Approximation Strategies for Routing Edge Disjoint Paths in Complete Graphs

Adrian Kosowski*

Gdańsk University of Technology
Department of Algorithms and System Modeling
kosowski@sphere.eti.pg.gda.pl

Abstract. The paper deals with the well known Maximum Edge Disjoint Paths Problem (MaxEDP), restricted to complete graphs. We propose an off-line 3.75-approximation algorithm and an on-line 6.47-approximation algorithm, improving earlier 9-approximation algorithms due to Carmi, Erlebach and Okamoto (Proceedings WG'03, 143–155). Next, it is shown that no on-line algorithm for the considered problem is ever better than a 1.50-approximation. Finally, the proposed approximation techniques are adapted for other routing problems in complete graphs, leading to an off-line 3-approximation (on-line 4-approximation) for routing with minimum edge load, and an off-line 4.5-approximation (on-line 6-approximation) for routing with a minimum number of WDM wavelengths.

1 Introduction

The fundamental networking problem of establishing point-to-point connections between pairs of nodes in order to handle communication requests has given rise to numerous path routing problems in graph theory. The topology of the network is modeled in the form of a graph whose vertices correspond to nodes, while edges represent direct physical connections between nodes. This paper deals with the well established problem of handling the maximum possible number of communication requests without using a single physical link more than once, known as the *Maximum Edge Disjoint Paths Problem* (MaxEDP). We focus on the construction of approximation algorithms for the *NP*-hard MaxEDP problem in complete graphs, which are used to model networks with direct connections between all pairs of nodes. Two basic algorithmic approaches are considered — *off-line algorithms*, which compute a routing for a known set of requests provided at input, and *on-line algorithms*, which have to handle requests individually, in the order in which they appear.

Problem definition. The physical architecture of the network is given in the form of an undirected graph $G = (V, E)$, where V denotes the set of nodes, while

* Research supported by the State Committee for Scientific Research (Poland) Grant No. 4 T11C 047 25.

P. Flocchini and L. Gąsieniec (Eds.): SIROCCO 2006, LNCS 4056, pp. 130–142, 2006.
© Springer-Verlag Berlin Heidelberg 2006

E represents the set of connections between them. A sequence of edges $P = (e_1, e_2, \ldots, e_l) \in E^l$, such that $e_i = \{v_i, v_{i+1}\}$ for some two vertices $v_i, v_{i+1} \in V$, is called a *path* of length $l = |P|$ in G, with endpoints v_1 and v_{l+1}. The symbol $P_{\{u,v\}}$ is used to denote any path in G with endpoints $u, v \in V$. A pair of paths P_1 and P_2 is called *conflicting* if there exists an edge $e \in E$ such that $e \in P_1$ and $e \in P_2$. For a given set of paths R in graph G, the *conflict graph* $Q(R)$ is a simple graph with vertex set R and edges connecting all pairs of vertices corresponding to paths from set R which conflict in G.

An *instance* I in network G is defined as any multiset of pairs $\{u, v\}$, $u, v \in V$, $u \neq v$, such that each element of I represents a single *communication request* between a pair of nodes. An equivalent representation of instance I may be given in the form of the *instance multigraph* $H(I) = (V, I)$, where communication requests are treated as edges of $H(I)$. A *routing* R of instance I in network G is a multiset of paths in G, such that there is a one-to-one correspondence between paths $P_{\{u,v\}} \in R$ and elements $\{u, v\} \in I$. The set of all routings of instance I is denoted as $\mathcal{R}(I)$. For use in further considerations, we define the following parameters for any routing R:

- *dilation* $\mathtt{d}(R)$, defined as the length of the longest path in routing R: $\mathtt{d}(R) = \max_{P \in R} |P|$,
- *edge load* $\pi(R)$, given by the formula: $\pi(R) = \max_{e \in E} |\{P \in R : e \in P\}|$.

A routing R is said to consist of *edge disjoint paths* if $\pi(R) = 1$, or equivalently, if conflict graph $Q(R)$ has no edges. A formal definition of the MaxEDP problem, expressed in these terms, is given below.

Maximum Edge Disjoint Paths Problem [MaxEDP]

Input: Instance I in graph G.
Solution: A set of pairwise edge-disjoint paths R_{OPT}, such that $R_{\text{OPT}} \in \mathcal{R}(I_{\text{OPT}})$ for some instance $I_{\text{OPT}} \subseteq I$.
Goal: Maximise the cardinality of R_{OPT}.

Notation. Throughout the paper, the complete graph with vertex set V is denoted K_V. Unless otherwise stated, we will assume that the MaxEDP problem is considered for the instance I in complete graph $G = K_V = (V, E)$. The optimal solution to the MaxEDP problem is some routing $R_{\text{OPT}} \in \mathcal{R}(I_{\text{OPT}})$ ($I_{\text{OPT}} \subseteq I$), while approximation algorithms yield a solution denoted as $R_S \in \mathcal{R}(I_S)$ ($I_S \subseteq I$), of not greater cardinality than R_{OPT}. Approximation ratios are understood in terms of upper bounds on the ratio $\frac{|I_{\text{OPT}}|}{|I_S|}$. The number of elements of a set or multiset, and also the length of a path, is written as $|P|$. The symbols Δ_H and χ'_H are used to denote the maximum vertex degree and the chromatic index of multigraph H, respectively.

