

Adapted and Constrained Dijkstra for Elastic Optical Networks

Ireneusz Szcześniak

AGH University of Science and Technology
Department of Telecommunications
al. Mickiewicza 30
30-059 Krakow
Poland

Bożena Woźna-Szcześniak

Jan Długosz University in Częstochowa
Faculty of Mathematics and Natural Sciences
Institute of Mathematics and Computer Science
al. Armii Krajowej 13/15
42-200 Częstochowa
Poland

Abstract—We present an optimal and efficient algorithm for finding a shortest path in an elastic optical network. The algorithm is an adaptation of the Dijkstra shortest path algorithm, where we take into account the spectrum continuity and contiguity constraints, and a limit on the path length. The adaptation redefines the node label in the Dijkstra algorithm, allows for revisiting nodes even at a higher cost for different slices, avoids loops, and prunes worse labels. The algorithm is generic and agnostic of a specific spectrum allocation policy, as it finds the largest set of available slices from which slices can be allocated in any way. We describe and motivate the algorithm design, and point to our freely-available implementation using the Boost Graph Library. We carried out 8100 simulation runs for large, random and realistic networks, and found that the probability of establishing a connection using the proposed algorithm can be even twice as large as the probability of establishing a connection using the edge-disjoint shortest paths, and the Yen K shortest paths.

Index Terms—elastic optical networks, Dijkstra, constrained routing, Yen algorithm, edge-disjoint shortest paths, simulation, Gabriel graph

I. INTRODUCTION

Elastic optical networks (EON), a.k.a., the flex-grid networks, are considered the successor of wavelength-division multiplexing (WDM) networks. In recent years, EONs have been intensely researched by both the industry and the academia.

In EONs, the optical spectrum (the erbium window) is divided into thin spectrum *slices* (of, e.g., 6.25 GHz width), as opposed to coarse fixed-grid channels (of, e.g., 25 GHz width) of wavelength-division multiplexing (WDM) networks. In EONs, *contiguous* slices are concatenated to form a *slot*. Slots are tailored for a specific demand, unlike WDM channels, thus making EONs more spectrum-efficient than WDM networks.

The routing of a single connection in an EON is the single most important operation of a network management system (NMS), and one of the many research problems of the EON design, planning, and operation. For the NMS based on the Generalized Multi-Protocol Label Switching (GMPLS), the path computation element (PCE) is responsible for solving the routing and spectrum assignment (RSA) problem for the given demand, i.e., finding the path and the slices.

EONs should deal with dynamic traffic, where connections frequently arrive and do not last long as opposed to the incumbent WDM traffic. Furthermore, given the growing optical networks, the ever-increasing need for bandwidth and connection agility, further increased by the requirements of the fifth generation (5G) wireless networks, the PCE is expected to solve the RSA problem fast and well.

The RSA problem can be constrained by, for example, the path length, i.e., we can require the solution to be shorter than the given limit. This is an acceptable and desired limitation, since we may need to limit the path length for a number of reasons: the connection takes too much of network resources, the latency is too large, or the quality of the optical signal is low. The constriction limits the search space, thus perhaps making the problem tractable, i.e., solvable in reasonable time. To the best of our knowledge, it has not been proven whether a constriction makes the RSA problem tractable or not.

Our novel contribution is the algorithm which quickly and optimally solves the constrained RSA problem for a single demand. The algorithm is an adaptation and constriction of the Dijkstra shortest path algorithm. We show its effectiveness in comparison to routing with the edge-disjoint shortest paths and the Yen K shortest paths. The high-quality, high-performance implementation of the algorithm using the Boost Graph Library (BGL) is available at [1] under the General Public License (GPL).

Dijkstra is a principal graph algorithm, amenable to various adaptations due to its simple and clever design. Dijkstra is efficient and optimal, and follows the label-setting paradigm, as opposed to the label-correcting paradigm [2]. At first look, our adaptation seems to divorce the label-setting paradigm in favor of the label-correcting paradigm, because we allow for revisiting nodes, which Dijkstra does not do, and which is a hallmark of the label-correcting algorithms. But this is not so, the proposed algorithm is still a label-setting algorithm.

