and $F \subseteq E(G)$, we denote by $G[F]$ the subgraph of G induced by the edges in F together with their endpoints. Given a subset $S \subseteq V(G)$, we define $N_G[S]$ to be the set of vertices of $V(G)$ at distance at most 1 from at least one vertex of S . If $S = \{v\}$, we simply use the notation $N_G[v]$. We also define $N_G(v) = N_G[v] \setminus \{v\}$. The *degree* of a vertex $v \in V(G)$ is defined as $\deg_G(v) = |N_G(v)|$. A graph is *cubic* if all its vertices have degree 3. The *maximum degree* of G is defined as $\Delta(G) = \max_{v \in V(G)} \deg_G(v)$. A *matching* in a graph is a set of disjoint edges, and a *vertex cover* is a set of vertices that contains at least one endpoint of every edge. The *girth* of a graph is the length of a shortest cycle. Given an edge $e = \{u, v\}$, by *subdividing* e we denote the operation of deleting the edge $e = \{u, v\}$, adding a new vertex w , and making it adjacent to both u and v .

Complexity and approximation algorithms. Given an NP-hard minimization problem Π , we say that a polynomial-time algorithm \mathcal{A} is an α -approximation algorithm for Π , with $\alpha \geq 1$, if for any instance of Π , algorithm \mathcal{A} finds a feasible solution with cost at most α times the cost of an optimal solution. For instance, a maximal matching constitutes a 2-approximation algorithm for the MINIMUM VERTEX COVER problem. In complexity theory, the class APX (Approximable) contains all NP-hard optimization problems that can be approximated within a constant factor. The subclass PTAS (Polynomial Time Approximation Scheme) contains the problems that can be approximated in polynomial time within a ratio $1 + \varepsilon$ for *any* fixed $\varepsilon > 0$. In some sense, these problems can be considered to be *easy* NP-hard problems. Since, assuming $P \neq NP$, there is a strict inclusion of PTAS in APX (for instance, $\text{MINIMUM VERTEX COVER} \in \text{APX} \setminus \text{PTAS}$), an APX-hardness result for a problem implies the non-existence of a PTAS unless $P = NP$.

3 Hardness results for general graphs

In this section we prove that, unless $P = NP$, (d, p) -TR does not admit a PTAS for any $d, p \geq 2$, even if the underlying graph G has maximum degree at most 3 and the lightpaths have length $\mathcal{O}(d)$. Before this, we need two technical results to be used in the reductions.

MINIMUM VERTEX COVER is known to be APX-hard in cubic graphs [1]. By a simple reduction, we prove in the following lemma that MINIMUM VERTEX COVER is also APX-hard in a class of graphs with degree at most 3 and high girth, which will be used in the sequel.

Lemma 1. *MINIMUM VERTEX COVER is APX-hard in the class of graphs \mathcal{H} obtained from cubic graphs by subdividing each edge twice.*

Thomassen proved [20] that the edges of any cubic graph can be two-colored such that each monochromatic connected component is a path of length at most 5. In addition, the aforementioned coloring can be found in polynomial time [20]. Note that in such a coloring of a cubic graph, each vertex appears exactly once as an endpoint of a path, and exactly once as an internal vertex of another path. We next show that this result can be easily extended to graphs with maximum degree at most 3.

Lemma 2. *The edges of any graph with maximum degree at most 3 can be two-colored such that each monochromatic connected component is a path of length at most 5.*

We are now ready to announce the main results of this section. For the sake of presentation, we first present in Proposition 1 the result for $d = p = 2$, and then we show in Theorem 1 how to extend the reduction to any fixed $d, p \geq 2$.

