

The maximum edge-disjoint paths problem in complete graphs

Adrian Kosowski

Department of Algorithms and System Modeling, Gdańsk University of Technology, Narutowicza 11/12, 80-952 Gdańsk, Poland

Abstract

In this paper, we consider the undirected version of the well known maximum edge-disjoint paths problem, restricted to complete graphs. We propose an off-line 3.75-approximation algorithm and an on-line 6.47-approximation algorithm, improving the earlier 9-approximation algorithm proposed by Carmi, Erlebach, and Okamoto [P. Carmi, T. Erlebach, Y. Okamoto, Greedy edge-disjoint paths in complete graphs, in: Proc. 29th Workshop on Graph Theoretic Concepts in Computer Science, in: LNCS, vol. 2880, 2003, pp. 143–155]. Moreover, we show that for the general case, no on-line algorithm is better than a $(2 - \varepsilon)$ -approximation, for all $\varepsilon > 0$. For the special case when the number of paths is within a linear factor of the number of vertices of the graph, it is established that the problem can be optimally solved in polynomial time by an off-line algorithm, but that no on-line algorithm is better than a $(1.5 - \varepsilon)$ -approximation. Finally, the proposed techniques are used to obtain off-line and on-line algorithms with a constant approximation ratio for the related problems of edge congestion routing and wavelength routing in complete graphs.

© 2008 Elsevier B.V. All rights reserved.

Keywords: Graph algorithms; Edge-disjoint paths; Complete graph

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 E represents the set of connections between them. A sequence of edges $P = (e_1, e_2, \dots, e_l) \in E^l$, 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

E-mail address: kosowski@sphere.pl.

$l = |P|$ in G , with endpoints v_1 and v_{l+1} . The symbol $\mathcal{P}_{\{u,v\}}$ is used to denote the set of all paths 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 the vertex set R and edges connecting all pairs of vertices corresponding to those 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 elements $\{u, v\} \in I$ and paths $P \in R$ with endpoints u and v , $P \in \mathcal{P}_{\{u,v\}}$. 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* $d(R)$, defined as the length of the longest path in routing R : $d(R) = \max_{P \in R} |P|$,
- *edge congestion* $\pi(R)$, given by the formula: $\pi(R) = \max_{e \in E} \pi_e(R)$, where $\pi_e(R) = |\{P \in R : e \in P\}|$,
- *wavelength count* $w(R)$, defined as the chromatic number of the conflict graph: $w(R) = \chi(Q(R))$; the inequality $w(R) \geq \pi(R)$ holds for any routing [3].

A routing R is said to consist of *edge-disjoint paths* if $\pi(R) = 1$, or equivalently, if the 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 edge-disjoint paths R_S , such that $R_S \in \mathcal{R}(I_S)$ for some instance $I_S \subseteq I$.

Goal: Maximise the cardinality of R_S .

Notation. Throughout the paper, the complete graph with vertex set V is denoted by K_V . Unless otherwise stated, we will assume that the MAXEDP problem is considered for the instance I in the 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 cardinality not greater 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 number of edges incident to a vertex v in multigraph H is called its degree and denoted as $\deg_H v$. The symbols Δ_H and χ'_H are used to denote the maximum vertex degree and the chromatic index of the multigraph H , respectively.

State-of-the-art results. In the case of general networks, the MAXEDP problem is closely related to a family of unsplittable flow problems. As a consequence MAXEDP is NP-hard, difficult to approximate in polynomial time within 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 difficult even to approximate within $O(|E|^{\frac{1}{2}-\varepsilon})$, for any $\varepsilon > 0$ [9]. Both the directed and undirected versions are approximable within a factor of $O(|E|^{\frac{1}{2}})$ [14].

When the graph G is the complete graph K_V , the MAXEDP problem, though still 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 models of computation, owing to Carmi, Erlebach, and Okamoto [4]. A comparison of the known approximation algorithms is provided in Table 1.

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 general case of 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 a 2-approximation. A summary of the most important new results concerning the MAXEDP problem is given in Table 2. 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.



Table 1

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

| Principle of operation | Model | Approx. ratio | Dilation | Reference |
|------------------------------------|----------|---------------|----------|------------|
| Shortest-path-first variant of BGA | Off-line | 54 | | [5], 2001 |
| Set tripartition | Off-line | 27 | | [5], 2001 |
| BGA with $L = 4$ | On-line | 17 | ≤ 4 | [10], 2002 |
| BGA with $L = 4$ | On-line | 9 | ≤ 4 | [4], 2003 |
| BGA with $L = 2$ | On-line | 6.47 | ≤ 2 | Theorem 17 |
| Routing by edge coloring | Off-line | 3.75 | ≤ 2 | Theorem 6 |

Table 2

New complexity results for the MAXEDP problem in complete graphs

| Instance restriction | Off-line complexity | | On-line complexity | |
|--------------------------------------|---------------------|--------------------|-------------------------------------------------|--------------------------|
| $\Delta_{H(I)} \leq \frac{ V }{12}$ | $O(V I)$ | Proposition 4 | $O(V)$ per request | Corollary 18 |
| $ I < V $ | $O(V)$ | Proposition 3 | $O(V)$ per request | Corollary 18 |
| $ I < k V $, const k | $O(V ^2)$ | Theorem 10 | no $(1.5 - \varepsilon)$ -approx. | Theorem 19 |
| $ I < V ^s$, const $s \in (1, 2)$ | PTAS, NPH | Theorems 12 and 13 | | |
| General case | 3.75-approx. | Theorem 6 | 6.47-approx. no $(2 - \varepsilon)$ -approx. | Theorem 17 Theorem 20 |

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 $[a, b]$. In further considerations we use the following result.