The article is organized as follows. In Section II we briefly review related works, in Section III we define the research problem, in Section IV we describe the algorithm, and in Section V we report on the simulation results. Finally, Section VI concludes the article.

II. RELATED WORKS

The RSA problem is reported to be NP-complete, so along with linear programming formulations, there have been heuristic algorithms proposed for real-sized networks [3]. In [4], these algorithms are categorized into one-stage algorithms, which route and assign spectrum in one stage, and two-stage algorithms, which do it in two separate stages.

In [3], the authors propose a one-stage heuristic algorithm, which is a constrained Yen K shortest path algorithm. The algorithm prunes the path deviations incapable of supporting a demand. The algorithm resorts to the Dijkstra algorithm to compute shortest paths.

In [5], the authors report a two-stage algorithm for routing with Yen K shortest paths. First, the Yen algorithm computes K shortest paths also using the Dijkstra algorithm, and next they try to establish a connection along these paths. The authors also proposed a one-stage algorithm, termed a modified Dijkstra algorithm, which is a typical constriction of the Dijkstra algorithm, where a candidate path is rejected if it cannot support a demand due to the lack of slices.

Another two-stage algorithm is routing with edge-disjoint shortest paths. First, all edge-disjoint shortest paths are found with the Dijkstra algorithm, and next they try to establish a connection along these paths.

All these algorithms, unlike ours, fail to find a shortest path capable of supporting a demand, when there is a shorter path incapable of supporting a demand, because that shorter path decoys Dijkstra into a dead end.

III. PROBLEM STATEMENT

Given:

- directed multigraph $G = (V, E)$, where $V = \{v_i\}$ is a set of nodes, and $E = \{e_i\}$ is a set of edges,
- attribute c of edge e_i , i.e., $e_i.c$, which gives non-negative cost (length) of edge e_i ,
- attribute SSC of edge e_i , i.e., $e_i.SSC$, which gives the set of available slices, with Ω being the set of all slices,
- maximal path cost (length) m ,
- demand $d = (s, t, n)$, where s is the source node, t is the target node, and n is the number of contiguous slices required.

Sought:

- shortest path $p = (e_1, \dots, e_i, \dots, e_l)$ for demand d in graph G , where e_i is the i -th edge of path p ,
- largest set of slices Σ , which can support demand d .

The objective is to find the *largest* set of slices (SSC) along the shortest path (SP). An SSC can describe any slices available on an edge or along a path, contiguous or not. The sought SSC is the largest possible, which can support demand d with n slices, i.e., it can have any number of contiguous spectrum fragments, each with at least n slices.

This problem formulation allows the algorithm to be more generic and agnostic of the *spectrum allocation policy*, because once the largest SSC is found, any policy can be used to allocate slices, e.g., first or fittest.

IV. PROPOSED ALGORITHM

We adapted and constrained the shortest path Dijkstra algorithm to find an SP in EONs. The adaptation is novel, and the constriction is trivial. The adaptation keeps track of the SSCs along the found paths, while the constriction limits the length of an SP. The found paths are the SPs capable of supporting a given SSC, though they are usually not the SPs in the graph.

The Dijkstra algorithm is a *label-setting* algorithm in that once a node is visited, its label is set, and does not change, but the label is updated by the *edge relaxation*, if the given edge yields a better than known label. In label-setting algorithms, a label is associated with every node, and gives information on what cost and how to reach the given node from the source. In Dijkstra a label is the pair of cost and a preceding node.

Our label, however, is more elaborate, since it has to describe more elaborate data. We define a label as a *tuple of cost, a preceding edge, and an SSC*. For instance, label $(1, e_1, \{1, 2\})$ says that a node is reached with cost 1 along edge e_1 and with the SSC of $\{1, 2\}$. To allow for multigraphs, in the tuple we keep a preceding edge, not a node.

In Dijkstra, node labels converge to their optimum by edge relaxation, which updates a label when a better one is found. Dijkstra compares two labels: a candidate one, and a known one. The candidate label is better if it offers to reach the given node at a lower cost than the known one.