Proposition 1. *$(2, 2)$ -TR does not admit a PTAS unless $P = NP$, even if G has maximum degree at most 3 and the lightpaths have length at most 7.*

Proof: The reduction is from MINIMUM VERTEX COVER (VC for short) in the class of graphs \mathcal{H} obtained from cubic graphs by subdividing each edge twice, which does not admit a PTAS by Lemma 1 unless $P = NP$. Note that by construction any graph in \mathcal{H} has girth at least 9. Given a graph $H \in \mathcal{H}$ as instance of VERTEX COVER, we proceed to build an instance of $(2, 2)$ -TR. We set $G = H$, so G has maximum degree at most 3.

To define the two sets of lightpaths L_1 and L_2 , let $\{E_1, E_2\}$ be the partition of $E(H)$ given by the two-coloring of Lemma 2. Therefore, each connected component of $H[E_1]$ and $H[E_2]$ is a path of length at most 5. Each such path in $H[E_1]$ (resp. $H[E_2]$) will correspond to a lightpath in L_1 (resp. L_2), which we proceed to define. A key observation is that, as the paths of the two-coloring have length at most 5, if any endpoint v of a such path P had one neighbor in $V(P)$, it would create a cycle of length at most 6, a contradiction to the fact that the girth of H is at least 9. Therefore, as the vertices of H have degree 2 or 3, any endpoint v of a path P has at least one neighbor in $V(H) \setminus V(P)$.

We are now ready to define the lightpaths. Let P be a path with endpoints u, v , and let u' (resp. v') be a neighbor of u (resp. v) in $V(H) \setminus V(P)$, such that $u' \neq v'$ (such distinct vertices u', v' exist by the above observation and by the fact that H has girth at least 9). The lightpath associated with P consists of the concatenation of $\{u', u\}$, P , and $\{v, v'\}$. Therefore, the length of each lightpath is at most 7. This completes the construction of the instance of $(2, 2)$ -TR. Observe that since we assume that $d = 2$, regenerators must be placed in such a way that all the internal edges of a lightpath (that is, all the edges except the first and the last one) have a regenerator in at least one of their endpoints. We can assume without loss of generality that no regenerator serves at the endpoints

of a lightpath, as the removal of such regenerators does not alter the feasibility of a solution. Note that in our construction, each vertex of G appears as an internal vertex in at most two lightpaths, one (possibly) in L_1 and the other one (possibly) in L_2 , so we can assume that $\text{reg}(v) \leq 1$ for any $v \in V(G)$.

We now claim that $\text{OPT}_{\text{VC}}(H) = \text{OPT}_{(2,2)\text{-TR}}(G, \{L_1, L_2\})$.

Indeed, let first $S \subseteq V(H)$ be a vertex cover of H . Placing one regenerator at each vertex belonging to S defines a feasible solution to $(2,2)$ -TR in G with cost $|S|$, as at least one endpoint of each internal edge of each lightpath contains a regenerator. Therefore, $\text{OPT}_{\text{VC}}(H) \geq \text{OPT}_{(2,2)\text{-TR}}(G, \{L_1, L_2\})$.

Conversely, suppose we are given a solution to $(2,2)$ -TR in G using r regenerators. Since E_1 and E_2 are a partition of $E(G) = E(H)$ and the set of internal edges of the lightpaths in L_1 (resp. L_2) is equal to E_1 (resp. E_2), the regenerators placed at the endpoints of the internal edges of the lightpaths constitute a vertex cover of H of size at most r . Therefore, $\text{OPT}_{\text{VC}}(H) \leq \text{OPT}_{(2,2)\text{-TR}}(G, \{L_1, L_2\})$.

Summarizing, since $\text{OPT}_{\text{VC}}(H) = \text{OPT}_{(2,2)\text{-TR}}(G, \{L_1, L_2\})$ and any feasible solution to $\text{OPT}_{(2,2)\text{-TR}}(G, \{L_1, L_2\})$ using r regenerators defines a vertex cover of H of size at most r , the existence of a PTAS for $(2,2)$ -TR would imply the existence of a PTAS for VERTEX COVER in the class of graphs \mathcal{H} , which is a contradiction by Lemma 1, unless $\text{P} = \text{NP}$. \square

Theorem 1. *(d,p) -TR does not admit a PTAS for any $d \geq 2$ and any $p \geq 2$ unless $\text{P} = \text{NP}$, even if the underlying graph G satisfies $\Delta(G) \leq 3$ and the lightpaths have length at most $\lceil \frac{7d}{2} \rceil$.*

4 Approximation algorithms for general graphs

We have seen in Section 3 that (d,p) -TR does not admit a PTAS for $d, p \geq 2$ unless $\text{P} = \text{NP}$. In this section we complement this result with a constant-factor approximation algorithm for (d,p) -TR in general graphs.