Proposition 1 ([8]). *The problem of finding an $[a, b]$ -factor with the maximum possible number of edges in multigraph $H = (V, I)$ can be solved in $O(|V||I| \log |V|)$ 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 2. *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 characterization of two classes useful in further considerations is given below.

Proposition 3. *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 $d(R_{\text{OPT}}) \leq 2$, 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 minimum 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 $I \setminus \{\{v_1, v_2\}\}$ in the complete graph $K_{V \setminus \{u\}}$. 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 the instance $I \setminus \{\{u, v\}\}$ in $K_{V \setminus \{u\}}$. The described approach may easily be implemented in the form of an algorithm with $O(|V|)$ time complexity. \square

Observe that the claim of Proposition 3 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 10).

Proposition 4. *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 $d(R_{\text{OPT}}) \leq 2$, can be determined in $O(|V||I|)$ time.*

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

1. Routing R_{OPT} consists of edge-disjoint paths only.
2. The length of any path added to R_{OPT} is at most 2.
3. Each vertex of graph K_V is the center of at most $\frac{|V|}{12}$ paths of length 2 in R_{OPT} .

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. \square

2.2. An off-line 3.75-approximation algorithm

The idea of the approximation algorithms presented in this paper is based on the following observation.

Lemma 5. *Given an instance I in graph K_V and an edge-coloring of multigraph $H(I)$ using at most $|V|$ colors, in $O(|I|)$ time it is possible to determine:*

- (i) routing $R \in \mathcal{R}(I)$ such that $d(R) \leq 2$, $\pi(R) \leq 2$, and $w(R) \leq 3$,
- (ii) routing $R_S \in \mathcal{R}(I_S)$ for some $I_S \subseteq I$, such that $|I_S| \geq \frac{2}{5}|I|$, $d(R_S) \leq 2$, and R_S uses edge-disjoint paths only.

Proof. Let c_e denote the color assigned to an edge $e \in I$ in the given edge-coloring of $H(I)$. Since c_e is a value from the range $[1, |V|]$, it may be treated as an identifier of some vertex in graph K_V . Define routing R of instance I in graph K_V as follows: $R = \{(\{v_1, c_e\}, \{c_e, v_2\}) : e = \{v_1, v_2\} \in I\}$ (see Fig. 1 for an illustration). No vertex of $H(I)$ may ever be incident to two edges from I of the same color; therefore each edge $\{v_1, v_2\}$ of graph K_V belongs to at most two paths of routing R — one path in which v_1 is an end vertex and v_2 is a central vertex (a color of an edge in I), and another path in which the roles of vertices v_1 and v_2 are reversed. Routing R thus fulfills the following conditions: $d(R) \leq 2$ and $\pi(R) \leq 2$. Consequently, each path of R may only conflict with at most two other paths, and the conflict graph $Q(R)$ is of degree bounded by $\Delta_{Q(R)} \leq 2$. Graph $Q(R)$ is thus a set of isolated vertices, paths and cycles; hence we immediately have $w(R) = \chi(Q(R)) \leq 3$, which completes the proof of clause (i).

Next, notice that the three vertex cycle C_3 is a connected component of $Q(R)$ only if some three paths form a triangle, i.e. $P_1, P_2, P_3 \in R$ 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

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.




Fig. 1. An illustration of Lemma 5 for instance $I = \{(1, 2), \{1, 2\}, \{1, 3\}, \{2, 3\}, \{3, 4\}\}$ in the complete graph K_4 : (a) an edge coloring of multigraph $H(I)$, (b) a routing R of instance I in graph K and its conflict graph $Q(R)$ (the independent set of paths forming the sought routing R_S is marked in bold).

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 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 routing $R_S \in \mathcal{R}(I_S)$ sought in clause (ii) is now obtained by indicating a maximum independent set R_S in conflict graph $Q(R)$. Graph $Q(R)$ has $|R|$ vertices, and once all cycles C_3 have been eliminated, the independent set R_S consists of at least $\frac{2}{3}|R|$ vertices (or equivalently, $|I_S| \geq \frac{2}{3}|I|$). \square

Theorem 6. *There exists a 3.75-approximation algorithm with $O(|I||V| \log |V|)$ runtime for the MAXEDP problem in complete graphs. The dilation of the returned solution is not greater than 2.*

Proof. Let I be an arbitrary instance in the complete graph K_V , and let $I_{OPT} \subseteq I$ 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 $O(|I||V| \log |V|)$ time. Moreover, by Corollary 2 we have $|I_{OPT}| \leq |I^*|$.

