Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ANDREAS BLEY

On the Complexity of Vertex–Disjoint Length–Restricted Path Problems

Preprint SC 98-20 (August 1998)

On the Complexity of Vertex–Disjoint Length–Restricted Path Problems

Andreas Bley

August 11, 1998

Abstract

Let $G = (V, E)$ be a simple graph and s and t be two distinct vertices of G. A path in G is called ℓ-bounded for some $\ell \in \mathbb{N}$, if it does not contain more than ℓ edges. We study the computational complexity of approximating the optimum value for two optimization problems of finding sets of vertex-disjoint ℓ-bounded s, t-paths in G.

First, we show that computing the maximum number of vertex-disjoint ℓ-bounded s, t-paths is $\mathcal{APX}$–complete for any fixed length bound $\ell \geq 5$.

Second, for a given number $k \in \mathbb{N}$, $1 \leq k \leq |V| - 1$, and non-negative weights on the edges of G, the problem of finding k vertex-disjoint ℓ-bounded s, t-paths with minimal total weight is proven to be $\mathcal{NPO}$–complete for any length bound $\ell \geq 5$. Furthermore, we show that, even if G is complete, it is $\mathcal{NP}$–hard to approximate the optimal solution value of this problem within a factor of $2^{\langle\phi\rangle^\epsilon}$ for any constant $0 < \epsilon < 1$, where $\langle\phi\rangle$ denotes the encoding size of the given problem instance ϕ.

We prove that these results are tight in the sense that for lengths $\ell \leq 4$ both problems are polynomially solvable, assuming that the weights satisfy a generalized triangle inequality in the weighted problem.

All results presented also hold for directed and non-simple graphs. For the analogous problems where the path length restriction is replaced by the condition that all paths must have length equal to ℓ or where vertex-disjointness is replaced by edge-disjointness we obtain similar results.

Keywords: disjoint paths, length bounded paths, approximation, reducibility, completeness

Mathematical Subject Classification (1991): 68Q25, 90C27, 05C38, 05C40

1 Introduction

Due to the trend to high-capacity but sparse networks, survivability has become a major issue in the design of telecommunication networks over the last decades. Motivated by this, there has been a large amount of research in this area. Many different models of network survivability as well as lots of different cost and capacity models have been studied; for a survey see [AGW97]. In a variety of these models survivability is enforced by demanding a certain number of disjoint paths between each pair of communicating nodes. Usually, these paths have to fulfill additional requirements, e. g., transition time or router number limits, that can be modeled by path length restrictions.

We study two problems related to such survivability models in this article, namely finding maximum sets of vertex-disjoint length-restricted paths between two communicating nodes and finding weight-minimal such sets containing a given number of paths.

For both problems the input consists of a simple undirected graph $G = (V, E)$ and two distinct vertices s and t of G. We call an s,t-path in G *ℓ-bounded* for some length restriction $\ell \in \mathbb{N}$ if it contains ℓ or less edges.

The first problem studied in this article, the Max-Disjoint-ℓ-Bounded-Paths problem denoted by $\mathsf{MDBP}(\ell)$, is the problem of finding the maximum number of vertex-disjoint ℓ-bounded paths between s and t in G.

The second problem considered here is the Weighted-Disjoint-ℓ-Bounded-Paths problem, denoted by $\mathsf{WDBP}(\ell)$. In this problem we are additionally given non-negative weights on the edges of G and a number $k \in \mathbb{N}$,$1 \leq k \leq |V| - 1$. The aim is to find a set of k vertex-disjoint ℓ-bounded s,t-paths with minimum total weight.

It was shown by Lovàsz, Neumann-Lara, and Plummer in [LNLP78] that for $\ell \geq 5$ there are no (tight) Menger-type results for vertex-disjoint paths of bounded length. The ratio between the maximum number of vertex-disjoint ℓ-bounded s,t-paths and the number of vertices that have to be removed from G to increase the distance between s and t above ℓ or disconnect them is two already for $\ell = 5$ and increases with the size of G. Itai, Perl, and Shiloach proved in [IPS82] that the problem of deciding whether there exist k vertex-disjoint ℓ-bounded s,t-paths in G is $\mathcal{NP}$–complete for $\ell \geq 5$ and polynomially solvable for $\ell \leq 4$. A two phase primal heuristic solving the Max-Disjoint-ℓ-Bounded-Paths problem to optimality for $\ell \leq 4$ but without approximation guarantee for length bounds $\ell \geq 5$ (except the trivial factor of $\ell - 1$), was proposed by Perl and Ronen in [PR84].

In this article we study the computational complexity of approximating the optimal solution values for the two optimization problems defined above. In Sections 2.1 and 2.2 we prove that for $\ell \geq 5$ the Max-Disjoint-ℓ-Bounded-Paths and the Weighted-Disjoint-ℓ-Bounded-Paths problem are $\mathcal{APX}$–complete and $\mathcal{NPO}$–complete [Pap94, BC94], respectively. Furthermore, we show that it is $\mathcal{NP}$–complete to approximate the latter for $\ell \geq 5$ within a factor of $2^{\langle G \rangle^\epsilon}$ for any constant $\epsilon > 0$ or within $2^{\langle \phi \rangle^\epsilon}$ for any $0 < \epsilon < 1$, even if G is a complete graph. $\langle G \rangle$ and $\langle \phi \rangle$ denote the encoding size of the graph G and of the entire $\mathsf{WDBP}(\ell)$ instance, respectively. According to the definition of the class $\mathcal{NPO}$ of $\mathcal{NP}$ optimization problems, we assume that there is no solution of total weight zero for $\mathsf{WDBP}(\ell)$. If the graph G is not necessarily complete, it is $\mathcal{NP}$–complete to approximate the problem for $\ell \geq 5$ within $2^{\langle \phi \rangle^\epsilon}$ even for any fixed non-negative ϵ, i.e., $\mathsf{WDBP}(\ell)$ is not in $\mathtt{exp}$–$\mathcal{APX}$ [CKST96] for $\ell \geq 5$. For length bounds $\ell \leq 4$ the Max-Disjoint-ℓ-Bounded-Paths problem and the semi-metric Weighted-Disjoint-ℓ-Bounded-Paths problem are proven to be polynomially solvable in Section 3. In the last section we extended our results to the corresponding problems where the path length restriction is replaced by the condition that all paths must have length equal to ℓ and to the corresponding edge-disjoint paths problems.

2 Inapproximability Results

In this section we will show that both problems introduced in the previous section are computationally hard to approximate for length bounds $\ell \geq 5$.

It is not in the scope of this article to introduce all concepts of computational complexity theory needed in the following. We assume the reader to be familiar with the basic definitions and techniques in this area. For a formal definition of the complexity classes, an introduction to the reduction techniques, and the basic structural results used in this article we refer the reader to [OM87, PY88, PY91, BC94, Pap94, CKST96]. A regularly updated compendium of the current knowledge on the complexity of specific problems can be found in [CK95].

2.1 Max-Disjoint-5-Bounded-Paths

In [IPS82] Itai, Perl, and Shiloach proved that the problem of deciding whether a graph G contains a given number k of vertex-disjoint ℓ-bounded s,t-paths is $\mathcal{NP}$–complete for $\ell \geq 5$, presenting a reduction from a special version of 3-Sat. Although their construction is sufficient to prove the $\mathcal{NP}$–completeness of the decision problem, it cannot be used to show the $\mathcal{APX}$–completeness of $\mathsf{MDBP}(\ell)$. The construction only transforms feasible solutions of the above decision problem to feasible solutions of 3-Sat. To prove the $\mathcal{APX}$–completeness of $\mathsf{MDBP}(\ell)$, we need a reduction that also maps approximate solutions of this problem to approximate solutions of Max-3-Sat. Using gadgets similar to those introduced by Itai et. al., we construct such approximation preserving reductions to prove the $\mathcal{APX}$–completeness of $\mathsf{MDBP}(\ell)$ as well as the $\mathcal{NPO}$–completeness of $\mathsf{WDBP}(\ell)$ for $\ell \geq 5$.