State-of-the-art results. In the case of general networks G, the MaxEDP problem is closely related to a family of unsplittable flow problems. In consequence MaxEDP is *NP*-hard, difficult to approximate in polynomial time within

Table 1. New complexity results for the MaxEDP problem in complete graphs

Instance restriction	*Off-line complexity*		*On-line complexity*	
$\Delta_{H(I)} \leq \frac{\lvert V\rvert}{12}$	$O(\lvert V\rvert^3)$	Prop. 3	$O(\lvert V\rvert)$ per request	Cor. 3
$\lvert I\rvert < \lvert V\rvert$	$O(\lvert V\rvert)$	Prop. 2	$O(\lvert V\rvert)$ per request	Cor. 3
$\lvert I\rvert < k\lvert V\rvert$, const $k > 0$	$O(\lvert V\rvert^3)$	Thm. 2	not approx. within 1.50 for $\lvert I\rvert \geq 3\lvert V\rvert$	Thm. 8
$\lvert I\rvert < \lvert V\rvert^s$, const $s \in (1, 2)$	*NPH*, *PTAS*	Thm. 3, 4		
general case	3.75-approximation	Thm. 1	6.47-approximation	Thm. 7

a constant factor, and difficult to approximate within a factor of $O(\log^{\frac{1}{3}-\varepsilon}|E|)$, for any $\varepsilon > 0$ (unless $NP \subseteq ZPTIME(n^{\text{poly}\log n})$, [1]). The variant of MaxEDP defined for directed graphs is even difficult to approximate within $O(|E|^{\frac{1}{2}-\varepsilon})$, for any $\varepsilon > 0$ [12]. Both the directed and undirected version are approximable within a factor of $O(|E|^{\frac{1}{2}})$ [15].

When graph G is the complete graph K_V, the MaxEDP problem, though remaining *NP*-hard, becomes approximable within a constant factor. The best known approximation ratio was equal to 9 both in the off-line and on-line model of computation, due to Carmi, Erlebach and Okamoto [4]. A comparison of known approximation algorithms is provided in Table 2 at the end of the paper.

Our contribution and outline of the paper. In Section 2 we deal with the off-line MaxEDP problem in complete graphs, providing a 3.75-approximation algorithm based on the simple combinatorial concept of edge-coloring. Moreover, we show that for instances with significantly fewer than $|V|^2$ requests, the problem is either polynomially solvable, or admits a polynomial time approximation scheme. For the on-line version of the problem, in Section 3 we provide a 6.47-approximation algorithm, and show that no algorithm is better than 1.50-approximate, even for restricted instances. A summary of the most important new results concerning the MaxEDP problem is given in Table 1. Finally, in Section 4 we discuss the application of similar approximation techniques to other routing problems in complete graphs, and remark on their implementation in a distributed setting.

2 The Off-Line MaxEDP Problem in Complete Graphs

In the off-line routing model, it is assumed that all pairs of vertices forming the routed instance are initially known and all paths are determined by the routing algorithm at the same time.

2.1 Preliminaries: Bounds on Solution Cardinality

Factors in a multigraph. Let F_v be a set of nonnegative integers defined for each vertex $v \in V$. An *F-factor* in multigraph $H = (V, I)$ is a set of edges of H such that the number of edges from this set which are incident to vertex v

belongs to F_v. An $[a,b]$-*factor* is defined as an F-factor such that each set F_v consists of all integers from the range $[a,b]$. An $[a,b]$-factor with the maximum number of edges may be found efficiently by reduction to a minimum weighted perfect matching problem.

Proposition 1 ([16],[11]). *The problem of finding an $[a,b]$-factor with the maximum possible number of edges in multigraph $H = (V, I)$ can be solved in $O(|I|^3)$ time.*

Let I be an instance in graph K_V. Consider an instance I_{OPT} yielding an optimal solution to the MAXEDP problem for instance I. It is immediately evident that any vertex $v \in V$ can belong to at most $\deg_{K_V} v = |V| - 1$ requests of I_{OPT}, hence I_{OPT} is a $[0, |V| - 1]$-factor in $H(I)$ and we have the following bound.

Corollary 1. *The cardinality of the optimal solution to the* MAXEDP *problem for I is bounded from above by the size of the maximum $[0, |V|-1]$-factor in $H(I)$.*

Instances admitting an edge-disjoint routing. It is interesting to note that relatively wide classes of instances can be entirely routed using edge disjoint paths and in polynomial time. A short characterisation of two classes useful in further considerations is given below.

Proposition 2. *If $|I| < |V|$, then the entire instance I can be routed in K_V by edge disjoint paths, and a solution $R_{\text{OPT}} \in \mathcal{R}(I)$ to the* MAXEDP *problem, such that $\mathtt{d}(R_{\text{OPT}}) \leq 2$, can be determined in $O(|V|)$ time.*