Theorem 2. *For any fixed $d, p \geq 2$, there is a polynomial-time approximation algorithm for the (d,p) -TR problem with ratio $\min\{p, H_{d,p} - 1/2\}$, where $H_{d,p} = \sum_{i=1}^{d \cdot p} \frac{1}{i}$.*

Note that for big d, p , $H_{d,p} \approx \ln d + \ln p + 1/2$, so comparing both approximation ratios, we have that $p < \ln d + \ln p$ when $d = \Omega(2^p)$.

5 The case of the path

In this section we focus on the case where the network topology is a path, which is one of the most important topologies in real networks, as well as one of the most natural and apparently simplest underlying graphs to study. Clearly, hardness results obtained for this topology will carry over all topologies. We first present in Section 5.1 an NP-hardness result, and then we present in Section 5.2 polynomial-time optimal algorithms for two families of instances.

5.1 Hardness result

The model of [10] turns out to be polynomial-time solvable when the underlying topology is a tree. Surprisingly enough, in this section we show that our model remains NP-hard even if the network is a path, for any $d \geq 2$ and any $p \geq 2$.

Theorem 3. *(d, p) -TR is NP-hard in paths for any $d, p \geq 2$, even if each vertex is the endpoint of at most 10 lightpaths and there are two edges of the path such that any lightpath uses at least one of them.*

5.2 Polynomial-time solvable cases

In this section we present polynomial-time optimal algorithms in path networks for two restricted sets of instances, namely when all the lightpaths go through a common edge, and when the *load* of the path (that is, the maximum number of lightpaths in any set L_i crossing an edge of the path) is bounded by a logarithmic function of the input size.

Edge instances. In an *edge instance* there is an edge $e \in E(G)$ that is used by all the lightpaths. Note that the instance built in the reduction used in the proof of Theorem 3 contains two edges intersecting all the lightpaths.

Proposition 2. *For any fixed $d, p \geq 2$, there is a polynomial-time algorithm solving the (d, p) -TR problem for edge instances in a path where all the lightpaths share the first edge.*

Bounded load. From Theorem 3 it follows that (d, p) -TR remains hard in paths even if each vertex is the endpoint of at most 10 lightpaths. It turns out that if we further impose that not only the number of lightpaths per vertex is bounded, but also the load of the path is bounded by an appropriate function of the size of the instance, then the problem is solvable in polynomial time. Intuitively, this special case of instances is in the opposite extreme of the edge instances, where there is an edge with unbounded load.

Proposition 3. *For any fixed $d, p \geq 2$, (d, p) -TR is polynomial-time solvable in paths if the load is $\mathcal{O}\left(\frac{\log |I| - \log p}{2p \cdot \log d}\right) = \mathcal{O}(\log |I|)$, where $|I|$ is the size of the instance.*

6 More general settings

In this section we generalize the (d, p) -TR problem in two natural directions. Namely, in Section 6.1 we allow the number p of traffic patterns to be unbounded, and in Section 6.2 we introduce a parameter k that bounds the number of regenerators that can be placed at a vertex. Technologically, the latter constraint captures the fact of having a bounded number of ROADMs per vertex, as the number of wavelengths (and therefore, the number of regenerators) an ROADM can handle is usually not too big (see Section 1.1).