In the following we show that approximating the maximum number of vertex-disjoint ℓ-bounded s,t-paths is $\mathcal{APX}$–complete for $\ell \geq 5$. First of all, it is easy to see that for any *fixed* length bound ℓ the Max-Disjoint-ℓ-Bounded-Paths problem $\mathsf{MDBP}(\ell)$ is in $\mathcal{APX}$.

Lemma 2.1 $\mathsf{MDBP}(\ell)$ *is in* $\mathcal{APX}$ *for each* $\ell \in \mathbb{N}$.

Proof. Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a set of vertex-disjoint ℓ-bounded s,t-paths, such that there exists no further ℓ-bounded s,t-path in G that is disjoint to all $P \in \mathcal{P}$. Obviously, such a set of paths can be found in polynomial time by repeatedly applying a depth-bounded depth first search starting in s. Let $\mathcal{P}' = \{P'_1, \dots, P'_{k'}\}$ be a maximum set of vertex-disjoint ℓ-bounded s,t-paths. Since each $P'_i \in \mathcal{P}'$, except maybe the direct edge st, intersects with at least one $P_j \in \mathcal{P}$ and all paths in $\mathcal{P}'$ are vertex-disjoint, we have $|\mathcal{P}'| \leq |\{\text{internal vertices in } \mathcal{P}|\} \leq (\ell-1)|\mathcal{P}|$ if G does not contain the direct edge st, and $|\mathcal{P}'| - 1 \leq |\{\text{internal vertices in } \mathcal{P}|\} \leq (\ell-1)(|\mathcal{P}|-1)$ otherwise. Hence, $\mathcal{P}$ is an $(\ell-1)$-approximate solution. □

Second, to show that $\mathsf{MDBP}(\ell)$ is $\mathcal{APX}$–complete for all length bounds $\ell \geq 5$, we will prove that $\mathsf{MDBP}(5)$ is $\mathcal{APX}$–complete. Then, for $\ell > 5$ the $\mathcal{APX}$–completeness follows by replacing the edges emanating from s in the graph constructed in this reduction by a simple path of length $\ell - 4$. We prove the $\mathcal{APX}$–completeness of $\mathsf{MDBP}(5)$ by constructing an L–reduction[PY88, PY91] from Max-3-Sat(3), which is known to be $\mathcal{APX}$–complete due to Papadimitriou and Yannakakis [PY88, PY91] and Ausiello et. al. [ACG+98].

The Max-3-Sat problem is, given a set X of boolean variables and a collection C of disjunctive clauses of at most three literals per clause, to find a truth assignment for X satisfying the maximum number of clauses. In our proof we use a restricted version Max-3-Sat(3) of Max-3-Sat, where each variable occurs at most three times. Even with this restriction the Max-3-Sat problem is still $\mathcal{APX}$–complete, see [ACG+98].

We need the following simple lemma to prove that the reduction we are going to construct below is approximation preserving

Lemma 2.2 *Let* ϕ *be an instance of* Max-3-Sat *with* r *clauses and let* $Opt(\phi)$ *denote the value of an optimal solution of* ϕ*. Then* $Opt(\phi) \geq \frac{1}{2} r$.

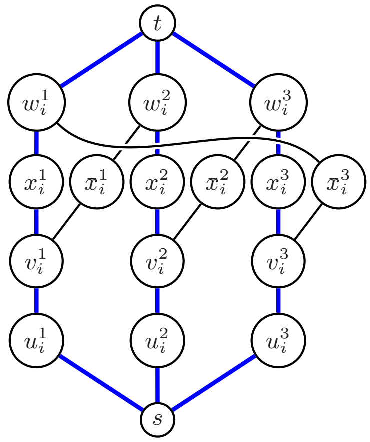

Figure 1: Variable graph G_i

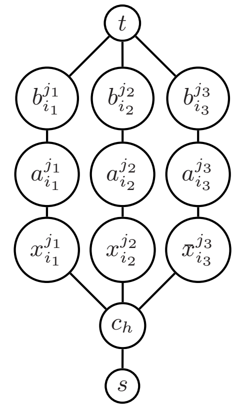

Figure 2: Clause graph H_h

Proof. Let x be an arbitrary truth assignment for the variables in ϕ. If x satisfies less than $r/2$ clauses, then the assignment $-$x, where each variable in X is negated, satisfies more than $r/2$ clauses. On the other hand, an optimal assignment cannot satisfy more than r clauses. □

Theorem 2.3 MDBP(5) *is $\mathcal{APX}$–complete.*

Proof. The proof consists of two parts. First, for a given instance ϕ of Max-3-Sat(3), where q and r are the number of variables and the number of clauses in ϕ, respectively, we construct a graph G that contains $3q + r$ vertex-disjoint 5-bounded s,t-paths if and only if all clauses in ϕ are satisfiable. In the second part, we show that there are at least $3q$ vertex-disjoint 5-bounded s,t-paths in G and that, for each truth assignment that satisfies r' clauses, there is a set of $3q + r'$ such paths in G, and vice versa.

Let ϕ be the given instance of Max-3-Sat(3) with variables $\mathrm{x}_1, \ldots, \mathrm{x}_q$ and clauses $C_1, \ldots, C_r$. For notational convenience we assume that the occurrences of each variable x_i are denoted by x_i^j, $0 \le j \le 2$ (or $0 \le j \le 1$ if x_i occurs only two times in ϕ).

The graph G is constructed by clipping together $q+r$ gadget graphs, one for each variable and one for each clause in ϕ.

For each variable x_i in ϕ, we construct a "variable graph" G_i as shown in Figure 1. This graph contains two s,t-paths for each $0 \le j \le 2$, namely $P_i^j = [s, u_i^j, v_i^j, x_i^j, w_i^j, t]$ and $\bar{P}_i^j = [s, u_i^j, v_i^j, \bar{x}_i^j, w_i^{j'}, t]$, with $j' = (j + 1 \mod 3)$. Obviously, these P_i^j and $\bar{P}_i^j$ are the only 5-bounded s,t-paths in G_i. Furthermore, all three P_i^j and all three $\bar{P}_i^j$ paths are vertex-disjoint, respectively, and any collection of three vertex-disjoint 5-bounded s,t-paths in G_i must consist either of all P_i^j paths or of all $\bar{P}_i^j$ paths.

We interpret choosing path P_j^i in G_i as setting x_i^j to *true* in ϕ (if this occurrence of x_i exists in ϕ) and choosing path $\bar{P}_i^j$ as setting it to *false*. With this interpretation, having three vertex-disjoint 5-bounded s,t-paths in G_i corresponds to setting all occurrences of variable x_i in ϕ uniformly.

For each clause C_h in ϕ, we introduce a "clause graph" H_h as shown in Figure 2. This graph contains the vertices s, t and c_h. Furthermore, for each x_i^j that occurs in clause C_h, this graph contains the vertices a_i^j, b_i^j, and either vertex x_i^j or vertex $\bar{x}_i^j$ (of G_i), depending whether x_i^j occurs negated or un-negated in C_h. If it occurs negated, then H_h contains vertex

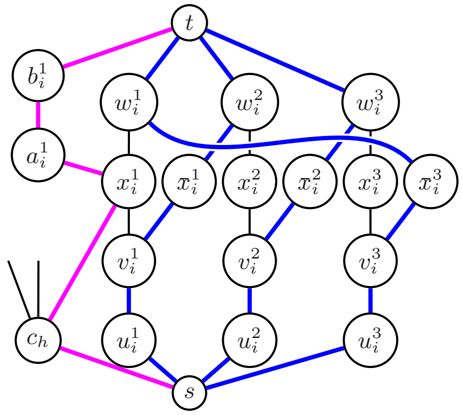

Figure 3: Clipping of the G_i and H_h subgraphs

x_i^j, if it occurs un-negated H_h contains vertex $\bar{x}_i^j$. For each x_i^j occurring in C_h, these vertices form an s,t-path of length five, either $Q_i^j := [s, c_h, x_i^j, a_i^j, b_i^j, t]$ or $\bar{Q}_i^j := [s, c_h, \bar{x}_i^j, a_i^j, b_i^j, t]$. Since all these paths have the vertex c_h in common, there is at most one "vertex-disjoint" s,t-path within each H_h, which has length five.