Proof. The proof is constructive and proceeds by induction with respect to $|V|$. For $|V| = 2$, we have $|I| \leq 1$ and the proposition is obviously true. Next, let $|V| > 2$ be fixed and let $u \in V$ be a vertex belonging to the smallest number of requests in I, i.e. such that u is of minimal degree in $H(I)$. Since $|I| < |V|$, it is evident that $\deg_{H(I)} u = 0$ or $\deg_{H(I)} u = 1$. In the former case, we select an arbitrary request $\{v_1, v_2\} \in I$, and return the solution to the MAXEDP problem for I in K_V in the form of path $(\{v_1, u\}, \{u, v_2\})$ added to the solution to MAXEDP for instance $I \backslash \{\{v_1, v_2\}\}$ in complete graph $K_{V \backslash \{u\}}$. Thus $|R_{\text{OPT}}| = 1 + (|I| - 1) = |I|$ by the inductive assumption. In the latter case, let $\{u, v\} \in I$ be the only request involving vertex u. The sought routing then consists of the single-edge path $(\{u, v\})$ added to the solution to MAXEDP for instance $I \setminus \{\{u, v\}\}$ in $K_{V \backslash \{u\}}$. The described approach may easily be implemented in the form of an algorithm with $O(|V|)$ time complexity. □

Observe that the thesis of Proposition 2 does not hold if $|I| = |V|$ (it suffices to consider an instance composed of $|V|$ requests between a fixed pair of vertices). Nevertheless, if $|I| \in O(|V|)$ the problem can be solved in polynomial time (see Theorem 2).

Proposition 3. *If $\Delta_{H(I)} \leq \frac{|V|}{12}$, then the entire instance I can be routed in K_V by edge disjoint paths, and a solution $R_{\text{OPT}} \in \mathcal{R}(I)$ to the* MAXEDP *problem, such that $\mathtt{d}(R_{\text{OPT}}) \leq 2$, can be determined in $O(|V||I|)$ time.*

Proof. First, let us observe that the size of any instance I fulfilling the assumptions of the theorem is bounded by $|I| \leq \frac{|V|}{2} \cdot \frac{|V|}{12}$. The sought routing $R_{\text{OPT}} \in \mathcal{R}(I)$ consisting of edge disjoint paths can be formed by sequentially assigning paths to requests from I (in arbitrary order), in such a way as to preserve the following conditions:

1. The length of any path added to R_{OPT} is at most 2.
2. Each vertex of graph K_V is the center of at most $\frac{|V|}{12}$ paths.

It suffices to show that the described construction of routing R_{OPT} is always possible. Suppose that at some stage of the algorithm R_{OPT} fulfills conditions 1 and 2, and the next considered request is $\{v_1, v_2\}$. Vertex v_1 is the endpoint of at most $\frac{|V|}{12} - 1$ paths and the center of at most $\frac{|V|}{12}$ paths already belonging to R_{OPT}, thus at least $\frac{3|V|}{4}$ edges of K_V incident to v_1 do not belong to any path of R_{OPT}. The same is true for vertex v_2. Thus we immediately have that the set U of vertices connected to both v_1 and v_2 by edges unused in R_{OPT} is of cardinality $|U| \geq \frac{3|V|}{4} + \frac{3|V|}{4} - |V| = \frac{|V|}{2}$. Since routing R_{OPT} currently consists of fewer than $|I| \leq \frac{|V|}{2} \cdot \frac{|V|}{12}$ paths, by the pigeonhole principle there must exist a vertex $u \in U$ such that u is the center of fewer than $\frac{|V|}{12}$ paths from R_{OPT}. Therefore the request $\{v_1, v_2\}$ may be fulfilled by adding path $(\{v_1, u\}, \{u, v_2\})$ [1] to routing R_{OPT}, thus preserving the assumptions of the construction, which completes the proof. □

2.2 An Off-Line 3.75-Approximation Algorithm

Theorem 1. *There exists a* 3.75*-approximation algorithm for the* MaxEDP *problem in complete graphs with* $O(|I|^3)$ *runtime. The dilation of the returned solution is not greater than* 2.

Proof. Let I be an arbitrary instance in complete graph K_V, and let $I_{\text{OPT}} \subseteq I$ be a subset of the considered instance whose routing is an optimal solution to the MaxEDP problem. We denote by $H^* = (V, I^*)$ a multigraph $H^* \subseteq H(I)$ with the maximum possible number of edges, such that $\Delta_{H^*} < |V|$. Since the edge set of multigraph H^* is in fact a maximum $[0, |V|-1]$-factor in $H(I)$, by Proposition 1 multigraph H^* can be determined in $O(|I|^3)$ time. Moreover, by Corollary 1 we have $|I_{\text{OPT}}| \leq |I^*|$.

We will now show that there exists an algorithm with $O(|I|^3)$ runtime which finds a routing $R_S \in \mathcal{R}(I_S)$ composed of edge disjoint paths, such that $I_S \subseteq I^* \subseteq I$ and the obtained solution is a 3.75-approximation of the optimal MaxEDP solution, $|I_S| \geq \frac{|I^*|}{3.75} \geq \frac{|I_{\text{OPT}}|}{3.75}$. Instance I_S is constructed as a subset of the edge set of multigraph H^*. Since $\Delta_{H^*} < |V|$, by a well known result due to Shannon [10], the chromatic index χ'_{H^*} is bounded by $\chi'_{H^*} \leq \frac{3\Delta_{H^*}}{2} < \frac{3|V|}{2}$, and an edge coloring of multigraph H^* using not more than $\frac{3|V|}{2}$ colors can be obtained in $O(|I|^3)$ time. Without loss of generality we may assume that colors are labelled

[1] Throughout the paper, we assume that edges of the form $\{v, v\}$ which appear in notation when enumerating edges of paths should be treated as nonexistent.

with integers from the range $\{1, \ldots, \frac{3|V|}{2}\}$, in such a way that a color with a smaller label is never assigned to fewer edges than a color with a larger label. Let I_C denote the subset of edges from I^* colored with colors from the range $\{1, \ldots, |V|\}$. Due to the adopted ordering of the color labels, we immediately have $|I_C| \geq \frac{2}{3}|I^*|$. For each edge $\{v_1, v_2\} \in I_C$, let $c_{\{v_1,v_2\}}$ denote the color assigned to this edge, which is an integer from the range $\{1, \ldots, |V|\}$, and as such may be treated as an identifier of some vertex in graph K_V (see Fig. 1 for an exemplary illustration).