When trying to relax a candidate edge, we also compare two labels, but we take into account not only the costs, but also the SSCs of the labels. Label l_1 is better than or equal to label l_2 , denoted by $l_1 \leq l_2$, if $\text{cost}(l_1) \leq \text{cost}(l_2)$ and $SSC(l_1) \supseteq SSC(l_2)$.

In Dijkstra, a node has a single label, while we allow a node to have a set of labels, provided that no label is better than or equal to some other label, i.e., for any labels l_i and l_j of a given node, $l_i \leq l_j$ is false. Our edge relaxation takes care of that.

As to the constriction of the path length, during a node visit, Dijkstra traverses the out-edges of the node to find candidate labels, and we require a candidate label to be dropped if its cost exceeds the limit m .

We had to limit the path length, because we had to narrow the solution search space. Otherwise, the algorithm can, as in some cases we ran into, keep going through a very large search space for days, and not find a solution.

A. Adaptation

The adaptation takes into account the spectrum continuity and contiguity constraints. The following three observations shaped the adaptation.

1) *Revisit nodes*: In Dijkstra, a node is visited once for a single label. We, however, allow for revisiting nodes for multiple labels, because one of them yields, if possible, an SP capable of supporting a given demand.

We show an example in Fig. 1 to motivate node revisiting, where the edge label gives a length and an SSC available on an edge, e.g., $(1, \{1, 2\})$ says the edge is of length 1 and slices




Fig. 1: Example for node revisiting and looping.

1 and 2 are available. We are searching for an SP with two slices from node s to node t .

In the first step of Dijkstra, node s is visited, and node i is discovered along the two parallel edges e_1 and e_2 , but the label is not updated for the longer edge e_2 . In the second step, node i is visited, and its label is set. Now the final label for node i is known: $(1, e_1, \{1, 2\})$. The problem is that node t cannot be discovered, because the spectrum continuity constraint would be violated: the SSC of edge e_3 is $\{2, 3\}$, node i was reached with SSC $\{1, 2\}$, and the demand requires two slices.

Revisiting nodes solves this problem. We allow for revisiting a node even at a higher than known cost. In Dijkstra, in contrast, a node is visited only once at the lowest cost.

Continuing with the example, and allowing for node revisiting, now node i is discovered along both parallel edges e_1 and e_2 , and none of the discoveries is discarded. Then node i is visited along edge e_1 with label $(1, e_1, \{1, 2\})$, and then revisited along edge e_2 with label $(2, e_2, \{2, 3\})$, thus allowing node t to be discovered, end eventually visited with label $(3, e_3, \{2, 3\})$.

For simplicity, we illustrated node revisiting with parallel edges, but could have also used parallel paths. In the example, for instance, edge e_2 can be replaced with two edges and a node between them.

2) *Avoid loops:* Revisiting nodes may cause the search to find paths with loops. For instance, considering the same example in Fig. 1, when visiting node i , we discover node s and later revisit it, thus finding the loop (e_1, e_2) . In Dijkstra, loops are avoided by the edge relaxation, which accepts only labels of lower cost: since edge weights are non-negative, loops cannot decrease cost, and so they will not be allowed by edge relaxation. The thing is, that we need to allow for revisiting even at higher costs, but also need to avoid loops.

To avoid loops, and still to allow for node revisiting, an edge can be relaxed even at a cost higher than the cost of any node label, provided the candidate label offers an SSC not already included in SSCs of the node labels. Therefore, a node is visited and possibly revisited always at the lowest cost for an SSC not included in the SSCs of previous visits. And so, a node can have a set of labels, but no label is better than or equal to some other label, i.e., for any labels l_i and l_j of a given node, $l_i \leq l_j$ is false.

For example, in Fig. 1, the initial label for node s is $l_1 = (0, e_\emptyset, \Omega)$. The null edge e_\emptyset , which is not present in graph G , marks the beginning of an SP. When visiting node i with label $(1, e_1, \{1, 2\})$, node s is discovered along edge e_2 with label $l_2 = (3, e_2, \{2\})$, but the edge will not be relaxed, because




Fig. 2: Example for purging worse labels.