We will now show that there exists an algorithm with $O(|I||V| \log |V|)$ 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, $|I_S| \geq \frac{|I^*|}{3.75} \geq \frac{|I_{OPT}|}{3.75}$. Instance I_S is constructed as a subset of the edge set of multigraph H^* . Since $\Delta_{H^*} < |V|$, by a well known result owing to Shannon [13], 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 $O(|I||V| \log |V|)$ time. Without loss of generality we may assume that the colors are labelled with integers from the range $\{1, \dots, \frac{3|V|}{2}\}$, 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 $\{1, \dots, |V|\}$. Due to the adopted ordering of the color labels, we immediately have $|I_C| \geq \frac{2}{3}|I^*|$. Moreover, since instance I_C fulfills the assumptions of Lemma 5, we can efficiently construct an instance $I_S \subseteq I_C$ such that $|I_S| \geq \frac{2}{3}|I_C|$ and I_S admits a routing R_S using edge-disjoint paths. Since $I_S \subseteq I_C \subseteq I^* \subseteq I$, we can treat R_S as the suboptimal solution to the MAXEDP problem, obtaining the following bound:

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

which completes the proof of the approximation ratio of the designed algorithm. \square

It is interesting to note that although the off-line MAXEDP problem in complete graphs is NP-hard even for relatively small instances (Theorem 13), the conjecture that it is APX-hard still remains open [4], and the only non-approximability result concerns the on-line problem (Theorem 20). 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 with relatively few requests, and study the increasing difficulty of the problem with the increase of the bound on $|I|$. Considerations start with some auxiliary notation and lemmas.

Let I be a given instance in the graph K_V and let set $T \subseteq V$ be defined as the set of all vertices belonging to more than $\frac{|V|}{24}$ requests, $T = \{v \in V : \deg_{H(I)} v > \frac{|V|}{24}\}$. The symbol E^* is used to denote the set of edges of subgraph $K_{V \setminus T} \subseteq K_V$.

Lemma 7. *The size of set T fulfills the bound $|T| \leq 48 \frac{|I|}{|V|}$.*

Proof. By the handshaking lemma we may write:

$$2|I| = \sum_{v \in V} \deg_{H(I)} v \geq \sum_{v \in T} \deg_{H(I)} v \geq \frac{|V|}{24} |T|,$$

and the claim follows directly. \square

Lemma 8. *Assume that $|I| \leq \frac{|V|^2}{1248}$. If an instance $I_S \subseteq I$ admits a routing $R_S \in \mathcal{R}(I_S)$ whose paths are edge-disjoint with respect to $E \setminus E^*$ (but may share edges in E^*), then I_S also admits a routing whose paths are edge-disjoint with respect to all the edges in E . Such a routing may be constructed in $O(|I||V|)$ time.*

Proof. Consider the following construction of an instance I^* in graph $K_{V \setminus T}$. For successive paths $P \in R_S$, we add to I^* a request consisting of the first and of the last vertex from $V \setminus T$ which appears in P (requests of the form $\{v, v\}$ are left out). Observe that it now suffices to show that I^* admits an edge-disjoint routing in $K_{V \setminus T}$, since then routing R_S may be appropriately modified to obtain the sought routing of I_S which is edge-disjoint in E . More precisely, we will prove that for all $v \in V \setminus T$ we have $\deg_{H(I^*)} v \leq \frac{1}{12} |V \setminus T|$, and therefore that the claim holds by Proposition 4.

Let $v \in V \setminus T$ be arbitrarily chosen. The construction of a request $\{u, v\} \in I^*$, for any $u \in V \setminus T$, made use of some path $P \in R_S$, which may have one of two possible arrangements:

1. Vertex v is an endpoint of P . In this case, path P was used to satisfy some request from instance I_S , containing vertex v .
2. Vertex v is not an endpoint of P . By definition of instance I^* , there must then exist an edge $\{v, w\}$ belonging to path P such that $w \in T$.

Thus, $\deg_{H(I^*)} v = d_1 + d_2$, where values d_1 and d_2 denote the numbers of requests $\{u, v\}$ corresponding to Cases 1 and 2 above, respectively. In order to bound from above the value d_1 , observe that from the description of Case 1 we have $d_1 \leq \deg_{H(I_S)} v$. Taking into account the definition of set T , we obtain:

$$d_1 \leq \deg_{H(I_S)} v \leq \deg_{H(I)} v \leq \frac{|V|}{24}.$$

An upper bound on the value of d_2 is obtained from the edge-disjointness condition for the set of paths R_S ; clearly, d_2 cannot exceed the number of connections between v and vertices from set T :

$$d_2 \leq |T|.$$

Combining the last two inequalities gives:

$$\deg_{H(I^*)} v = d_1 + d_2 \leq \frac{|V|}{24} + |T|. \quad (1)$$

Rewriting the assumption $|I| \leq \frac{|V|^2}{1248}$ as $\frac{|I|}{|V|} \leq \frac{|V|}{1248}$ and taking into account Lemma 7, we have $|T| \leq \frac{|V|}{26}$, which immediately implies the following relation:

$$\frac{|V|}{24} + |T| \leq \frac{|V|}{12} - \frac{|T|}{12} = \frac{|V \setminus T|}{12}. \quad (2)$$

Inequalities (1) and (2) directly lead to the sought bound $\deg_{H(I^*)} v \leq \frac{|V \setminus T|}{12}$.

Taking into account Proposition 4, the new edge-disjoint routing can be computed in a time proportional to the total length of all paths of R_S , which is bounded from above by $O(|I||V|)$. \square

We now consider the multigraph G' formed from K_V by contracting all vertices from $V \setminus T$ into a single vertex x . Formally, G' has vertex set $T \cup \{x\}$, all pairs of vertices from T are connected by single edges, and vertex x is

connected with each vertex from T by exactly $|V \setminus T|$ edges. Observe that there exists a natural bijection f from the subset of edges $E \setminus E^*$ of K_V to the set of all edges of multigraph G' which preserves the identifiers of vertices within the set T . The given instance I in K_V can also easily be converted into an instance I' in G' through a one-to-one correspondence of requests, by replacing all occurrences of vertices $v \in V \setminus T$ by occurrences of vertex x . Requests of the form $\{x, x\}$ which appear in I' should be treated as immediately fulfilled.