Let us now consider routing R_C of instance I_C in graph K_V, defined as follows: $R_C = \{(\{v_1, c_{\{v_1,v_2\}}\}, \{c_{\{v_1,v_2\}}, v_2\}) : \{v_1, v_2\} \in I_C\}$. No vertex of H^* may ever be incident to two edges from I_C of the same color, therefore each edge $\{v_1, v_2\}$ of graph K_V belongs to at most two paths of routing R_C — one path, in which v_1 is an end vertex and v_2 is a central vertex (an edge color in I_C), and another path in which the functions of vertices v_1 and v_2 are reversed. Routing R_C thus fulfills the following conditions: $\mathtt{d}(R_C) \leq 2$ and $\pi(R_C) \leq 2$. Consequently, each path of R_C may only conflict with at most two other paths, and the conflict graph $Q(R_C)$ is of degree bounded by $\Delta_{Q(R_C)} \leq 2$. Graph $Q(R_C)$ is thus a set of isolated vertices, paths and cycles. Notice that the three vertex cycle C_3 is a connected component of $Q(R_C)$ only if some three paths form a triangle, i.e. $P_1, P_2, P_3 \in R_C$ and $P_1 = (\{v_1, v_3\}, \{v_3, v_2\})$, $P_2 = (\{v_2, v_1\}, \{v_1, v_3\})$, $P_3 = (\{v_3, v_2\}, \{v_2, v_1\})$, for some three vertices $v_1, v_2, v_3 \in V$. Such a structure may however be easily eliminated by removing paths P_1, P_2, P_3 from R_C and replacing them by the following three paths: $P_1' = (\{v_1, v_2\})$, $P_2' = (\{v_2, v_3\})$, $P_3' = (\{v_3, v_1\})$, which satisfy the same set of requests and whose conflict graph consists of three isolated vertices.

The sought suboptimal solution R_S to the MaxEDP problem is now obtained by indicating a maximum independent set R_S in conflict graph $Q(R_C)$. Graph $Q(R_C)$ has $|R_C|$ vertices, and once all cycles C_3 have been eliminated the independent set R_S consists of at least $\frac{2}{5}|R_C|$ vertices (or equivalently, $|I_S| \geq \frac{2}{5}|I_C|$). Therefore, we finally obtain the following bound:

$$\frac{|I_{\mathrm{OPT}}|}{|I_S|} \leq \frac{|I^*|}{|I_S|} = \frac{|I^*|}{|I_C|}\frac{|I_C|}{|I_S|} \leq \frac{3}{2} \cdot \frac{5}{2} = 3.75$$

which completes the proof of the approximation ratio of the designed algorithm. □

It is interesting to note that although the off-line MaxEDP problem in complete graphs is *NP*-hard even for relatively small instances (Theorem 3), the conjecture that it is *APX*-hard still remains open [4], and the only inapproximability result concerns the on-line problem (Theorem 8). In fact, in the following subsection we show that for all instances of sufficiently bounded size, the off-line MaxEDP problem is not *APX*-hard.

2.3 Problem Complexity for Bounded Instances

We now deal with the MaxEDP problem restricted to instances I such that $|I| < |V|^s$ for some $s < 2$, and study the increasing difficulty of the problem with the increase of the bound on $|I|$.

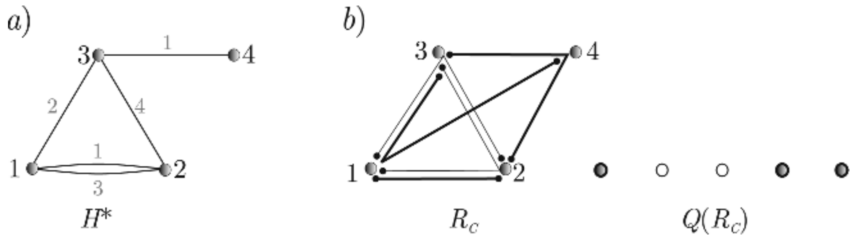

Fig. 1. Construction of an approximate solution to the MaxEDP problem for instance $I = \{\{1,2\},\{1,2\},\{1,3\},\{2,3\},\{3,4\}\}$ in complete graph K_4: *a*) an edge coloring of multigraph H^* (in the considered case $I_C = I^* = I$), *b*) a routing R_C of instance I_C in graph K_4 and its conflict graph $Q(R_C)$ (the independent set of paths forming the sought routing R_S is marked in bold)

Theorem 2. *An optimal solution to the* MaxEDP *problem in complete graphs can be determined in $O(|V|^3)$ time if the size of the input instance is bounded by $|I| \leq k|V|$, for any constant value of parameter $k > 0$.*

Proof. Let $T \subseteq V$ be defined as the set of all vertices belonging to more than $\frac{|V|}{24}$ requests, $T = \{v \in V : \deg_{H(I)} v > \frac{|V|}{24}\}$. Suppose that $|V| \geq 1248k$ (the problem for all smaller graphs may be solved by exhaustive search).