$l_1 \leq l_2$, thus avoiding a loop.

3) *Purge labels:* When we relax an edge, we add a new label for the node, but we may also need to purge worse labels. The purging of worse labels is illustrated by the example in Fig. 2. When visiting node s , node i is discovered along edge e_1 with label $l_1 = (1, e_1, \{1, 2\})$, and node i has label l_1 only. Next, node i is discovered along edge e_2 with label $l_2 = (1, e_2, \{1, 2, 3\})$. Label l_2 is better than l_1 , because the SSC of l_2 includes the SSC of l_1 , and both labels are of the same cost. We purge l_1 from the set of labels of node i , and now node i has label l_2 only.

B. Algorithm

Algorithm 1 presents the complete algorithm with the typical of the Dijkstra algorithm structure, where the main loop processes the priority queue elements. Our priority queue Q stores the elements $q = (c, e)$, which are pairs of cost c , and edge e . The queue is sorted according only to the increasing cost c of the elements, without the consideration of edge e , which is associated with the cost to tell what to process.

We require Q to store only unique elements: pushing the same element q many times results in just one element q in the queue. This property is required to process in one iteration of the main loop all the labels with the same cost and edge, but a different SSC. Q can be implemented as a set to guarantee this property.

We initialize $L_s = \{(0, e_\emptyset, \Omega)\}$ to make all slices available at node s at cost 0. The null edge e_\emptyset , which is not present in graph G , marks the beginning of an SP. L_v is a set of labels of node v , and $L = \{L_v\}$ is the set of sets of node labels. Next, we put element $(0, e_\emptyset)$ to the queue to boot the search.

In every iteration of the main loop, we pop from Q element $q = (c, e)$, and visit node $v = e.target$ reached along edge e at cost c . The target node of an edge is given by attribute *target*, i.e., $e.target$, with the special case of $e_\emptyset.target = s$. If $v = t$, then we found a solution and break the main loop.

In the main loop we iterate over two nested loops. One loop iterates over all SSCs S of labels in L_v with cost c and edge e , and the other loop iterates over all outgoing edges e' of node v , in order to discover a neighbor node $v' = e'.target$ along edge e' at cost $c' = c + e'.cost$ with SSC $S' = S \cap e'.SSC$. We continue working with S' and e' , if $c' \leq m$ and S' can support demand d , i.e., S' has at least n contiguous slices.

Next, we check whether edge e' can be relaxed with candidate label l' , i.e., whether node v' has no label l better than or equal to label l' . If so, then 1) we purge every label l of node v' if $l' \leq l$, 2) add label l' to the set of labels of node v' , 3) push element (c', e') to Q . Edge relaxation replenishes

Algorithm 1In: $G = (V = \{v_i\}, E = \{e_i\}), W(e_i), S(e_i), m, d = (s, t, n)$ Out: $p = (e_1, \dots, e_i, \dots, e_l), \Sigma = \{\sigma_i\}$

```

 $L_s = \{(0, e_\emptyset, \Omega)\}$ 
push  $(0, e_\emptyset)$  to  $Q$ 
while  $Q$  is not empty do
   $q = (c, e) = \text{pop}(Q)$ 
   $v = e.target$ 
  if  $v == t$  then
    break the while loop
  end if
   $SSSC = \{l.SSC : l \in L_v \text{ and } l.c == c \text{ and } l.e == e\}$ 
  for all  $S \in SSSC$  do
    for all  $e' \in$  outgoing edges of  $v$  do
       $S' = S \cap S(e')$ 
       $c' = c + W(e')$ 
      if  $c' \leq m$  and  $S'$  can support  $d$  then
         $v' = e'.target$ 
         $l' = (c', e', S')$ 
        if  $\nexists l \in L_{v'} : l \leq l'$  then
           $L_{v'} = L_{v'} \setminus \{l : l \in L_{v'} \text{ and } l' \leq l\}$ 
           $L_{v'} = L_{v'} \cup \{l'\}$ 
          push  $(c', e')$  to  $Q$ 
        end if
      end if
    end for
  end for
end while
return  $(p, \Sigma) = \text{trace}(L, s, t)$ 

```

the queue, and the algorithm keeps iterating until destination node t is reached, or the queue is empty.