Lemma 9. *Assume that $|I| \leq \frac{|V|^2}{1248}$. A subset of requests $I_S \subseteq I$ admits an edge-disjoint routing in K_V if and only if the corresponding subset of requests $I'_S \subseteq I'$ admits an edge-disjoint routing in G' . Given either of these routings, the other may be computed in $O(|I||V|)$ time.*

Proof. (\Rightarrow) Let $R_S \in \mathcal{R}(I_S)$ be an edge-disjoint routing in K_V . For each path $P \in R_S$ in the graph K_V which satisfies some request from I_S , we construct a path P' in G' satisfying the corresponding request from I'_S , by performing the following operations. First, all edges from set E^* are removed from P . Next, transformation f is applied to the remaining edges of P to obtain a sequence of edges in G' . Finally, all cycles are removed from this sequence, giving the sought path P' . It is easy to see that the obtained routing R'_S does indeed satisfy instance I'_S . Moreover, since transformation f does not merge edges into each other, routing R'_S is clearly edge-disjoint in G' .

(\Leftarrow) Let $R'_S \in \mathcal{R}(I'_S)$ be an edge-disjoint routing in G' . For each path $P' \in R'_S$ in G' which satisfies some request from I'_S , the corresponding path P in K_V is obtained as follows. First, a transformation f^{-1} is applied to all edges of P' . The obtained sequence of edges is either the sought path P , or can be converted into P by adding exactly one edge from set E^* . The set of all paths P obtained in this way is some routing R_S which clearly satisfies instance I_S . Moreover, routing R_S is edge disjoint with respect to the set of edges $E \setminus E^*$. Thus, by Lemma 8, I_S also admits a routing which is edge disjoint with respect to all edges of K_V , which completes the proof of the equivalence.

Since both the above transformations can be performed in a time proportional to the total length of all paths of the routing, the $O(|I||V|)$ complexity bound follows directly. \square

Theorem 10. *An optimal solution to the MAXEDP problem in complete graphs can be determined in $O(|V|^2)$ time if the size of the input instance is bounded by $|I| \leq k|V|$, for any constant value of parameter $k > 0$.*

Proof. Without loss of generality we may assume that $|V| \geq 1248k$, since for smaller values of $|V|$ the problem can be solved by exhaustive search. Thus, the assumptions of Lemma 9 are fulfilled, and the problem is reduced to solving MAXEDP for the instance I' in multigraph G' . We now replace the parallel edges of multigraph G' by single edges with an associated capacity value corresponding to the original multiplicity of the edge. Multigraph G' reduces to the complete graph $K_{T \cup \{x\}}$ and the edge capacity function is given as $c(\{u, v\}) = 1$ for all $u, v \in T$, and $c(\{u, x\}) = |V \setminus T|$ for all $u \in T$. The considered MAXEDP problem may now be reformulated as follows: in graph $K_{T \cup \{x\}}$ we seek a maximum sized routing $R'_S \in \mathcal{R}(I'_S)$, where $I'_S \subseteq I'$, such that the number of paths using each edge does not exceed its capacity, $\pi_{\{u, v\}}(R'_S) \leq c(\{u, v\})$ for all $u, v \in T \cup \{x\}$. Such a network optimization problem admits a simple Integer Linear Programming formulation. For $u, v \in T \cup \{x\}$, let $d(\{u, v\})$ denote the number of times the request $\{u, v\}$ appears in instance I' . With each distinct path P in graph $K_{T \cup \{x\}}$ we associate a variable y_P , representing the number of times path P appears in the ILP solution, and state the problem as follows:

$$\begin{aligned} \text{Maximise: } & \sum_P y_P \\ \text{Subject to: } & \sum_{P: \{u, v\} \in P} y_P \leq c(\{u, v\}) \wedge \sum_{P \in \mathcal{P}_{\{u, v\}}} y_P \leq d(\{u, v\}), \forall u, v \in T \cup \{x\} \\ & y_P \in \mathbb{Z}, y_P \geq 0 \end{aligned}$$

Now, observe that by Lemma 7 we have $|T| \leq 48k$; hence graph $K_{T \cup \{x\}}$ has $O(1)$ vertices, and consequently also $O(1)$ distinct paths. Thus, the above ILP problem on $O(1)$ variables may be solved in $O(|V|)$ time by a well known result of Lenstra [11]. The overall $O(|V||I|) = O(|V|^2)$ time complexity of the solution is determined by the procedure in Lemma 9. \square

When constrained to instances I such that $|I| \leq |V|^s$, for any $1 < s < 2$, the MAXEDP problem becomes NP-hard (see Theorem 13), but it is possible to construct a polynomial time approximation scheme. Before showing the latter result we recall a simple observation about cardinalities of maximum factors in multigraphs (the proof is attributed to folklore).

Lemma 11. Let I_a and I_b be a maximum $[0, a]$ -factor and a maximum $[0, b]$ -factor in multigraph $H(I) = (V, I)$, respectively, where $2b > a \geq b$. Then $\frac{|I_a|}{|I_b|} \leq \frac{a}{2b-a}$.