Property. The size of the solution $R_{\text{OPT}} \in \mathcal{R}(I_{\text{OPT}})$ to the MaxEDP *problem for instance I remains unchanged even if paths need not be disjoint with respect to edges from the edge set E^* of subgraph $K_{V\setminus T} \subseteq K_V$.* Indeed, let R^* be a routing of a maximal possible instance $I^* \subseteq I$ such that no edge from $E \setminus E^*$ belongs to more than one path of R^*. We create an instance I^{**} in graph $K_{V\setminus T}$ by successively considering all paths $P \in R^*$, and adding to I^{**} a request consisting of the first and the last vertex from $V \setminus T$ which appears in P. We now proceed to establish that it is possible to reroute instance I^{**} in $K_{V\setminus T}$ using edge disjoint paths, leading to the conclusion that $I_{\text{OPT}} = I^*$ is a valid solution to MaxEDP in K_V. Let $v \in V \setminus T$ be arbitrarily chosen. By definition of set T, we have $\deg_{H(I^*)} v \leq \deg_{H(I)} v \leq \frac{|V|}{24}$. Since each vertex $v \in V \setminus T$ is obviously connected to at most $|T|$ vertices from T, we immediately have $\deg_{H(I^{**}\setminus I^*)} v \leq |T|$. Combining the last two inequalities we obtain $\deg_{H(I^{**})} v \leq \frac{|V|}{24} + |T|$. Since $2k|V| \geq 2|I| = \sum_{v\in V} \deg_{H(I)} v \geq \sum_{v\in T} \deg_{H(I)} v \geq \frac{|V|}{24}|T|$, we have $|T| \leq 48k$. Taking into account the assumption $|V| \geq 1248k$, we finally obtain $\deg_{H(I^{**})} v \leq \frac{|V|}{24} + |T| \leq \frac{|V|}{24} + 48k = \frac{1}{12}(\frac{|V|+1248k}{2} - 48k) \leq \frac{1}{12}|V \setminus T|$, which means that by Proposition 3 instance I^{**} can be routed in $K_{V\setminus T}$ by means of edge disjoint paths, closing the proof of the property.

Let $I' \subseteq I$ denote the set of requests from I with at least one vertex in T, and let $I'_{\text{OPT}} \subseteq I'$ be a maximal subset of I' which can be routed by edge disjoint paths in K_V. The instance $I'_{\text{OPT}} \cup (I \setminus I') \subseteq I$ is therefore a maximal subset of I which can be routed by paths conflicting only within the edge set of graph $K_{V\setminus T}$, and by the proven Property such a routing can be converted to a correct solution to the MAXEDP problem for I. The problem of finding $I_{\text{OPT}} \subseteq I$ is thus reduced to finding $I'_{\text{OPT}} \subseteq I'$. Furthermore, when considering instance I' all vertices from set $V \setminus T$ may be regarded as indistinguishable (after once more relaxing the edge disjointness condition within $K_{V\setminus T}$). Thus graph K_V may be reduced to the multigraph G' formed by connecting each of the vertices of K_T to one additional vertex u using exactly $|V \setminus T|$ edges. In order to solve the MAXEDP problem for instance I' in G', we consider all possible arrangements of paths in the edge set of K_T, taken over all routings of all subsets of instance I'. Note that the number of such arrangements is bounded, since $|T| \in O(1)$. For a fixed arrangement of paths in the edge set of K_T, the MAXEDP problem for instance I' in G' can be easily reduced to the MAXEDP problem for a related instance in the multistar $G' \setminus K_T$. The latter problem can in turn be solved in $O(|I|^3) = O(|V|^3)$ time, using a generalisation of a technique from [7] (the solution proceeds by reduction to the problem of finding a maximal $[0, |V \setminus T|]$-factor in a multigraph). This procedure determines the complexity of the entire algorithm; the final rerouting step within $K_{V\setminus T}$ only requires $O(|V||I|) = O(|V|^2)$ time by Proposition 3. □

Theorem 3. *The* MAXEDP *problem in complete graphs is NP-hard even for instances of size bounded by $|I| \leq |V|^s$, for any value of parameter $s > 1$.*

Proof (sketch). The proof proceeds by reduction from the MAXEDP problem in complete graphs with cardinality restriction $|I| \leq |V|^2$, which was shown to be *NP*-hard in [8]. Let $s = 1 + \varepsilon$, $\varepsilon > 0$. Consider an arbitrary subset of vertices $V' \subseteq V$ of cardinality equal to at most $|V|^\varepsilon$. Let I' be any instance in $K_{V'}$. We define instance I in K_V as follows: $I = I' \cup \{\{u, v\} : u \in V', v \in V \setminus V'\}$; for sufficiently large $|V|$ we have $|I| \leq |V|^s$. The proof is complete when we observe that an optimal solution R_{OPT} to the MAXEDP problem for instance I in graph K_V is always equal to the union of two sets of paths: the set of all one-edge paths connecting vertices from $K_{V'}$ with vertices from $K_{V\setminus V'}$, and some optimal solution R'_{OPT} to the MAXEDP problem for instance I' in graph $K_{V'}$. In particular, we have: $|R_{\text{OPT}}| = |R'_{\text{OPT}}| + |V'|(|V| - |V'|)$. □