We interpret choosing path Q_i^j (or $\bar{Q}_i^j$) as satisfying clause C_h by setting x_i^j to *false* (or *true*, respectively). Note that, as long as we have three vertex-disjoint 5-bounded paths within G_i, this interpretation agrees with that of choosing path P_i^j or $\bar{P}_i^j$ in G_i.

Finally, all these gadget graphs are clipped together in s, t, and the appropriate x_i^j and $\bar{x}_i^j$ vertices, see Figure 3.

Formally, the instance $\psi = \psi(\phi)$ of MDBP(5) is defined as the problem of finding the maximum number of vertex-disjoint 5-bounded paths from s to t in G. It is clear that the above construction can be accomplished in polynomial time.

To transform the solutions of ψ back to solutions of ϕ preserving the approximation quality, we need the following two lemmas.

Lemma 2.4 *Let $\mathcal{P}$ be an arbitrary set of vertex-disjoint 5-bounded s,t-paths in G. Then we can find (in polynomial time) a set $\mathcal{P}'$ of vertex-disjoint 5-bounded s,t-paths in G with $|\mathcal{P}'| \geq |\mathcal{P}|$ that contains either all three P_i^j or all three $\bar{P}_i^j$ paths within each G_i.*

Proof. If $\mathcal{P}$ contains three paths within each G_i, then these paths must be either all P_i^j or all $\bar{P}_i^j$ paths.

Now, suppose there is an i, $0 \leq i \leq q$, such that $\mathcal{P}$ contains two or less paths within G_i. We construct a new set $\mathcal{P}'$ containing either all three P_i^j or all three $\bar{P}_i^j$ paths, modifying only those paths in $\mathcal{P}$ that are contained within or do intersect with G_i.

First of all, we make sure that all paths of $\mathcal{P}$ containing one of the edges $c_h x_i^j$ or $c_h \bar{x}_i^j$ for some h and j are Q_i^j or $\bar{Q}_i^j$ paths, respectively. The only other 5-bounded paths in $\mathcal{P}$ that might contain these edges are the paths $[s, c_h, x_i^j, w_i^j, t]$ and $[s, c_h, \bar{x}_i^j, w_i^{j'}, t]$, respectively. But these paths can be replaced by Q_i^j or $\bar{Q}_i^j$ without decreasing the number of paths or destroying the vertex-disjointness or 5-boundedness of $\mathcal{P}$. After these exchanges, any 5-bounded path in $\mathcal{P}$ is either contained within G_i or intersects G_i in exactly one vertex x_i^j or $\bar{x}_i^j$.

In the second step, we simply drop all paths from $\mathcal{P}$ that are completely contained within G_i. Then we compute a maximum set of parallel vertex-disjoint 5-bounded s,t-paths within G_i, i.e., a set containing either only P_i^j paths or only $\bar{P}_i^j$ paths, that do not intersect with the remaining paths in $\mathcal{P}$. This can be done in constant time, because G_i is of fixed size. Since each literal in ϕ is either negated or un-negated, but never both, for each $j \in \{0,1,2\}$ only one of the vertices x_i^j and $\bar{x}_i^j$ can be contained in some clause graph. Hence, for each such j at least one of x_i^j and $\bar{x}_i^j$ is not contained in the remaining paths in $\mathcal{P}$ and, consequently, we will always find at least two parallel paths in G_i that are disjoint to $\mathcal{P}$.

If we find three such paths, we are done. Otherwise, assuming w.l.o.g. we found the two paths P_i^1 and P_i^2, $\mathcal{P}$ must contain Q_i^3. Then, in the third step, we can replace Q_i^3 by P_i^3 without decreasing the number of paths or destroying the vertex-disjointness or 5-boundedness in $\mathcal{P}$ and end up with all three P_i^j paths. □

Lemma 2.5

(i) Let x *be a truth assignment for the variables in ϕ that satisfies r' clauses. Then there is a solution $\mathcal{P}$ of ψ that contains exactly $3q + r'$ paths. In particular, $\mathcal{P}$ contains either all three P_i^j or all three $\bar{P}_i^j$-paths in each G_i.*

(ii) Let $\mathcal{P}$ be a solution of ψ and $r' := |\mathcal{P}| - 3q$. Then there is a truth assignment x *for the variables in ϕ that satisfies at least r' clauses.*

Proof.
(i): Given the truth assignment x for ϕ, we construct a set $\mathcal{P}$ of vertex-disjoint 5-bounded s,t-paths in G as follows: For each i, $0 \le i \le q$, we choose all three P_i^j paths, $0 \le j \le 2$, if variable x_i is *true* in the assignment, otherwise we choose all three $\bar{P}_i^j$ paths. For each clause satisfied by this assignment, we choose one of its literals that evaluates to *true*. If the chosen literal is x_i^j, i.e., the corresponding variable occurs un-negated, then we add the path$\bar{Q}_i^j$ to $\mathcal{P}$, if it is $\bar{\mathrm{x}}_i^j$ we add Q_i^j. Since $\bar{Q}_i^j$ is chosen only if P_i^j was (or Q_i^j and $\bar{P}_i^j$, respectively), all paths in $\mathcal{P}$ are vertex-disjoint. Furthermore, if x satisfies r' clauses of ϕ, then $\mathcal{P}$ contains $3q + r'$ paths.

(ii): By Lemma 2.4 we may assume that $\mathcal{P}$ contains either all P_i^j or all $\bar{P}_i^j$ paths in each G_i and, hence, $|\mathcal{P}| \ge 3q$. Given such a set of paths, we define the truth assignment x as follows: For each i, $0 \le i \le q$, we set the boolean variable x_i of ϕ to *true* if $\mathcal{P}$ contains all P_i^j paths within G_i, to *false* otherwise.

Since each of the remaining r' paths in $\mathcal{P}$ is either an $\bar{Q}_i^j$ or an Q_i^j path that does not intersect with the corresponding P_i^j or $\bar{P}_i^j$ paths, the literal that corresponds to the j-th occurrence of variable x_i in ϕ evaluates to *true*. By the construction of G, there can be at most one "vertex-disjoint" path through each vertex c_h, i.e., all these Q_i^j and $\bar{Q}_i^j$ paths correspond to literals in different clauses of ϕ. Hence, at least r' clauses of ϕ are satisfied by this assignment. □

Using this lemma, is easy to prove that the reduction constructed above is an L–reduction. We have to show that there are two constants $\alpha, \beta > 0$ (independent of ϕ), such that $Opt(\psi) \le \alpha \cdot Opt(\phi)$ and $|Opt(\phi) - \#$ clauses statisfied by x$| \le \beta \cdot |Opt(\psi) - |\mathcal{P}||$, for each solution $\mathcal{P}$ of ψ and the corresponding solution x of ϕ.