6.1 Unbounded number of sets of lightpaths

If p is part of the input, then (d, p) -TR contains as a particular case the model studied in [10] (the so-called *location* problem, denoted $\text{RPP}/\infty/+$ in [10]). Indeed, if each set of lightpaths consists of a single lightpath (that is, when p is the number of lightpaths), then the objective is to place the minimum number of regenerators such that each lightpath is satisfied. Therefore, the hardness results stated in [10] also apply to this more general setting, in particular an approximation lower bound of $\Omega(\log(d \cdot p))$ unless NP can be simulated in subexponential time. Note that this hardness bound matches the approximation ratio given by Theorem 2. Nevertheless, note also that the approximation algorithm presented in Theorem 2 runs in polynomial time only for bounded p .

We now reformulate the problem studied in [10] using our terminology. Let $d \geq 1$ be a fixed integer.

d -REGENERATORS LOCATION (d -RL)

Input: An undirected graph $G = (V, E)$ and a set of lightpaths L .

Output: A function $\text{reg} : V \times L \rightarrow \{0, 1\}$ such that each lightpath $\ell \in L$ is d -satisfied.

Objective: Minimize $\sum_{v \in V} \text{reg}(v)$, where $\text{reg}(v) = \max_{\ell \in L} \text{reg}(v, \ell)$.

Note that in the above problem, $\text{reg}(v) \in \{0, 1\}$. We now focus on the case $d = 2$ of d -RL.

Remark 1 *Given an instance of 2-RL in a graph G , the problem can be reduced to a MINIMUM VERTEX COVER problem in a subgraph of G . Indeed, given a set of lightpaths L , remove the first and the last edge of each lightpath, and let H be the subgraph of G defined by the union of the edges in the modified lightpaths. It is then clear that the minimum number of regenerators to 2-satisfy all the lightpaths in L equals the size of a minimum vertex cover of H .*

By Remark 1 and König's theorem [7], it follows that 2-RL can be solved in polynomial time in bipartite graphs. This result extends the results of [10] for $d = 2$, where it is proved that for any $d \geq 2$, d -RL is polynomial-time solvable in trees and rings. Finally, it also follows from Remark 1 that 2-RL admits a PTAS in planar graphs [4] and, more generally, in any family of minor-free graphs [6].

6.2 Bounded number of regenerators per vertex

From a technological point of view, it makes sense to introduce a parameter k that limits the number of regenerators that can be used at a single vertex. Adding this restriction to the d -RL problem, we get the following problem, which is actually the so-called k -location problem and denoted $\text{RPP}/k/+$ in [10].

Again, we restate the problem using our terminology. Let $d, k \geq 1$ be two fixed integers.

(d, k) -REGENERATORS LOCATION $((d, k)$ -RL)

Input: An undirected graph $G = (V, E)$ and a set of lightpaths L .

Output: A function $\text{reg} : V \times L \rightarrow \{0, 1\}$ such that each lightpath $\ell \in L$ is d -satisfied and $\text{reg}(v) \leq k$, where $\text{reg}(v) = \sum_{\ell \in L} \text{reg}(v, \ell)$.

Objective: Minimize $|\{v \in V \mid \text{reg}(v) > 0\}|$.

We now resolve two questions that were left open in [10]. Namely, it is proved in [10] that given an instance of $(3, 1)$ -RL, it is NP-complete to decide whether there exists a feasible solution for it, which in particular implies that the $(3, 1)$ -RL problem itself is NP-hard to approximate within any ratio. In the following we prove that, surprisingly, the situation changes for $d = 2$ and $k = 1$. More precisely, it is in P to decide whether there exists a feasible solution for an instance of $(2, 1)$ -RL, while finding an optimal one is NP-hard.