Finally, function *trace* traces back an SP found, if any, based on the node labels L . We do not present the algorithm for tracing back an SP, since it is rather easy.

V. SIMULATIVE STUDIES

We evaluate the performance of the proposed algorithm with simulations, and compare it to the performance of routing with the edge-disjoint paths and the Yen K shortest paths. We also show, based on the simulation results, that the proposed algorithm efficiently solves the constrained RSA problem, which suggests that this problem is tractable, though we offer no proof.

For comparison, we use the edge-disjoint paths and the Yen K shortest paths, because they are very different: the edge-disjoint shortest paths do not share even a single edge, while the Yen K shortest paths can differ with a single edge only. To find edge-disjoint paths, we search for an SP in a graph with the edges of the previous SPs disabled. The Yen K shortest paths are found with the well-known Yen algorithm. We do not limit the number of edge-disjoint paths, because at most it equals to the degree of the source or target nodes, which is a small number. We limit, however, the Yen shortest paths to at

most $K = 10$, because Yen can produce a very large number of paths.

Having either the edge-disjoint shortest paths or the Yen K shortest paths, we try to route a demand as follows. We start with the first SP, and calculate the largest available SSC along the path by intersecting all the SSCs available on every edge of the shortest path. If the largest SSC cannot support the demand, we try the next shortest path, until there are no more shortest paths.

Having found a path with the largest SSC, a spectrum allocation policy allocates n slices for demand d from the largest SSC. We consider only the *fittest* and the *first* spectrum allocation policies. The fittest policy allocates n slices in the fittest fragment of the largest SSC, which can support demand d , i.e., the smallest fragment with at least n slices. The first policy allocates n slices in the first fragment (with slices of the smallest numbers) of the largest SSC, which can support demand d .

To be general, we formulate and solve the problem for a directed multigraph, but we use it for an undirected graph, which models an EON.

A. Simulation setting

We generate a set of random graphs with random traffic to obtain reliable statistical results for various populations of interest, because we find studying a specific topology (e.g., Polish PIONIER or NSFNet) with some specific traffic case rather inconclusive. We use Gabriel graphs, because they have been shown to model the properties (e.g., the node degree) of the transport networks very well [6].

We randomly generate 50 Gabriel graphs, where each edge has 400 slices. Each graph has 100 nodes, which are uniformly distributed over an area 1000 km long and 1000 km wide. In generating Gabriel graphs, the number of edges cannot be directly controlled, as it depends on the location of nodes, and on the candidate edges meeting the conditions of the Gabriel graph. The statistics of the generated graphs are given in Table I. The limit on the path length is $m = 2000$ km, which is well above 1582 km, the length of the longest of all shortest paths in the generated graphs.

Demands arrive according to the exponential distribution with the rate of λ demands per day. The probability distribution of the demand holding time is also exponential with the mean of $\beta = 10$ days. The number of slices a demand requires follows the Poisson distribution with the mean of 10 slices.

We argue that the choice of a traffic model is irrelevant to our study as the traffic only produces the input data (i.e., the state of the graph) for the routing algorithms, and we chose the exponential and Poisson distributions to keep the discussion simple. The question is how the algorithms perform under the given utilization and fragmentation, regardless of how the utilization and fragmentation were obtained, which could have been equally well produced randomly. For the very same reason we do not incorporate into our study spectrum defragmentation.

value	min	average	max	variance
Number of links	160	179.2	194	48.52
Link length	1	97.95	347	2696.46
Node degree	1	3.584	8	1.2213
SP length	1	589.61	1582	78118.1
SP hops min	1	6.7634	22	10.8737

TABLE I: Statistics of the generated Gabriel networks.