Lemma 2.5 *(ii)* implies that $Opt(\psi) \leq Opt(\phi) + 3q$. With $q \leq 3r$ and $r \leq 2 \cdot Opt(\phi)$ (see Lemma 2.2), it follows that $Opt(\psi) \leq (1 + 3 \cdot 3 \cdot 2)Opt(\phi) = 19 \cdot Opt(\phi)$, i.e., the optimal solution value is blown up by a factor of at most 19 in this reduction. That the absolute errors of the solutions $\mathcal{P}$ of ψ and x of ϕ are equal follows directly from Lemma 2.5. Hence, the second inequality holds for $\beta = 1$. This completes the proof of Theorem 2.3. □

Corollary 2.6 MDBP(ℓ) *is $\mathcal{APX}$–complete for each $\ell \geq 5$.*

Proof. Replace every edge sv in the above construction by a path $[s, v_1, \ldots, v_{\ell-5}, v]$.

2.2 Weighted-Disjoint-5-Bounded-Paths

In this section we will show that the problem of finding a set of k vertex-disjoint ℓ-bounded s,t-paths of minimum total weight in a weighted graph G is $\mathcal{NPO}$–complete for length bounds $\ell \geq 5$. Furthermore, we will see that, unless $\mathcal{P} = \mathcal{NP}$, this problem cannot be approximated within a factor of $2^{\langle\phi\rangle^\epsilon}$ for *any* $\epsilon > 0$, i.e., it is not in exp–$\mathcal{APX}$. The reason for this strong inapproximability result is that it is already $\mathcal{NP}$–complete to find k such paths, regardless of the weights on the edges. But even if finding k vertex-disjoint ℓ-bounded s,t-paths is easy, for example, if G is a complete graph, the problem cannot be approximated within reasonable bounds. We will show that in these cases it is $\mathcal{NP}$–hard to approximate WDBP(ℓ) for $\ell \geq 5$ within a factor of $2^{\langle\phi\rangle^\epsilon}$ for any $0 < \epsilon < 1$.

It is easy to see that WDBP(ℓ) is in $\mathcal{NPO}$ for all $\ell \in \mathbb{N}$. As in the previous section, it is sufficient to prove that WDBP(5) is $\mathcal{NPO}$-complete to show that WDBP(ℓ) is $\mathcal{NPO}$–complete for all $\ell \geq 5$.

Theorem 2.7 WDBP(5) *is $\mathcal{NPO}$-complete.*

We construct a Strict reduction (see [OM87]) from Minimum-Weighted-3-Satisfiability to WDBP(5) to prove this theorem.

The Minimum-Weighted-3-Satisfiability problem is defined as follows: Given a set X of boolean variables, a collection C of disjunctive clauses of at most three literals per clause, and a non-negative integer weight for each variable in X, the aim is to find a truth assignment for X that satisfies all clauses in C and minimizes the sum of the weights of the *true* variables. Due to Orponen and Mannila [OM87], this problem is known to be $\mathcal{NPO}$–complete. As in the unweighted satisfiability problem, we can restrict to the case where each variable occurs at most three times, the problem will remain $\mathcal{NPO}$–complete (see [ACG+98]). Although this restriction is not necessary to prove Theorem 2.7, it allows us to reuse exactly the same graph gadgets we had in the proof of Theorem 2.3.

Proof [2.7]. Let ϕ be an instance of Minimum-Weighted-3-Satisfiability(3) with variables $x_1, \ldots, x_q$, clauses $C_1, \ldots, C_r$, and nonnegative weights w_i, $1 \leq i \leq q$, associated with the variables.

Except that we have to define edge weights here, the reduction is the same as the one used to prove Theorem 2.3. We construct a graph G as described there and define the weights of its edges as follows:

$$w_e := \begin{cases} w_i & \text{if } e = v_i^1 x_i^1 \text{ for some } i \in \{1, \ldots, q\}, \\ 0 & \text{otherwise.} \end{cases}$$

Formally, the instance $\psi = \psi(\phi)$ of $\mathsf{WDBP}(5)$ is defined as the problem of finding $3q + r$ vertex-disjoint 5-bounded paths from s to t in G with minimum total weight. Obviously, this reduction can be accomplished in polynomial time. Note that, since in MINIMUM-WEIGHTED-3-SATISFIABILITY(3) a feasible assignment must satisfy *every* clause, we have to find $3q + r$ vertex-disjoint 5-bounded s, t-paths in G, i.e., three paths within each G_i and one path within each H_h. A set $\mathcal{P}$ of less paths would not correspond to a feasible assignment in ϕ.

It is easy to see that

$$\delta_G(s) = \sum_{i=1}^{q} \delta_{G_i}(s) + \sum_{h=0}^{r} \delta_{H_h}(s) = 3q + r,$$

where $\delta_G(s)$ denotes the degree of vertex s with respect to (sub-) graph G. Hence, for each set $\mathcal{P}$ of $3q + r$ vertex-disjoint 5-bounded s, t-paths in G, every subgraph G_i must contain exactly three of these paths and every H_h exactly one of them. With Lemmas 2.4 and 2.5 it follows that $\mathcal{P}$ corresponds to a feasible solution x of ϕ. On the other hand, these Lemmas imply that for each feasible truth assignment x of ϕ there is a set $\mathcal{P}$ of $3q + r$ vertex-disjoint 5-bounded s, t-paths in G.

By our construction, $\mathcal{P}$ contains path P_i^1 if and only if variable x_i is *true* in the corresponding assignment x. Because, for each $i \in \{1, \ldots, q\}$, this path is the only 5-bounded s, t-path that contains edge $v_i^1 x_i^1$ in G, the total weight of $\mathcal{P}$ equals the total weight of this assignment. □

Corollary 2.8 $\mathsf{WDBP}(\ell)$ *is $\mathcal{NPO}$-complete for all $\ell \geq 5$.*

It is known that for each $\mathcal{NPO}$–complete problem there is some threshold $\epsilon > 0$, such that the problem cannot be approximated within $2^{\langle\phi\rangle^\epsilon}$ [CK95], where $\langle\phi\rangle$ denotes the coding size of the given problem instance ϕ. Since, in our case, already the problem of finding a feasible solution is $\mathcal{NP}$–hard [IPS82], the minimization problem $\mathsf{WDBP}(5)$ cannot be approximated within *any* such bound, unless $\mathcal{P} = \mathcal{NP}$.

Theorem 2.9 $\mathsf{WDBP}(\ell)$ *is not in* $\texttt{exp}$–$\mathcal{APX}$ *for any $\ell \geq 5$, unless $\mathcal{P} = \mathcal{NP}$.*

Note that, due to the polynomial bound on the computation time of the objective function for optimization problems in $\mathcal{NPO}$, a classification of approximability beyond $\texttt{exp}$–$\mathcal{APX}$ does not make sense in $\mathcal{NPO}$. If we are able to find at least a feasible solution whenever one exists, this solution will always be within a factor of $2^{\langle\phi\rangle^\epsilon}$ of the optimum value for some fixed $\epsilon > 0$.

Unfortunately, even if we are given an initial feasible solution for $\mathsf{WDBP}(\ell)$ or if we can compute one in polynomial time, for example, if G is a complete graph, the problem remains extremely hard to approximate.

Theorem 2.10 *It is $\mathcal{NP}$–complete to approximate* $\mathsf{WDBP}(\ell)$ *for complete graphs and $\ell \geq 5$ within a factor of*

(i) $2^{\langle G\rangle^\epsilon}$ *for any $\epsilon > 0$, or within*

(ii) $2^{\langle\phi\rangle^\epsilon}$ *for any $0 < \epsilon < 1$.*

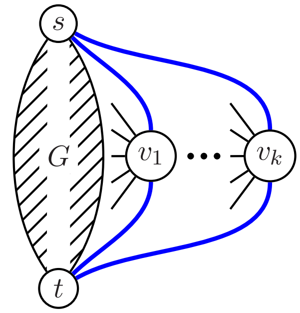

Figure 4: The weighted graph G'

Proof. We show that any algorithm that approximates $\mathsf{WDBP}(\ell)$ within such a factor could be used to find a set of k vertex-disjoint ℓ-bounded s,t-paths in a graph G, which is $\mathcal{NP}$–complete. The basic idea is to guarantee the existence of k such paths by adding k extra s,t-paths to G and assign astronomical weights to their edges.