Theorem 4. *The* MAXEDP *problem in complete graphs admits a polynomial time approximation scheme for instances of size bounded by $|I| \leq |V|^s$, for any value of parameter $s < 2$.*

Proof (sketch). Let $|I| = |V|^s$, where $s = 2 - \varepsilon$, $\varepsilon > 0$. The proof is in essence similar to that of Theorem 2. We adopt the same definition of set T, obtaining $|T| \leq 48|V|^{1-\varepsilon}$. In all considerations we assume $|V| \geq 1248^{\frac{1}{\varepsilon}}$, so that the Property stated in the proof of Theorem 2 also holds in this case. By this property, any subset of instance I such that each vertex from T is the endpoint of at most $|V \setminus T|$ paths can be routed in K_V using edge disjoint paths. This implies that any maximal $[0, |V \setminus T|]$-factor in $H(I)$ is a suboptimal solution I_S to the considered MAXEDP problem. On the other hand, the cardinality of the optimal solution $|I_{\text{OPT}}|$ is bounded from above by the size of the maximal $[0, |V| - 1]$-factor in $H(I)$ by Corollary 1. The sizes of the considered factors in $H(I)$ are closely related, which leads to the following bound: $\frac{|I_{\text{OPT}}|}{|I_S|} \leq \frac{|V|}{|V|-2|T|} \leq \frac{1}{1-96|V|^{-\varepsilon}}$. Thus, for any

$\delta > 0$ the considered approach achieves an approximation ratio of $1+\delta$ provided $|V| > (96(1+\max\{12, \delta^{-1}\}))^{\frac{1}{\varepsilon}}$, whereas the problem may be optimally solved by exhaustive search for all smaller values of $|V|$. □

A summary of the main results of the section is given in Table 1.

3 The On-Line MaxEDP Problem in Complete Graphs

On-line algorithms for the MaxEDP problem, which are considered in this paper, are treated as a special case of greedy algorithms. We assume that successive requests from instance I appear sequentially at input, becoming known to the algorithm only once the previous request has been processed. The decision taken at every step as to whether some path fulfilling the current request should be added to the constructed edge disjoint routing R_S is inadvertent and impossible to change at a later stage of the algorithm. Approximation ratios are calculated with respect to the best possible solution R_{OPT} in the off-line model.

3.1 An On-Line 6.47-Approximation Algorithm

A slight modification of the approximation algorithm provided for the off-line case (Theorem 1) allows for its on-line operation. In the considered approach, the algorithm sequentially processes requests from instance I, treating them as edges of multigraph $H(I)$, and at every step attempts to color the edge using a color from the range $\{1, \ldots, |V|\}$. A generalization of this problem was recently considered by Favrholdt and Nielsen [9], under the name of the *maximum k-edge-colorable subgraph problem* for a multigraph. They stated that any fair on-line algorithm (i.e. an algorithm which always colors an edge, if only a color from the range $\{1, \ldots, k\}$ is available) leads to a $\frac{1}{2\sqrt{3}-3}$-approximation of the solution. In fact, the obtained result was significantly stronger; we shall reformulate it here for easier use in further considerations.

Theorem 5 ([9]). *For any multigraph $H = (V, I)$, any fair on-line algorithm for the k-edge-colorable subgraph problem labels a subset of edges $I_C \subseteq I$ with colors $\{1, \ldots, k\}$, such that $|I_C| \geq (2\sqrt{3}-3)|I^{**}|$, where I^{**} denotes a maximal $[0, k]$-factor in H.*

In particular, the above theorem holds for $k = |V|$, thus using the notation from Theorem 1 we may write $|I_C| \geq (2\sqrt{3}-3)|I^*|$. As the coloring proceeds, the sought routing R_S may be incrementally constructed using an on-line independent set algorithm applied to graph $Q(R_C)$. Since graph $Q(R_C)$ only consists of cycles, paths and isolated vertices, we obtain $|I_S| \geq \frac{1}{3}|I_C|$. Combining the obtained relations leads to the bound:

$$\frac{|I_{\text{OPT}}|}{|I_S|} \leq \frac{|I^*|}{|I_S|} = \frac{|I^*|}{|I_C|}\frac{|I_C|}{|I_S|} \leq \frac{1}{2\sqrt{3}-3} \cdot 3 < 6.47$$

which may be expressed by means of the following statement.

Corollary 2. *There exists a* 6.47*-approximation algorithm for the on-line* MaxEDP *problem in complete graphs, requiring* $O(|V|)$ *time to process a single request. The dilation of the returned solution is not greater than* 2*.*

In fact, the algorithm resulting from the above considerations can be written in much simpler form, as described in the next subsection.

3.2 Performance Analysis of the BGA Algorithm

The *bounded length greedy algorithm* (BGA) is an on-line strategy for the MaxEDP problem, introduced in [13]. The basic principle of its operation is that at every step an attempt is made to route the current request by the shortest possible path P which does not contain any of the edges already belonging to R_S, and to add P to the solution R_S provided $|P| \leq L$, where L is a fixed parameter of the algorithm. The computed routing R_S therefore fulfills the bound $\mathtt{d}(R_S) \leq L$. The BGA strategy was last studied by Carmi, Erlebach and Okamoto [4], who bounded its approximation ratio for $L = 4$ using an unsplittable flow technique.