Proof. Consider the $[0, b]$ -factor I_b^* obtained from I_a as follows. For each vertex $v \in V$, we define $(I_a \setminus I_b^*)_v \subseteq I_a$ as a set of exactly $\max\{0, (\deg_{H(I_a)} v) - b\}$ arbitrarily chosen edges adjacent to v in I_a . Factor I_b^* is constructed as $I_b^* = I_a \setminus \bigcup_{v \in V} (I_a \setminus I_b^*)_v$. We may write:

$$\frac{|I_b^*|}{|I_a|} = 1 - \frac{\left| \bigcup_{v \in V} (I_a \setminus I_b^*)_v \right|}{|I_a|} \geq 1 - \frac{\sum_{v \in V} |(I_a \setminus I_b^*)_v|}{\frac{1}{2} \sum_{v \in V} \deg_{H(I_a)} v} \geq 1 - 2 \max_{v \in V} \frac{|(I_a \setminus I_b^*)_v|}{\deg_{H(I_a)} v}.$$

For each vertex v , we either have $\deg_{H(I_a)} v < b$ and $|(I_a \setminus I_b^*)_v| = 0$, or $a \geq \deg_{H(I_a)} v \geq b$ and $|(I_a \setminus I_b^*)_v| \leq a - b$. Therefore for all vertices we obtain $\frac{|(I_a \setminus I_b^*)_v|}{\deg_{H(I_a)} v} \leq \frac{a-b}{a}$, and finally $\frac{|I_b^*|}{|I_a|} \geq \frac{2b-a}{a}$. Since $|I_b| \geq |I_b^*|$ and $2b - a > 0$, the claim follows directly. \square

Theorem 12. 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. Let $|I| = |V|^s$, where $s = 2 - \varepsilon$, $\varepsilon > 0$. As in the proof of Theorem 10, we assume that the size of set $|V|$ is sufficiently large (this time $|V| \geq 1248^{\frac{1}{\varepsilon}}$), so that the assumptions of Lemma 9 hold. Once again, we will consider the instance I' in multigraph G' . Let $I_S \subseteq I$ be any maximum $[0, |V \setminus T|]$ -factor in $H(I)$. Since for all $v \in V$ we have $\deg_{H(I_S)} v \leq |V \setminus T|$, for the corresponding instance I'_S in G' and all vertices $v \in T$, we also obtain $\deg_{H(I'_S)} v \leq |V \setminus T|$. Recall that in G' all vertices from set T are connected to vertex x by exactly $|V \setminus T|$ edges each; this means that I'_S can easily be routed by edge-disjoint paths of length at most 2 using only these edges. Thus, by Lemma 9, instance I_S admits an edge-disjoint routing in K_V , which may be determined in $O(|I||V|)$ time.

We now proceed to prove that I_S is a sufficiently good approximation of the optimal solution to MAXEDP for instance I . Indeed, I_S is a maximum $[0, |V \setminus T|]$ -factor in $H(I)$, while by Corollary 2 the cardinality of an optimal solution $|I_{OPT}|$ is bounded from above by the size of the maximum $[0, |V| - 1]$ -factor in $H(I)$. Further, note that $|T| \leq 48|V|^{1-\varepsilon}$ by Lemma 7. Putting values $a = |V| - 1$ and $b = |V \setminus T| = |V| - |T| = |V|(1 - 48|V|^{-\varepsilon})$ in Lemma 11, the assumption $2b > a \geq b$ is fulfilled (since $|V| \geq 1248^{\frac{1}{\varepsilon}} > 96^{\frac{1}{\varepsilon}}$), and we obtain:

$$\frac{|I_{OPT}|}{|I_S|} \leq \frac{a}{2b - a} \leq \frac{|V|}{2|V|(1 - 48|V|^{-\varepsilon}) - |V|} = \frac{1}{1 - 96|V|^{-\varepsilon}}.$$

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

Theorem 13. The MAXEDP problem in complete graphs is NP-hard even for instances of size bounded by $|I| \leq |V|^s$, for any value of the parameter $s > 1$.

Proof. The proof proceeds by a reduction from the MAXEDP problem in complete graphs with the cardinality restriction $|I^*| \leq |V^*|^2$, which was shown to be NP-hard in [5]. Let $s = 1 + \varepsilon$, $\varepsilon > 0$. Given instance I^* in graph K_{V^*} , in polynomial time we construct instance I and graph K_V as follows. Select V in such a way that $V^* \subseteq V$ and $|V^*| \leq |V|^\varepsilon$. Instance I is defined as: $I = I^* \cup \{\{u, v\} : u \in V^*, v \in V \setminus V^*\}$. Observe that:

$$|I| = |I^*| + |V^*| \cdot (|V| - |V^*|) = (|I^*| - |V^*|^2) + |V^*||V| \leq |V|^{1+\varepsilon},$$

Thus the condition $|I| \leq |V|^s$ is fulfilled. The proof is complete when we notice that an optimal solution R_{OPT} to the MAXEDP problem for the 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 the instance I^* in graph K_{V^*} . In particular, we have: $|R_{OPT}^*| = |R_{OPT}| - |V^*|(|V| - |V^*|)$. \square

A summary of the main results of the section is given in Table 2.