We start with the proof of *(i)*. Let $n := |V|$, $m := |E|$, and $N := \langle G \rangle$. Suppose we have a polynomial time algorithm A that approximates $\mathsf{WDBP}(\ell)$ within a factor of 2^{N^ϵ} for some fixed $\epsilon > 0$. Given instance $\phi = (G, s, t, k, \ell)$ of the decision problem of finding k vertex-disjoint ℓ-bounded s,t-paths in G, an instance $\psi = (G', s, t, k, \ell, w)$ of $\mathsf{WDBP}(\ell)$ is defined as follows: We construct a new complete graph $G' = (V', E')$ with vertex set $V' := V \cup \{v_1, \ldots, v_k\}$, see Figure 4. We set the weights for the edges of G as follows:

$$w_e := \begin{cases} 1 & \text{if } e = uv \text{ with } u, v \in V, \\ 2^{N^\epsilon} \cdot m & \text{otherwise.} \end{cases}$$

Obviously, G' contains k vertex-disjoint ℓ-bounded s,t-paths.

Note that this transformation is polynomial. Although the values of the weights of the new edges are not polynomially bounded in N, their encoding size is. For the largest edge weight $w_{max} = 2^{N^\epsilon} m$ we have $\langle w_{max} \rangle \leq N^\epsilon + \langle m \rangle$.

It is easy to see that algorithm A, which is supposed to approximate ψ within a factor of 2^{N^ϵ} of the optimum value, could be used to solve ϕ: If A returns a solution of weight m or less, this solution contains only edges inherited from G, i.e., it is a feasible solution of ϕ. If the solution returned by A has weight larger than m, it must contain at least one of the new edges. But this implies that the solution's weight is even larger than $2^{N^\epsilon} m$ and, by the approximation guarantee of A, that no solution of weight less or equal to m exists. Hence, there are no k vertex-disjoint ℓ-bounded s,t-paths in G, since, otherwise, these paths would define a solution of ψ containing no new edge, which implies that $Opt(\psi) \leq m$.

Analogously, we can prove *(ii)*, assigning weight one to the edges of G' that are inherited from G and weight $2^{p^\epsilon} m$ to the new edges, where p is sufficiently large. With the same arguments as above, this will show the inapproximability within a factor of 2^{p^ϵ}. To choose an appropriate value for p, we need a bound on the encoding size of the new instance ψ of $\mathsf{WDBP}(\ell)$.

$$\begin{aligned} \langle \psi \rangle &= \langle G' \rangle + \langle \mathbf{w} \rangle \\ &\leq 2n + 4n^2 + 4n^2 \cdot \langle 2^{p^\epsilon} \cdot m \rangle \\ &\leq 2n + 4n^2 + 4n^2 \cdot p^\epsilon \log m \end{aligned}$$

Clearly, for each $\epsilon < 1$, there exists a polynomial $p(n,m)$ in n and m, such that $2n + 4n^2 + 4n^2 p(n,m)^\epsilon \log m \leq p(n,m)$. If we set the weights of the new edges in G to $2^{p(n,m)^\epsilon} m$, then $\langle\psi\rangle \leq p(n,m)$, i.e., the encoding size of ψ is polynomially bounded in ϕ, and ψ is not approximable within $2^{p(n,m)^\epsilon}$. □

Remark 2.11 *For complete graphs* WDBP(ℓ) *can be approximated within a factor of* $2^{\langle\phi\rangle}$.

Since in the Weighted-Disjoint-ℓ-Bounded-Paths the value of each solution is at most the sum of all edge weights, any feasible solution yields a $2^{\langle\phi\rangle}$ approximation.

Differently from other natural optimization problems on weighted graphs, as for example the traveling salesman problem or the steiner tree problem, it does not make sense to consider *metric* instances of the Weighted-Disjoint-ℓ-Bounded-Paths problem, i.e., instances where G is complete and the weights satisfy the triangle inequality. In such graphs an optimal solution is trivial: It will always contain the direct st-edge and the cheapest $k-1$ paths of length 2.

On the other hand, the weighted problem is still hard to approximate if we restrict to "semi–metric" instances, i.e., instances where G need not be complete, but for each cycle C in G and each edge $e \in C$ the inequality $w_e \leq \sum_{f \in C-e} w_f$ holds. Furthermore, even if we know an initial feasible solution, the same inapproximability results as for the complete Weighted-Disjoint-ℓ-Bounded-Paths problem hold.

Corollary 2.12 *Theorems 2.7 to 2.10 hold for semi–metric instances of* WDBP(ℓ)*, too.*

That Theorems 2.7 to 2.9 hold for semi-metric instances is trivial. The proof for the analogous version of Theorem 2.10, where G need not be complete but semi-metric and an initial feasible solution is given or can be found in polynomial time, can be easily obtained from the original one by a simple modification and is left to the reader.

3 Polynomially Solvable Cases

In the previous section we have seen that both the Max-Disjoint-ℓ-Bounded-Paths problem and the Weighted-Disjoint-ℓ-Bounded-Paths problem are computationally hard to approximate if the length bound ℓ is five or more. Fortunately, in practical applications the length bounds are usually very small. In the design of optical backbone telecommunication networks, for example, one often restricts to paths containing not more than three or four fiber-lines.

Trivially, both problems are polynomially solvable for $\ell \leq 2$. It was shown by Itai et. al. in [IPS82] that the Max-Disjoint-ℓ-Bounded-Paths problem is also polynomially solvable for lengths $\ell = 3$ and $\ell = 4$. As mentioned in the previous section, the metric Weighted-Disjoint-ℓ-Bounded-Paths problem is trivially solvable for any length bound $\ell \in \mathbb{N}$.

Proposition 3.1 *The metric* Weighted-Disjoint-ℓ-Bounded-Paths *problem is polynomially solvable for all length bounds* $\ell \in \mathbb{N}$.

In the following we will show that the semi-metric version of this problem is polynomially solvable for any fixed length bound $\ell \leq 4$. Recall that this problem is $\mathcal{NPO}$–complete for $\ell \geq 5$.

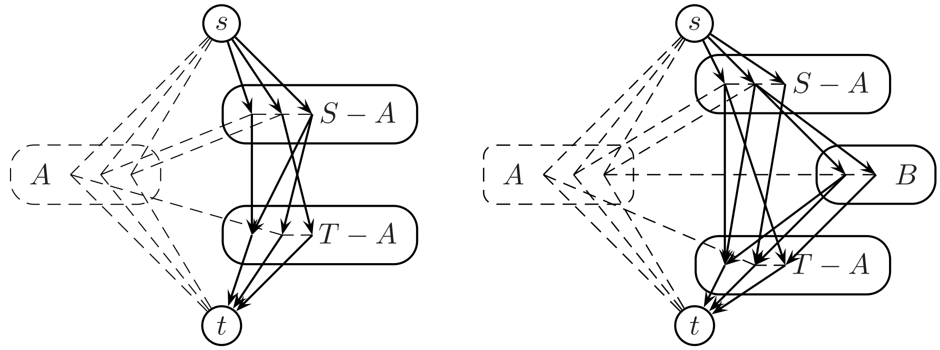

Figure 5: Graph G' for $\ell = 3$ and $\ell = 4$

Theorem 3.2 *The semi-metric* WEIGHTED-DISJOINT-ℓ-BOUNDED-PATHS *is polynomially solvable for all* $\ell \leq 4$.