Theorem 6 ([4]). *The* BGA *strategy with* $L = 4$ *is a* 9*-approximation on-line algorithm for the* MaxEDP *problem in complete graphs.*

However, it is interesting to note that further bounding of the parameter L may lead to algorithms for which a better approximation ratio can be proven.

Theorem 7. *The* BGA *strategy with* $L = 2$ *is a* 6.47*-approximation on-line algorithm for the* MaxEDP *problem in complete graphs.*

Proof (sketch). The proof is based on the observation that each step of BGA with $L = 2$ combines the properties of an on-line algorithm for the edge-colorable subgraph problem with those of an on-line independent set algorithm, thus implementing an approach very similar to that described in Subsection 3.1. A request $\{u, v\}$ can only be routed using BGA by a path $P = (\{u, w\}, \{w, v\})$ of length at most 2 via some vertex $w \in V$ if edge $\{u, v\}$ of multigraph $H(I)$ can be labeled with color $w \in \{1, \ldots, |V|\}$, and if path P does not conflict with any paths previously added to R_S. The only difference is that the $|V|$-edge-colorable subgraph of $H(I)$ implicitly found by the BGA algorithm need not correspond to that obtained by means of any fair algorithm, since in a step of BGA an edge of $H(I)$ is not colored whenever any color assignment is possible, but only when assigning a color contributes to the size of the resultant solution R_S. Careful analysis shows that this does not affect the overall approximation ratio which remains equal to 6.47 (Corollary 2). □

A further interesting property of the BGA strategy with parameter $L = 2$ is that it finds an edge disjoint routing of the whole instance I in the cases considered in Propositions 2 and 3.

Corollary 3. *If* $\Delta_{H(I)} \leq \frac{|V|}{12}$*, or* $|I| \leq |V| - 1$*, then the entire instance* I *can be routed in* G_V *by edge disjoint paths, and an optimal solution such that* $\mathtt{d}(R_{\text{OPT}}) \leq 2$ *is always determined by the* BGA *strategy with* $L = 2$*.*

3.3 Inapproximability Results

Whereas the complexity of finding a solution to the off-line MaxEDP problem in complete graphs still remains open, we now show that the on-line version is not approximable within a constant factor for sufficiently large instances.

Theorem 8. *There does not exist any on-line approximation algorithm for the* MaxEDP *problem in complete graphs with an approximation ratio smaller than* 1.50*, even when considering instances of size* $|I| < k|V|$*, for any* $k \geq 3$*.*

Proof. By contradiction, suppose that some on-line MaxEDP algorithm A has an approximation ratio not worse than 1.50. Given any graph K_V, let instance I begin with $|V| - 1$ requests of the form $\{u, v\}$, for some two distinguished vertices $u, v \in V$. At this point the routing R_S obtained by algorithm A consists of p paths, where $p \geq \frac{2}{3}(|V| - 1)$ (otherwise the instance is ended, and we have $|R_{\text{OPT}}| = |V| - 1 > 1.50|R_S|$). Instance I is now completed by presenting a further $2(|V| - 2)$ requests of the form $\{u, w\}$ and $\{v, w\}$, taken over all vertices $w \in V \setminus \{u, v\}$. Since the number of paths which end in any vertex (in particular, u or v) cannot exceed $|V| - 1$, the total number of paths eventually belonging to R_S is bounded by $|R_S| \leq p + 2((|V| - 1) - p) \leq \frac{4}{3}(|V| - 1)$, whereas $|R_{\text{OPT}}| = 2(|V| - 2) + 1 = 2(|V| - 1) - 1$, hence the ratio $\frac{|R_{\text{OPT}}|}{|R_S|}$ cannot be smaller than 1.50 for arbitrarily large values of $|V|$. □

Even in the on-line model, the gap remaining between the 1.50 inapproximability result of Theorem 8 and the 6.47-approximation algorithm from Theorem 7 is quite substantial. A partial attempt to bridge it may be performed by considering the inapproximability of specific classes of on-line algorithms. For example, the BGA algorithm and similar strategies are never better than 3-approximate for certain classes of instances [4].

4 Final Conclusions

The technique adopted in the proof of Theorem 1 — which may basically be thought of as *routing by edge coloring* — provides efficient approximation

Table 2. A comparison of presented approximation algorithms for the MaxEDP problem in complete graphs with previous results (updated from [4])

Principle of operation	*Model*	*Approximation ratio*	*Dilation*	*Reference*
Shortest-path-first variant of BGA	off-line	54		[8], 2001
Set tripartition	off-line	27		[8], 2001
BGA with $L = 4$	on-line	17	≤ 4	[13], 2002
BGA with $L = 4$	on-line	9	≤ 4	[4], 2003
BGA with $L = 2$	on-line	6.47	≤ 2	Thm. 7
Routing by edge coloring	off-line	3.75	≤ 2	Thm. 1

algorithms for a number of routing problems in complete graphs and similar extremely dense topologies. When applying this approach, the approximation ratio may vary depending on the considered problem, and is usually given in the form of the product of two parameters $M_1 \cdot M_2$, where M_1 denotes the relative loss in the first phase of the algorithm (determining an edge coloring), and M_2 is the relative loss in the second phase (post-processing the edge coloring).

For the MaxEDP problem, the applied techniques constitute a substantial improvement on earlier results (Table 2). We now give two more examples of routing problems for which fixed-ratio approximation algorithms can be similarly obtained.