Proposition 4. *Given an instance of $(2, 1)$ -RL, it can be decided in polynomial time whether there exists a feasible solution for it, while the $(2, 1)$ -RL problem itself (that is, finding an optimal solution) is NP-hard.*

7 Conclusions and further research

In this article we presented a theoretical study of the problem of placing regenerators in optical networks, so that on each lightpath we must put a regenerator every at most d hops. The cost is the total number of regenerators. We considered the case when p possible traffic patterns are given (each by a set of lightpaths), and the objective is to place the minimum number of regenerators satisfying each of these patterns. This setting arises naturally when designing real networks under uncertain traffic forecast. The problem is called (d, p) -TOTAL REGENERATORS problem, or (d, p) -TR for short. We now summarize our results and propose a number of lines for further research.

We proved that for any fixed $d, p \geq 2$, (d, p) -TR does not admit a PTAS unless $P = NP$, even if the network topology has maximum degree at most 3, by reduction from MINIMUM VERTEX COVER in cubic graphs. It would be interesting to determine which is the explicit approximation lower bound given by Theorem 1. The recent results of Austrin *et al.* [3] about the hardness of MINIMUM VERTEX COVER in graphs of bounded degree may shed some light on this question. We provided an approximation algorithm for (d, p) -TR with constant ratio $\ln(d \cdot p)$, by reducing it to MINIMUM SET COVER. Finding a polynomial-time approximation algorithm matching the hardness lower bound given by Theorem 1 seems to be a challenging task.

We proved that (d, p) -TR is NP-hard in paths for any $d, p \geq 2$, by reduction from the problem of whether the edges of a tripartite graph can be partitioned into triangles. It is easy to see that the proof of Theorem 3 can be adapted to the case when the network is a ring. In fact, the optimization version where the

objective is to find the maximum number of edge-disjoint triangles is APX-hard in tripartite graphs [2]. Nevertheless, the proof of Theorem 3 does not allow – at least, without major modifications – to prove that (d, p) -TR does not admit a PTAS in paths, although we believe that this is indeed the case. Therefore, the existence of a PTAS for (d, p) -TR in paths, trees, rings, or even planar graphs, remains open.

The NP-hardness result for paths holds even if there are two edges of the path such that each lightpath uses at least one of them. In order to better understand what makes the problem hard, we proved that when all lightpaths use the first (or the last) edge of the path, then (d, p) -TR becomes polynomial-time solvable for any $d, p \geq 2$. Between these two cases, it only remains to settle the complexity of the case when the edge shared by all lightpaths is an internal edge of the path, which could be polynomial or NP-hard.

Still in the path, but in the opposite extreme of the type of instances, we also proved that (d, p) -TR can be solved in polynomial time when the maximum number of lightpaths using an edge is logarithmically bounded by the size of the instance. It may be possible to extend our dynamic programming approach to trees with instances having this property, and even to graphs with bounded treewidth.

We generalized our model by allowing the number of sets of lightpaths to be unbounded, and by introducing a parameter k that bounds the number of regenerators that can be placed at a node. This way, the model studied in [10] becomes a particular case. We settled several complexity questions that were left open in [10] concerning the case $k = 1$ and $d = 2$. As future work, it seems to be of high importance to consider the parameter k in the original statement of our (d, p) -TR problem.

Acknowledgment. We would like to express our gratitude to Ori (Ornan) Gerstel from CISCO, Israel, for sharing with us the new trends concerning the placement of regenerators in optical networks.

References

1. P. Alimonti and V. Kann. Some APX-completeness results for cubic graphs. *Theoretical Computer Science*, 237(1-2):123–134, 2000.
2. O. Amini, S. Pérennes, and I. Sau. Hardness and Approximation of Traffic Grooming. *Theoretical Computer Science*, 410(38-40):3751–3760, 2009.
3. P. Austrin, S. Khot, and M. Safra. Inapproximability of Vertex Cover and Independent Set in Bounded Degree Graphs. In *Proc. of the 24th IEEE Computational Complexity Conference (CCC)*, pages 74–80, 2009.
4. B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.
5. S. Chen and S. Raghavan. The regenerator location problem. In *Proc. of the International Network Optimization Conference (INOC)*, 2007. Full version to appear in *Networks*.
6. E. Demaine, M. Hajiaghayi, and K.-i. Kawarabayashi. Algorithmic Graph Minor Theory: Decomposition, Approximation and Coloring. In *Proc. of the 46th Annual*

- IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 637–646, 2005.
7. R. Diestel. *Graph Theory*. Springer-Verlag, 2005.
 8. R. Fedrizzi, G. M. Galimberti, O. Gerstel, G. Martinelli, E. Salvadori, C. V. Saradhi, A. Tanzi, and A. Zanardi. A Framework for Regenerator Site Selection Based on Multiple Paths. In *Proceedings of IEEE/OSA Conference on Optical Fiber Communications (OFC)*, 2010. To appear.
 9. R. Fedrizzi, G. M. Galimberti, O. Gerstel, G. Martinelli, E. Salvadori, C. V. Saradhi, A. Tanzi, and A. Zanardi. Traffic Independent Heuristics for Regenerator Site Selection for Providing Any-to-Any Optical Connectivity. In *Proc. of IEEE/OSA Conference on Optical Fiber Communications (OFC)*, 2010. To appear.
 10. M. Flammini, A. Marchetti-Spaccamela, G. Monaco, L. Moscardelli, and S. Zaks. On the complexity of the regenerator placement problem in optical networks. In *Proc. of the 21st Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 154–162, 2009.
 11. M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
 12. O. Gerstel and H. Raza. Predeployment of Resources in Agile Photonic Networks. *IEEE/OSA Journal of Lightwave Technology*, 22(10):2236–2244, 2004.
 13. S. W. Kim and S. W. Seo. Regenerator placement algorithms for connection establishment in all-optical networks. *Communications, IEE Proceedings*, 148(1):25–30, 2001.
 14. S. K. Korotki. An Overview of the Global Network Expectation Model. In *Proceedings of IEEE/OSA Conference on Optical Fiber Communications (OFC)*, 2004.
 15. S. K. Korotki and K. N. Oikonomou. Scaling of Most-Likely Traffic Patterns of Hose- and Cost-Constrained Ring and Mesh Networks. In *Proceedings of IEEE/OSA Conference on Optical Fiber Communications (OFC)*, 2006.
 16. G. B. Mertzios, I. Sau, M. Shalom, and S. Zaks. Placing Regenerators in Optical Networks to Satisfy Multiple Sets of Requests. Technical Report CS-2010-07, Technion, Israel Institute of Technology, 2010. Available at www.cs.technion.ac.il/users/wwwb/cgi-bin/tr-info.cgi?2010/CS/CS-2010-07.
 17. C. S. R. Murthy and S. Sinha. Information Theoretic Approach to Traffic Adaptive WDM Networks. *IEEE/ACM Transactions on Networking*, 13(4):881–894, 2005.
 18. S. Pachnicke, T. Paschenda, and P. M. Krummrich. Physical Impairment Based Regenerator Placement and Routing in Translucent Optical Networks. In *Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference (Optical Society of America, paper OWA2)*, 2008.
 19. K. Sriram, D. Griffith, R. Su, and N. Golmie. Static vs. dynamic regenerator assignment in optical switches: models and cost trade-offs. In *Workshop on High Performance Switching and Routing (HPSR)*, pages 151–155, 2004.
 20. C. Thomassen. Two-Coloring the Edges of a Cubic Graph Such That Each Monochromatic Component Is a Path of Length at Most 5. *J. Comb. Theory, Ser. B*, 75(1):100–109, 1999.
 21. V. V. Vazirani. *Approximation algorithms*. Springer-Verlag, 2001.
 22. X. Yang and B. Ramamurthy. Dynamic routing in translucent WDM optical networks. In *Proc. of the IEEE International Conference on Communications (ICC)*, pages 955–971, 2002.
 23. X. Yang and B. Ramamurthy. Sparse Regeneration in Translucent Wavelength-Routed Optical Networks: Architecture, Network Design and Wavelength Routing. *Photonic Network Communications*, 10(1):39–53, 2005.