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 irrevocable 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 6) allows for its on-line operation. In the considered approach, the algorithm sequentially processes requests from instance I , treating them as edges of the multigraph $H(I)$, and at every step attempts to color the edge using a color from the range $\{1, \dots, |V|\}$. A more general version of this problem is called the *maximum k -edge-colorable subgraph problem*, where the goal is to color as many edges of the input multigraph as possible using colors $\{1, \dots, k\}$, where k is a number given at input. This problem was recently considered by Favrholt and Nielsen [6], who introduced the class of *fair on-line algorithms*, i.e. on-line algorithms which at every step consider a single edge e for coloring, and are required to assign some color from the range $\{1, \dots, k\}$ to e if at least one such color is available. They showed that any fair on-line algorithm leads to a $\frac{1}{2\sqrt{3}-3}$ -approximation of the solution. In fact, the obtained result is significantly stronger; we shall reformulate it here for easier use in further considerations.

Theorem 14 ([6]). *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 $\{1, \dots, k\}$, such that $|I_C| \geq (2\sqrt{3} - 3)|I^{**}|$, where I^{**} denotes a maximum $[0, k]$ -factor in H .*

In particular, the above theorem holds for $k = |V|$; thus using the notation from Theorem 6 we may write $|I_C| \geq (2\sqrt{3} - 3)|I^*|$. As the coloring proceeds, a routing R_C of instance I_C is naturally defined (see Lemma 5). The sought routing R_S may be incrementally constructed using the simple greedy 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:

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

The complexity of the on-line algorithm is determined by the coloring phase: for each request, we have to check whether any of the colors $\{1, \dots, |V|\}$ is available. This may be expressed in the form of the following statement.

Corollary 15. *There exists an on-line 6.47-approximation algorithm for the 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 a 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 [10]. 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 $d(R_S) \leq L$. 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 16 ([4]). *The BGA strategy with $L = 4$ is an on-line 9-approximation 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 17. *The BGA strategy with $L = 2$ is an on-line 6.47-approximation algorithm for the MAXEDP problem in complete graphs.*

Proof. The proof is based on the observation that for a given input instance I , any outcome of the routing process using BGA with $L = 2$ can also be reached by the algorithm described in Section 3.1. Given an ordering of requests in instance I , let I_1 be the set of requests which are assigned paths in the solution obtained by BGA and let I_2 be the set of all the remaining requests, $I = I_1 \cup I_2$. Consider the result of the algorithm from Section 3.1 for an ordering of requests of I which starts with all requests from I_1 , followed by all requests from I_2 . Without loss of generality, the on-line fair edge coloring subroutine may be defined so as to mimic the behavior of BGA, i.e. if a request $\{u, v\} \in I_1$ was routed by BGA using a path $P = (\{u, w\}, \{w, v\})$ of length at most 2 via some vertex $w \in V$, then the edge $\{u, v\}$ of multigraph $H(I)$ is labeled with color $w \in \{1, \dots, |V|\}$. The algorithm from Section 3.1 will therefore satisfy all requests in I_1 , obtaining exactly the same routing R_1 as that produced by BGA, and then will start processing the requests from I_2 . Let $\{u, v\} \in I_2$ be an arbitrary request. Since this request was not routed by BGA, it clearly means that there does not exist a path $P \in \mathcal{P}_{\{u,v\}}$ of length at most 2 which does not share any edges with paths from R_1 . Note, however, that the algorithm in Section 3.1 also only uses paths of length at most 2; therefore, it will not route this request either, and the routing obtained by both algorithms is exactly the same. This means that the worst-case performance of BGA is not worse than the worst-case performance of the algorithm from Section 3.1, hence the approximation ratio of BGA is at most 6.47 by Corollary 15. \square

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 3 and 4 (this is immediately clear from an analysis of their proofs).

Corollary 18. *If $\Delta_{H(I)} \leq \frac{|V|}{12}$, or $|I| \leq |V| - 1$, then the entire instance I can be routed in K_V by edge-disjoint paths, and an optimal solution such that $d(R_{\text{OPT}}) \leq 2$ is always determined by the BGA strategy with $L = 2$.*

3.3. Non-approximability results

Whereas the hardness of approximation of the off-line MAXEDP problem in complete graphs still remains an open question, we now show that the on-line version is not approximable within a factor of 1.5 even for instances of linear size with respect to the number of vertices of the graph, and is not approximable within a factor of 2 for general instances.

Theorem 19. *There does not exist any on-line approximation algorithm for the MAXEDP problem in complete graphs with an approximation ratio smaller than $1.5 - \varepsilon$ for any $\varepsilon > 0$, even when considering instances of size $|I| < k|V|$, for any $k \geq 3$.*

Proof. For contradiction, suppose that some on-line MAXEDP algorithm A has an approximation ratio smaller than $1.5 - \varepsilon$. 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 $|R_{\text{OPT}}| = |V| - 1 > 1.5|R_S|$). 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 $\frac{|R_{\text{OPT}}|}{|R_S|} \geq 1.5 - \frac{3}{4}(|V| - 1)^{-1}$ is not less than $1.5 - \varepsilon$ for sufficiently large values of $|V|$, $|V| \geq \frac{3}{4}\varepsilon^{-1} + 1$. \square

Theorem 20. *There does not exist any on-line approximation algorithm for the MAXEDP problem in complete graphs with an approximation ratio smaller than $2 - \varepsilon$, for any $\varepsilon > 0$.*