Proof. The cases $\ell = 1$ and $\ell = 2$ are trivial. In the following we present an algorithm that solves WDBP(ℓ) optimally for $\ell = 3$ and $\ell = 4$.

Let ϕ be a semi-metric instance of WDBP(ℓ) with $\ell = 3$ or $\ell = 4$. By S we denote the set of all neighbors of s, by T the set of all neighbors of t. Let $A := S \cap T$, and, if $\ell = 4$, let B be the set of all vertices in $G - (S \cup T \cup \{s, t\})$ that have a neighbor in both S and T. For $\ell = 3$ we can restrict to the case where $G - (S \cup T \cup \{s, t\})$ is empty, since no s, t-path containing a vertex in $G - (S \cup T \cup \{s, t\})$ can be shorter than four. Analogously, we can assume $G - (S \cup T \cup B \cup \{s, t\})$ is empty and that B is a stable set if $\ell = 4$.

It is easy to see that there is an optimal solution $\mathcal{P}^*$ of ϕ, such that any vertex $v \in A$ is either contained in a path of length two in $\mathcal{P}^*$ or does not occur in $\mathcal{P}^*$ at all. Otherwise, the path containing v could be shortened to the path $[s, v, t]$ without increasing the total weight, using the fact that G is semi-metric. Hence, we can assume that A is a stable set and, with the same argument, that S and T are stable sets, too.

In the first step of our algorithm we compute the set $\mathcal{P}_{12}$ consisting of all s, t-paths of length one and two in G, i.e., the path $[s, t]$ if edge st exists and all paths $[s, v, t]$ with $v \in A$. We denote $k_{12} := |\mathcal{P}_{12}|$ and, for each j, $1 \leq j \leq \min\{k_{12}, k\}$, we denote the set of the j weight minimal paths of $\mathcal{P}_{12}$ by $\mathcal{P}_{12}^j$. Let G' be the digraph obtained from G after removing the paths in $\mathcal{P}_{12}$ and all vertices in A and orienting the remaining edges from s to t as shown in Figure 5. Obviously, each s, t-path in G' corresponds to an s, t-path in G that has length three or four an does not contain any vertex from A, and vice versa.

In the second step, we compute the maximum number k_{34} of vertex-disjoint ℓ-bounded s, t-paths in G'. For each i, $1 \leq i \leq \min\{k_{34}, k\}$, we determine a set $\mathcal{P}_{34}^i$ of exactly i vertex-disjoint ℓ-bounded s, t-paths of minimum total weight in G', applying minimum cost flow techniques [Suu74]. Clearly, within each $\mathcal{P}_{34}^i$ all paths are of length three or four. Notice that if $k_{12} + k_{34} < k$, then there is no feasible solution for ϕ and we can stop.

Finally, for each i, $max\{0, k - k_{34}\} \leq i \leq \min\{k_{34}, k\}$, we construct a solution $\mathcal{P}^i$ of ϕ consisting of the paths in $\mathcal{P}_{12}^i$ and in $\mathcal{P}_{34}^{k-i}$ (their corresponding paths in G). If there exists a feasible solution for ϕ, then, clearly, for all these i the sets $\mathcal{P}^i$ contain k vertex-disjoint ℓ-bounded s, t-paths. Furthermore, $\mathcal{P}^i$ is an optimal solution for at least one of these i. This is easy to see:

Let $\mathcal{P}^*$ be an optimal solution of ϕ. We may assume that any vertex $v \in A$ is either contained in a path of length two in $\mathcal{P}^*$ or does not occur in $\mathcal{P}^*$ at all. Let $\mathcal{P}^*_{12}$ be the set of all paths of length one or two in $\mathcal{P}^*$, $\mathcal{P}^*_{34}$ be the set consisting of all paths of length three or four in $\mathcal{P}^*$, and $j = |\mathcal{P}^*_{12}|$. Then $w(\mathcal{P}^*_{12}) \geq w(\mathcal{P}^j_{12})$ and $w(\mathcal{P}^*_{34}) \geq w(\mathcal{P}^{k-j}_{34})$, because $\mathcal{P}^j_{12}$ and $\mathcal{P}^{k-j}_{34}$ contain j vertex-disjoint s,t-paths of length one and two in G and $k-j$ vertex-disjoint s,t-paths of length three and four in G' with minimum total weight, respectively. Hence, $w(\mathcal{P}^*) \geq w(\mathcal{P}^j)$ and $\mathcal{P}^j$ is an optimal solution, too. □

Whether WDBP(3) and WDBP(4) are polynomially solvable for arbitrary edge weights is still open.

4 Related Problems

In this last section we extend our results to several related problems. In the area of network design, for example, one considers not only vertex-disjoint paths between communicating nodes to model survivability, but edge-disjoint paths as well. Furthermore, in many cases it is easier to model the network by a directed graph than by an undirected one, especially when complex traffic routing is considered.

For the following unweighted maximization problems the underlying decision problems have been studied by Itai et. al. in [IPS82]. Using the same techniques as in Section 2.1, we modified the graph gadgets introduced there in order to prove the $\mathcal{NP}$–completeness of the decision problems to obtain appropriate gadgets for approximation preserving reductions for the optimization problems. Most of the results stated below can be proven analogously to the ones in Sections 2 and 3. Except for Theorem 4.1, we will only sketch the basic ideas and modifications necessary.

4.1 Exact Path Lengths

The first group of problems we consider here is obtained by requiring that the paths are of length equal to ℓ instead of being ℓ-bounded. We denote the resulting problems Max-Disjoint-Exact-ℓ-Length-Paths and Weighted-Disjoint-Exact-ℓ-Length-Paths by MDEP(ℓ) and WDEP(ℓ), respectively.

In [IPS82] it was shown that the problem of deciding whether a given number k of vertex-disjoint s,t-paths of length equal to ℓ exists in G is $\mathcal{NP}$–complete for $\ell \geq 4$ and polynomially solvable for $\ell \leq 3$. The $\mathcal{APX}$- and $\mathcal{NPO}$–completeness of MDEP(ℓ) and WDEP(ℓ), respectively, can be proven for lengths $\ell \geq 4$. The proofs are identical to those of the length bounded problems, except we use the gadgets graphs shown in Figure 6. The corresponding version of Theorem 2.9 for WDEP(ℓ) with $\ell \geq 4$ follows directly, the analogous version of Theorem 2.10 is proven by adding extra paths of length exactly ℓ instead of length two, as in the proof of Theorem 2.10, to the given graph.

For lengths $\ell \leq 3$ both MDEP(ℓ) and WDEP(ℓ) are solvable in polynomial time. The cases $\ell = 1$ and $\ell = 2$ are trivial, MDEP(3) and WDEP(3) can be reduced to a maximum matching or a minimum weight cardinality k matching problem (see [Edm65, Gab90]), respectively: Each s,t-path of length three in G corresponds to exactly one edge vw, where v is a neighbor of s and w is a neighbor of t, and vice versa. Hence, computing a maximum matching in the graph given by these edges yields a maximum set of such paths. If we set the weight for each

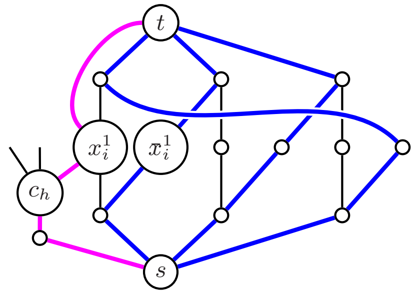

Figure 6: H_h and G_i for MDEP(4)

of these edges to the weight of its corresponding path, we can solve WDEP(3) for arbitrary weights using a minimum weight cardinality k matching algorithm.