The edge load routing problem. For a given instance I in graph K_V, we consider the problem of finding a routing $R_{\text{OPT}} \in \mathcal{R}(I)$, such that edge load $\pi(R_{\text{OPT}})$ is the minimum possible [2, 3]. In order to construct an approximation approach with respect to $\pi(R_S)$ within K_V, observe that multigraph $H(I)$ can always be efficiently edge-colored with at most $1.5(|V|-1)\,\pi(R_{\text{OPT}})$ colors in the off-line model, or $2(|V|-1)\,\pi(R_{\text{OPT}})$ colors in the on-line model. By applying a similar approach as that in the proof of Theorem 2, it is easy to see that the instance corresponding to any $(|V|-1)$-edge-colorable subgraph of $H(I)$ can always be routed with load at most 2, both in the off-line and the on-line model. Thus we have $M_2 = 2$ and $M_1 = 1.5$ (off-line) or $M_1 = 2$ (on-line), finally obtaining an off-line 3-approximation algorithm and an on-line 4-approximation algorithm for edge load routing in complete graphs.

The WDM wavelength count routing problem. This modification of the edge load routing problem is of special importance from the point of view of application in so called *all-optical wavelength division multiplexing* (WDM) networks [2, 5, 6]. For a given instance I in graph K_V, the sought routing $R_{\text{OPT}} \in \mathcal{R}(I)$ should minimize the value of a parameter called *WDM wavelength count* $\mathtt{w}(R_{\text{OPT}})$, defined as the chromatic number of conflict graph $Q(R_{\text{OPT}})$. The proposed construction of an approximation algorithm with respect to $\mathtt{w}(R_S)$ is nearly the same as for bounded edge load, the only difference being that in the second stage of the algorithm $(|V|-1)$-edge-colorable subgraphs of $H(I)$ can always be routed using 3 wavelengths. Therefore in this case we have $M_2 = 3$ and $M_1 = 1.5$ (off-line) or $M_1 = 2$ (on-line), yielding an off-line 4.5-approximation algorithm and an on-line 6-approximation algorithm for the considered problem.

Finally, let us remark on a general property of the approximate solutions obtained using the proposed approach: in all cases the dilation is bounded by a value of 2. Using paths with at most 1 intermediary node between the communicating pair of endpoints is advantageous from the point of view of resource usage, and additionally simplifies the routing process. Indeed, if the on-line version of the routing algorithm is considered in a distributed setting, each node can independently decide whether it may participate in the routing of a given communication request. Thus each request can be processed in $O(1)$ synchronous rounds, achieving a time-optimal routing process.

Acknowledgement. The author would like to express his gratitude to the anonymous referees for their numerous helpful comments and suggestions for the improvement of this paper.

References

1. M. Andrews, L. Zhang, Hardness of the undirected edge-disjoint paths problem. *Proc. STOC'05* (2005), 276–283.
2. B. Beauquier, J.C. Bermond, L. Gargano, P. Hell, S. Pèrennes and U. Vaccaro, Graph problems arising from wavelength routing in all-optical networks. *Proc. WOCS'97* (1997), Geneve, Switzerland.
3. J. Białogrodzki, Path Coloring and Routing in Graphs. In: *Graph Colorings*, M. Kubale ed., Contemporary Math. 352, AMS (2004), USA, 139–152.
4. P. Carmi, T. Erlebach, Y. Okamoto, Greedy edge-disjoint paths in complete graphs. *Proc. WG'03*, LNCS 2880 (2003), 143–155.
5. S. Choplin, L. Narayanan, J. Opatrny, Two-Hop Virtual Path Layout in Tori. *Proc. SIROCCO'04*, LNCS 3104 (2004), 69–78.
6. T. Erlebach, K. Jansen, The complexity of path coloring and call scheduling. *Theoret. Comp. Sci.* 255 (2001), 33–50.
7. T. Erlebach, K. Jansen, The Maximum Edge-Disjoint Paths Problem in Bidirected Trees. *SIAM J. Discret. Math.* 14 (2001), 326–355.
8. T. Erlebach, D. Vukadinović, New results for path problems in generalized stars, complete graphs, and brick wall graphs. *Proc. FCT'01*, LNCS 2138 (2001), 483–494.
9. L.M. Favrholdt, M.N. Nielsen, On-line edge-coloring with a fixed number of colors. *Algorithmica* 35 (2003), 176–191.
10. S. Fiorini, R.J. Wilson: *Edge-Colourings of Graphs*, Pittman (1977), USA.
11. H.N. Gabow, Data structures for weighted matching and nearest common ancestors with linking. *Proc. SODA'90* (1990), 434–443.
12. V. Guruswami *et al*, Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. *J. Comput. Syst. Sci.* 67 (2003), 473–496.
13. P. Kolman, C. Scheideler, Improved bounds for the unsplittable flow problem. *Proc. SODA'02* (2002), 184–193.
14. B. Ma, L. Wang, On the inapproximability of disjoint paths and minimum Steiner forest with bandwidth constraints, *J. Comput. Syst. Sci.* 60 (2000), 1–12.
15. A. Srinivasan, Improved approximations for edge-disjoint paths, unsplittable flow, and related routing problems, *Proc. FOCS'97* (1997), 416–425.
16. W.T. Tutte, A short proof of the factor theorem for finite graphs. *Canad. J. Math.* 6 (1954), 347–352.