Proof. Let A be any on-line approximation algorithm for the MAXEDP problem. Assume that instance I in graph K_V is put forward by an adaptive adversary, which at every step arbitrarily selects an element from some set S of

distinct requests, presents it as a request to algorithm A, and updates set S depending on the outcome of the routing process. The instance is terminated when set S is empty. Specifically, let set S initially contain exactly one request $\{u, v\}$, for some two distinguished vertices $u, v \in V$. Suppose that the i -th element, $i \geq 1$, of instance I presented by the adversary is a request $\{u_i, v_i\} \in S$. If algorithm A routes request $\{u_i, v_i\}$ by some path P_i , then the adversary removes $\{u_i, v_i\}$ from S , and inserts all edges of P_i as requests into S . However, if algorithm A does not route request $\{u_i, v_i\}$, then the adversary removes $\{u_i, v_i\}$ from S only if $i \geq |V|$.

First, observe that instance I always terminates after a finite number of steps. Next, let $E_R \subseteq E$ denote the set of all edges belonging to paths of the routing R_S obtained by algorithm A for the instance I . By the definition of the insertion operation for set S , an edge $e \in E$ appears at least once as a request in instance I if and only if $e = \{u, v\}$ or $e \in E_R$. Consequently, we have $|R_{\text{OPT}}| \geq |E_R|$, since requests corresponding to edges from E_R can be routed by paths of length one. Next, observe that any request other than $\{u, v\}$ presented to algorithm A corresponds to an edge already belonging to a path of routing R_S ; hence R_S contains at most one path of length one. This in turn implies that $|E_R| \geq 2|R_S| - 1$, and consequently:

$$|R_{\text{OPT}}| \geq 2|R_S| - 1.$$

Finally, taking into account that throughout the first $|V|$ requests of I set S remains non-empty, we have $|I| \geq |V|$, and thus by Proposition 3 $|R_{\text{OPT}}| \geq |V| - 1$. We obtain the following expression:

$$\frac{|R_{\text{OPT}}|}{|R_S|} \geq 2 - \frac{1}{R_{\text{OPT}}} \geq 2 - \frac{1}{|V| - 1}.$$

Therefore, for all values of $|V| \geq \varepsilon^{-1} + 1$, we have $\frac{|R_{\text{OPT}}|}{|R_S|} \geq 2 - \varepsilon$, which completes the proof. \square

Even in the on-line model, the gap remaining between the 2-non-approximability result of Theorem 20 and the 6.47-approximation algorithm from Theorem 17 is quite substantial. A partial attempt to bridge it may be performed by considering the non-approximability 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 remarks

The technique adopted in the proof of Theorem 6—which may basically be thought of as *routing by edge coloring*—provides efficient approximation 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, see (Table 1). We now give two more examples of routing problems for which fixed-ratio approximation algorithms can be similarly obtained.

The edge congestion 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 congestion $\pi(R_{\text{OPT}})$ is the minimum possible [2,3]. Approximation ratios are considered in terms of upper bounds on the ratio $\frac{\pi(R_S)}{\pi(R_{\text{OPT}})}$.

Theorem 21. *There exists an off-line $(3 + \frac{1}{\text{OPT}})$ -approximation algorithm with $O(|I|^2)$ runtime for the edge congestion routing problem in complete graphs. The dilation of the returned solution is not greater than 2.*

Proof. Let I be an arbitrarily chosen instance in graph K_V , and let $R_{\text{OPT}} \in \mathcal{R}(I)$ be any routing of I with minimum possible edge congestion. Note that for each vertex $v \in V$, each of the $\deg_{H(I)} v$ paths with one endpoint in v contributes by exactly 1 to the load of some edge incident to v . Thus, by the pigeon-hole principle the following inequality holds for any routing of instance I , and therefore also for routing R_{OPT} : $\pi(R_{\text{OPT}}) \geq \lceil \frac{\Delta_{H(I)}}{|V|-1} \rceil$. We will now present an algorithm for constructing routing $R_S \in \mathcal{R}(I)$ such that $\pi(R_S) \leq 3\pi(R_{\text{OPT}}) + 1$. First, obtain an edge-coloring of multigraph $H(I)$ using at most $\frac{3}{2}\Delta_{H(I)}$ colors in $O(|I|^2)$ time [7]. Let I_i , for $1 \leq i \leq \lceil \frac{3}{2} \frac{\Delta_{H(I)}}{|V|} \rceil$, denote the subset of instance I consisting of those requests which were assigned a color from the range $[(i-1)|V| + 1, i|V|]$

in the edge-coloring of $H(I)$; obviously, $I = \bigcup_{i=1}^{\lceil \frac{3}{2} \frac{\Delta_{H(I)}}{|V|} \rceil} I_i$. Since by definition of I_i multigraph $H(I_i) = (V, I_i)$ has a known edge coloring using at most $|V|$ colors, by Lemma 5 there exists a routing $R_i \in \mathcal{R}(I_i)$ such that $\pi(R_i) \leq 2$.