4.2 Edge-Disjoint Paths

Replacing the vertex-disjointness constraint by edge-disjointness, we obtain another group of problems similar to MDBP(ℓ) and WDBP(ℓ). We denote the corresponding edge-disjoint paths problems Max-Edge-Disjoint-ℓ-Bounded-Paths and Weighted-Edge-Disjoint-ℓ-Bounded-Paths by MEDBP(ℓ) and WEDBP(ℓ), respectively. If, additionally, the paths are required to be of length equal to ℓ instead of being ℓ-bounded, we denote the respective problems by MEDEP(ℓ) and WEDEP(ℓ).

One can show $\mathcal{APX}$- and $\mathcal{NPO}$–completeness for MEDBP(ℓ) and WEDEP(ℓ), respectively, for all $\ell \geq 5$. The corresponding exact path length problems MEDBP(ℓ) and WEDEP(ℓ) are $\mathcal{APX}$- and $\mathcal{NPO}$–complete for all lengths $\ell \geq 4$. The proofs are similar to the vertex-disjoint case. We use the gadgets shown in Figure 7 for the problems with bounded length and those in Figure 8 for the problems with exact path length. Literally following the proofs of Theorems 2.9 to 2.12 yields the same inapproximability results as in the vertex-disjoint cases for WEDBP(ℓ) with $\ell \geq 5$ and for WEDEP(ℓ) with $\ell \geq 4$.

Showing that all four edge-disjoint paths problems are polynomially solvable for lengths $\ell \leq 3$ takes a little more effort. This answers an open question in [IPS82], namely whether the problem of deciding if a given number of edge-disjoint s,t-paths of length equal to three exists in a graph is polynomially solvable. We will give a detailed proof here.

Theorem 4.1 *The problems* MEDEP(ℓ)*,* MEDBP(ℓ)*,* WEDEP(ℓ)*, and* WEDBP(ℓ) *are polynomially solvable for each* $\ell \leq 3$*.*

Proof. Again, the cases $\ell = 1$ and $\ell = 2$ are trivial. Consider the case $\ell = 3$.

Let ϕ be the given problem instance with graph $G = (V, E)$. If ϕ is an instance of MEDEP(3) or WEDEP(3), we call an s,t-path in G valid if it contains exactly three edges, if ϕ is an instance of MEDBP(3) or WEDBP(3), we call it valid if it contains no more than three edges. By S we denote the set of all neighbors of s, by T the set of all neighbors of t, and $B := S \cap T$. We can assume that G contains only these vertices, because any other vertex cannot be contained in a valid s,t-path. Furthermore, we may assume that G contains no edge that is not contained in any valid s,t-paths.

We define a digraph $D = (V', A)$ as follows: The vertex set V' contains the copies s' and t' of the vertices s and t of G, one copy v_S for each vertex $v \in S$, and one copy v_T for each

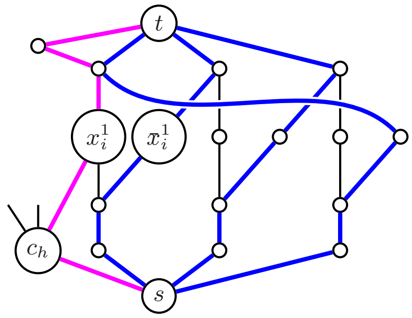

Figure 7: H_h and G_i for $\mathsf{MEDBP}(5)$

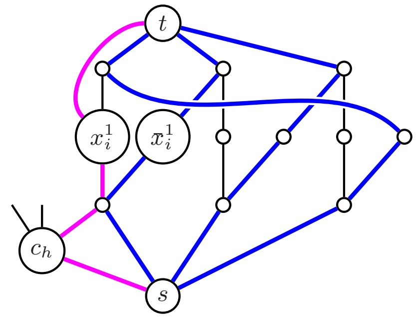

Figure 8: H_h and G_i for $\mathsf{MEDEP}(4)$

vertex $v \in T$. Note that for each vertex $v \in B$ there are now two copies v_S and v_T in V'. Furthermore, V' contains two extra vertices u_e and w_e for each edge $e \in [S,T]$:

$$
\begin{aligned}
V_S &:= \{v_S \mid v \in S\}\,, \\
V_T &:= \{v_T \mid v \in T\}\,, \\
V_U &:= \{u_e \mid e \in [S,T]\}\,, \\
V_W &:= \{w_e \mid e \in [S,T]\}\,, \\
V' &:= \{s', t'\} \cup V_S \cup V_T \cup V_U \cup V_W,
\end{aligned}
$$

where $[V_1, V_2] := \{e \mid e = uv \in E, u \in V_1, w \in V_2\}$. If all paths must be of length equal to three, i.e., ϕ is an instance of $\mathsf{MEDEP(3)}$ or $\mathsf{WEDEP(3)}$, the arc set A is defined as follows: For each edge $e = sv \in [s,S]$ or $e = vt \in [T,t]$ the set A contains an arc $a_e = (s', v_S)$ or $a_e = (v_T, t')$ between the the corresponding vertices, respectively. For all other edges $e \in E$, i.e., $e = vv' \in [S,T]$, A contains an arc $a_e = (u_e, w_e)$ between the two vertices associated with this edge and the arcs (v_S, u_e) and (w_e, v'_T) that represent the incidences of this edge e. If both v and v' are in B, then we have four incidence arcs (v_S, u_e), (v'_S, u_e), (w_e, v_T), and (w_e, v'_T) in D, see Figure 9.

$$
\begin{aligned}
A_{[s,S]} &:= \left\{a_e = (s', v_S) \mid e \in [s,S]\right\}, \\
A_{[S,T]} &:= \{a_e = (u_e, w_e) \mid e \in [S,T]\}\,, \\
A_{[T,t]} &:= \left\{a_e = (v', t') \mid e \in [s,S]\right\}, \\
A_S &:= \{(v_S, u_e) \mid v \in S, e \in \delta_G(v)\}\,, \\
A_T &:= \{(w_e, v_T) \mid v \in T, e \in \delta_G(v)\}\,, \\
A &:= A_{[s,S]} \cup A_{[S,T]} \cup A_{[T,t]} \cup A_S \cup A_T.
\end{aligned}
$$

Otherwise, if also paths of length one and two are allowed, then A additionally contains an arc (s', t') if edge st exists and an arc (v_S, v_T) between the two copies v_S and v_T of each vertex $v \in B$.

$$
\begin{aligned}
A_{st} &:= \{a_e = (s', t') \mid st \in E\}, \\
A_B &:= \{(v_S, v_T) \mid v \in B\}, \\
A &:= A_{[s,S]} \cup A_{[S,T]} \cup A_{[T,t]} \cup A_S \cup A_T \cup A_{st} \cup A_B.
\end{aligned}
$$

For an example see Figure 9.

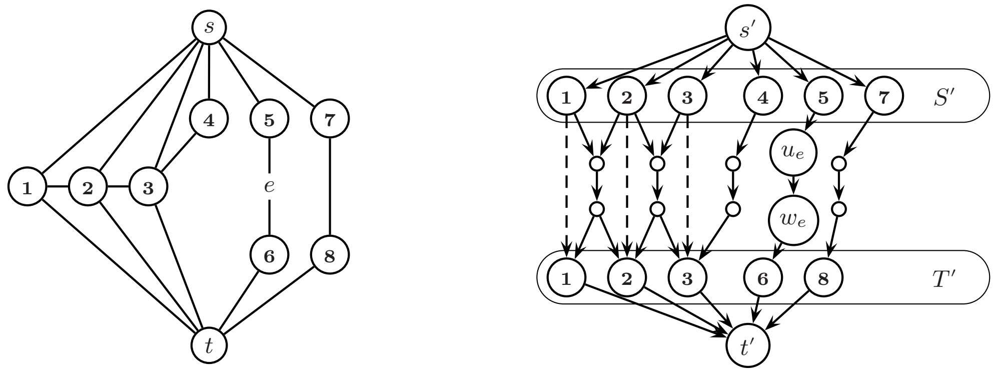

Figure 9: Example Graph G and corresponding Digraph D

Note that each edge in G has exactly one corresponding arc in D. Namely, the edges in $[s, S]$ correspond to the arcs in $A_{[s,S]}$, those in $[S, T]$ to the arcs in $A_{[S,T]}$ and so on. Each valid s, t-path in G corresponds to a directed s', t'-path in D, and vice versa. The arcs in A_B are contained in those s', t'-paths in D that correspond to paths of length two in G, the arc (s', t') to the direct st edge in G. Clearly, if a set of valid paths in G is edge-disjoint, so are the corresponding paths in D. On the other hand, if two paths share an edge in G, also their corresponding paths in D must share an arc, because there is only one correspondence of the common edge in D. If, in the weighted problems, we assign to the arcs in $A_{[s,S]}$, $A_{[S,T]}$, $A_{[T,t]}$, A_{st}, and A_B the weights of their corresponding edges in G and weight zero to all other arcs, the weight of a path in G equals the weight of its corresponding path in D.