The sought suboptimal solution $R_S \in \mathcal{R}(I)$ is defined as $R_S = \bigcup_{i=1}^{\lceil \frac{3}{2} \frac{\Delta_{H(I)}}{|V|} \rceil} R_i$. It is apparent that:

$$\pi(R_S) \leq \sum_{i=1}^{\lceil \frac{3}{2} \frac{\Delta_{H(I)}}{|V|} \rceil} \pi(R_i) \leq 2 \lceil \frac{3}{2} \frac{\Delta_{H(I)}}{|V|} \rceil \leq 3 \lceil \frac{\Delta_{H(I)}}{|V|} \rceil + 1 \leq 3\pi(R_{\text{OPT}}) + 1$$

which completes the proof of the approximation ratio of the algorithm. \square

In the above example, the value of the approximation ratio results from the inequality $\pi(R_S) \leq M_2 \lceil M_1 \pi(R_{\text{OPT}}) \rceil$, for values of approximation parameters $M_1 = 1.5$ and $M_2 = 2$, respectively. It is easy to see that applying the same approach in the on-line model does not affect the value $M_2 = 2$; however the simple online greedy edge-coloring algorithm for multigraph $H(I)$ may require $2\Delta_{H(I)}$ colors [7]; hence $M_1 = 2$. Consequently, we have the following corollary.

Corollary 22. *There exists an on-line 4-approximation algorithm for the edge congestion routing problem in complete graphs, requiring $O(|I|)$ time per request. The dilation of the returned solution is not greater than 2.*

The wavelength routing problem. The wavelength routing problem for complete graphs, motivated by applications in all-optical networks [12], is defined similarly to edge congestion routing, but using the wavelength count $w(R)$ as a minimization parameter. For a given instance I in graph K_V , we seek a routing $R_{\text{OPT}} \in \mathcal{R}(I)$, such that $w(R_{\text{OPT}})$ is the minimum possible [2,3]. Approximation ratios are expressed in terms of upper bounds on the ratio $\frac{w(R_S)}{w(R_{\text{OPT}})}$. The problem can be solved by an approach analogous to that described in the proof of Theorem 21, leading to an inequality of the form $w(R_S) \leq M_2 \lceil M_1 w(R_{\text{OPT}}) \rceil$. The value of M_1 remains unchanged and is equal to $M_1 = 1.5$ in the off-line case and $M_1 = 2$ in the on-line case. Moreover, since routings R_i which appear in the proof have a conflict graph with connected components which are cycles, paths, or isolated vertices, they may be colored in a greedy manner using 3 colors, and we have $M_2 = 3$. These observations may be written in the form of the following corollaries.

Corollary 23. *There exists an off-line $(4.5 + \frac{1.5}{\text{OPT}})$ -approximation algorithm with $O(|I|^2)$ runtime for the wavelength routing problem in complete graphs. The dilation of the returned solution is not greater than 2.*

Corollary 24. *There exists an on-line 6-approximation algorithm for the wavelength routing problem in complete graphs, requiring $O(|I|)$ time per request. The dilation of the returned solution is not greater than 2.*

Finally, let us remark on a general property of all 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 intermediate nodes 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.

References

- [1] M. Andrews, L. Zhang, Hardness of the undirected edge-disjoint paths problem, in: Proc. 37th ACM Symposium on Theory of Computing, 2005, pp. 276–283.
- [2] B. Beauquier, et al., Graph problems arising from wavelength routing in all-optical networks, in: Proc. 2nd IEEE Workshop on Optics and Computer Science, Geneva, Switzerland, 1997.
- [3] J. Białogrodzki, Path Coloring and Routing in Graphs, in: Graph Colorings, in: Contemporary Mathematics, vol. 352, AMS, Providence, USA, 2004, pp. 139–152.
- [4] P. Carmi, T. Erlebach, Y. Okamoto, Greedy edge-disjoint paths in complete graphs, in: Proc. 29th Workshop on Graph Theoretic Concepts in Computer Science, in: LNCS, vol. 2880, 2003, pp. 143–155.
- [5] T. Erlebach, D. Vukadinović, New results for path problems in generalized stars, complete graphs, and brick wall graphs, in: Proc. 13th International Symposium on Fundamentals of Computation Theory, in: LNCS, vol. 2138, 2001, pp. 483–494.
- [6] L.M. Favrholdt, M.N. Nielsen, On-line edge-coloring with a fixed number of colors, Algorithmica 35 (2003) 176–191.
- [7] S. Fiorini, R.J. Wilson, Edge-Colourings of Graphs, Pitman, London, UK, 1977.

- [8] H.N. Gabow, An efficient reduction technique for degree-constrained subgraph and bidirected network flow problems, in: Proc. 15th ACM Symposium on Theory of Computing, 1983, pp. 448–456.
- [9] V. Guruswami, et al., Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems, *Journal of Computer and System Sciences* 67 (2003) 473–496.
- [10] P. Kolman, C. Scheideler, Improved bounds for the unsplittable flow problem, in: Proc. 13th ACM-SIAM Symposium on Discrete Algorithms, 2002, pp. 184–193.
- [11] H.W. Lenstra, Integer Programming with a fixed number of variables, *Mathematics of Operations Research* 8 (1983) 538–548.
- [12] R. Raghavan, E. Upfal, Efficient routing in all-optical networks, in: Proc. 26th ACM Symposium on Theory of Computing, 1994, pp. 134–143.
- [13] C.E. Shannon, A theorem on coloring lines of a network, *Journal of Mathematical Physics* 28 (1949) 148–151.
- [14] A. Srinivasan, Improved approximations for edge-disjoint paths, unsplittable flow, and related routing problems, in: Proc. 38th IEEE Symposium on Foundations of Computer Science, 1997, pp. 416–425.