This implies that any set of k valid edge-disjoint s, t-paths in G corresponds to a set of k arc-disjoint directed s', t'-paths in D of the same weight, and vice versa. Hence, applying standard maximum flow or minimum cost flow techniques in D, all four edge-disjoint paths problems can be solved in polynomial time. □

It remains open whether MEDBP(4) and WEDBP(4) are polynomially solvable.

4.3 Directed and Non-Simple Graphs

For all of the above problems we can define the corresponding directed version. It is not hard to see that all results we obtained for undirected graphs also hold in the directed case. In the proofs of the appropriate versions of Theorem 2.3 and Theorem 2.7 each edge in the gadget graphs can be contained in only one "direction" in an s, t-path of bounded or exact length. Replacing the edges with arcs that have exactly this direction yields the proofs of the $\mathcal{APX}$– and $\mathcal{NPO}$–completeness for directed graphs. Note that we used gadget graphs different from those proposed by Itai et. al. in [IPS82] to show the $\mathcal{APX}$–completeness of MEDEP(ℓ) for $\ell \geq 4$. In our gadget graphs no edge can be contained in both directions in an s, t-paths of length four. Hence, these gadgets can be used to show that also the directed version MEDEP(ℓ) is $\mathcal{APX}$–complete for $\ell \geq 4$ and that its underlying decision problem is $\mathcal{NP}$–complete, a question that was left open in [IPS82].

On the other hand, the algorithms proposed to solve the polynomial cases of the undirected problems implicitly solved the corresponding directed counterparts utilizing standard matching or flow techniques. Hence, the the directed versions of these problems are polynomially solvable.

Allowing multiple edges (or arcs) and loops in the graph (or digraph) obviously does not affect the complexity of the vertex-disjoint paths problems. When looking for vertex-disjoint paths, we can remove all loops and all but the cheapest one from each set of parallel edges, except for the parallel st edges, that define independent s,t-paths.

The results proven for edge-disjoint paths problems do not change either, if multiple edges and loops are allowed. Trivially, all completeness and inapproximability results still hold. To see that the problems remain polynomially solvable for non-simple graphs and lengths $\ell \leq 3$, one can literally follow the proof of Theorem 4.1. Remember, that for each edge in G there is exactly one corresponding arc in D. Note that, however, the arcs that correspond to parallel edges in $[S,T]$ are not parallel in D but independent (u_e, w_e) arcs.

5 Conclusions

We have shown that the MAX-DISJOINT-ℓ-BOUNDED-PATHS problem is $\mathcal{APX}$–complete for $\ell \geq 5$. That this problem is polynomially solvable for $\ell \leq 4$ is due to [IPS82].

The WEIGHTED-DISJOINT-ℓ-BOUNDED-PATHS problem is proven to be $\mathcal{NPO}$–complete for length bounds $\ell \geq 5$, even if the underlying weighted graph is semi-metric. In case an initial feasible solution is given or can be found in polynomial time, f.e., if the underlying graph is complete, it is $\mathcal{NP}$–complete to approximate $\mathsf{WDBP}(\ell)$ for $\ell \geq 5$ within a factor of $2^{\langle G \rangle^\epsilon}$ for any constant $\epsilon > 0$, or within $2^{\langle \phi \rangle^\epsilon}$ for any $0 < \epsilon < 1$.

For length bounds $\ell \leq 4$ we can show that $\mathsf{WDBP}(\ell)$ can be solved in polynomial time for semi-metric weighted graphs. Whether the problem is polynomially solvable for $\ell = 3$ and $\ell = 4$, or at least approximable within reasonable bounds, remains open.

Finally, these results have been extended to the corresponding edge-disjoint paths problems and to the problems where the paths must have length exactly ℓ instead of at most ℓ. As a by-product, some open questions concerning the corresponding edge-disjoint paths problems have been answered. Namely, it was proven that the problem of finding the maximum number of edge-disjoint s,t-paths of length equal to ℓ is polynomially solvable for $\ell = 3$ in undirected graphs and $\mathcal{NP}$–hard for $\ell = 4$ in directed graphs. However, the complexity of the corresponding problem for 4-bounded edge-disjoint paths remains open.

References

[ACG+98] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti Spaccamela, and M. Protasi, *Approximate solution of NP-hard optimization problems*, Springer, to be published, 1998.

[AGW97] D. Alevras, M. Grötschel, and R. Wessäly, *Capacity and survivability models for telecommunications networks*, Tech. Report SC 97-24, Konrad-Zuse-Zentrum für Informationstechnik, Berlin, 1997.

[BC94] D. Bovet and P. Crescenzi, *Introduction to the theory of complexity*, Pretice Hall International, 1994.

[CK95] P. Crescenzi and V. Kann, *A compendium of NP optimization problems*, Tech. report, Dipartimento di Scienze dell' Informatione, Universitá di Roma "La Sapienza", 1995, latest version available via WWW: http://www.nada.kth.se/~viggo/problemlist/compendium.html.

[CKST96] P. Crescenzi, V. Kann, R. Silvestri, and L. Tresvisan, *Structure in approximation classes*, SIAM Jounal on Computing (submitted, 1996).

[Edm65] J. Edmonds, *Maximum matching and a polyhedron with 0,1-vertices*, Journal of Research of the National Bureau of Standards B (1965), no. 69B.

[Gab90] H. N. Gabow, *Data structures for weighted matching and nearest common ancestors with linking*, Proceedings of the 1^{st} Annual ACM–SIAM Symposium on Discrete Algorithms, 1990.

[IPS82] A. Itai, Y. Perl, and Y. Shiloach, *The complexity of finding maximum disjoint paths with length constraints*, Networks **12** (1982).

[LNLP78] L. Lovàsz, V. Neumann-Lara, and M. Plummer, *Mengerian theorems for paths of bounded length*, Periodica Mathematica Hungaria **9** (1978).

[OM87] P. Orponen and H. Mannila, *On approximation preserving reductions: Complete problems and robust measures*, Tech. Report C-1987-28, Department of Computer Science, University of Helsinki, 1987.

[Pap94] C. H. Papadimitriou, *Computational complexity*, Addison Wesley, 1994.

[PR84] Y. Perl and D. Ronen, *Heuristics for finding a maximum number of disjoint bounded paths*, Networks **14** (1984).

[PY88] C. H. Papadimitriou and M. Yannakakis, *Optimization, approximation and complexity classes*, Proc. 20^{th} ACM Symp. on the theory of computing (1988).

[PY91] C. H. Papadimitriou and M. Yannakakis, *Optimization, approximation and complexity classes*, J. Comput. System Sci. **43** (1991).

[Suu74] J. W. Suurballe, *Disjoint paths in a network*, Networks **4** (1974).