We define *network utilization* as the ratio of the number of the slices in use to the total number of slices on all edges. We cannot directly control the network utilization, but measure it in response to the offered load. And so to obtain different values of network utilization, we varied λ with 27 different values of 10, 12.5, 15, 17.5, 20, 25, 30, 35, 40, 45, 50, 55, 60, 70, 80, 90, 100, 150, 200, 300, 400, 500, 600, 700, 800, 900, and 1000 demand arrivals per day.

A *simulation run* simulates 100 days of a network in operation. For every simulated day, the following values are measured: (a) the instantaneous network utilization, (b) the average probability of establishing a connection during that day, (c) the instantaneous number of active connections, (d) the instantaneous amount of capacity served, (e) the average length of an established connection during that day, (f) the average number of slices of an established connection during that day, (g) the instantaneous number of spectrum fragments of an edge, (h) the average time taken by an SP search during that day. When a simulation finishes, the measured values are averaged and reported as the simulation results.

We want to get reliable results for a statistical population of simulation runs. A population is described by the routing algorithm used, the spectrum allocation policy used, and λ , and so there are 162 populations considered (three routing algorithms \times two spectrum allocation policies \times 27 values of λ). In a given population, all simulation runs have the same parameters, except the seed of a random number generator in order to generate different Gabriel graphs and different traffic. To get reliable results for a population, we carry out 50 simulation runs which are the population samples, and calculate the sample means of all the results reported by a simulation, except the average search time of which we take the minimum. In total there are 8100 simulation runs (162 populations \times 50 samples). We reckon the sample means reliably estimate the results of their populations, since their relative standard error is below 1%.

We do not calculate the sample mean of the average search time, but take the minimum, since we run simulations using a supercomputing infrastructure, and cannot control the specific hardware for our simulations, and how much the hardware is loaded with other jobs. Other processes running can heavily utilize memory, thus causing cache misses in our simulations, which severely degrade performance. The supercomputing infrastructure is composed of thousands of nodes equipped with the state-of-the-art multi-core processors of the AMD64 architecture.

B. Simulation results

Figures 3a-3d show the simulation results for routing with the proposed algorithm (solid curves), routing with the edge-disjoint paths (dashed curves), and routing with the Yen K shortest paths (dotted curves). Routing was carried out for two spectrum allocation policies: fittest (thick curves) and first (thin curves). Each figure has six curves for three routing types with two spectrum allocation policies. Each curve is plotted with 27 data points for different values of λ . Each data point represents a sample mean, except the time of the shortest path search, which represents a sample minimum. Since the relative standard errors of the sample means are below 1%, the error bars would be too small to plot.

Fig. 3a shows the probability of establishing a connection as a function of network utilization. We are interested in how a network performs for a given network state expressed by utilization. The proposed algorithm considerably outperforms the other routing types for all network loads. For the utilization of 30%, the proposed algorithm still has the probability of nearly 1, while for the other two types, their probabilities drop to about 0.75. For the load of 40%, the probability for the proposed algorithm is almost twice as large as that for the other two routing types.

Fig. 3b shows the time taken by a shortest path search as a function of network utilization, regardless of whether the search was successful or not. The algorithms for finding edge-disjoint paths and Yen K shortest paths do not take into account the slices available on edges, they do not depend on the utilization, and so their times are constant. Interestingly, the time taken by the proposed algorithm decreases as the utilization increases, because the search space gets narrower.

Fig. 3c shows the average length of an established connection as a function of network utilization. As the network is utilized more, the average length of a connection increases, because the proposed algorithm finds more circuitous paths, but still finds them, while the other routing types fail. The average length for all routing types drops as utilization keeps increasing, since demands with end nodes close to each other are more likely to be established.

Fig. 3d shows the average number of slices of an established connection as a function of network utilization. As expected, the average number decreases as utilization increases, because demands which require a smaller number of slices, are more likely to be established. At first glance, the proposed routing performs better for network utilization below 45%, and worse otherwise than the other two routing types. However, this is the average number of slices, provided a connection was established. As the other two routing algorithms offer lower probabilities of establishing a connection, they are more likely to establish short connections, which are more likely to succeed demanding a larger number of slices. The proposed algorithm is more likely to establish connections between pairs of distant nodes with a smaller number of slices, thus lowering the average number of slices.

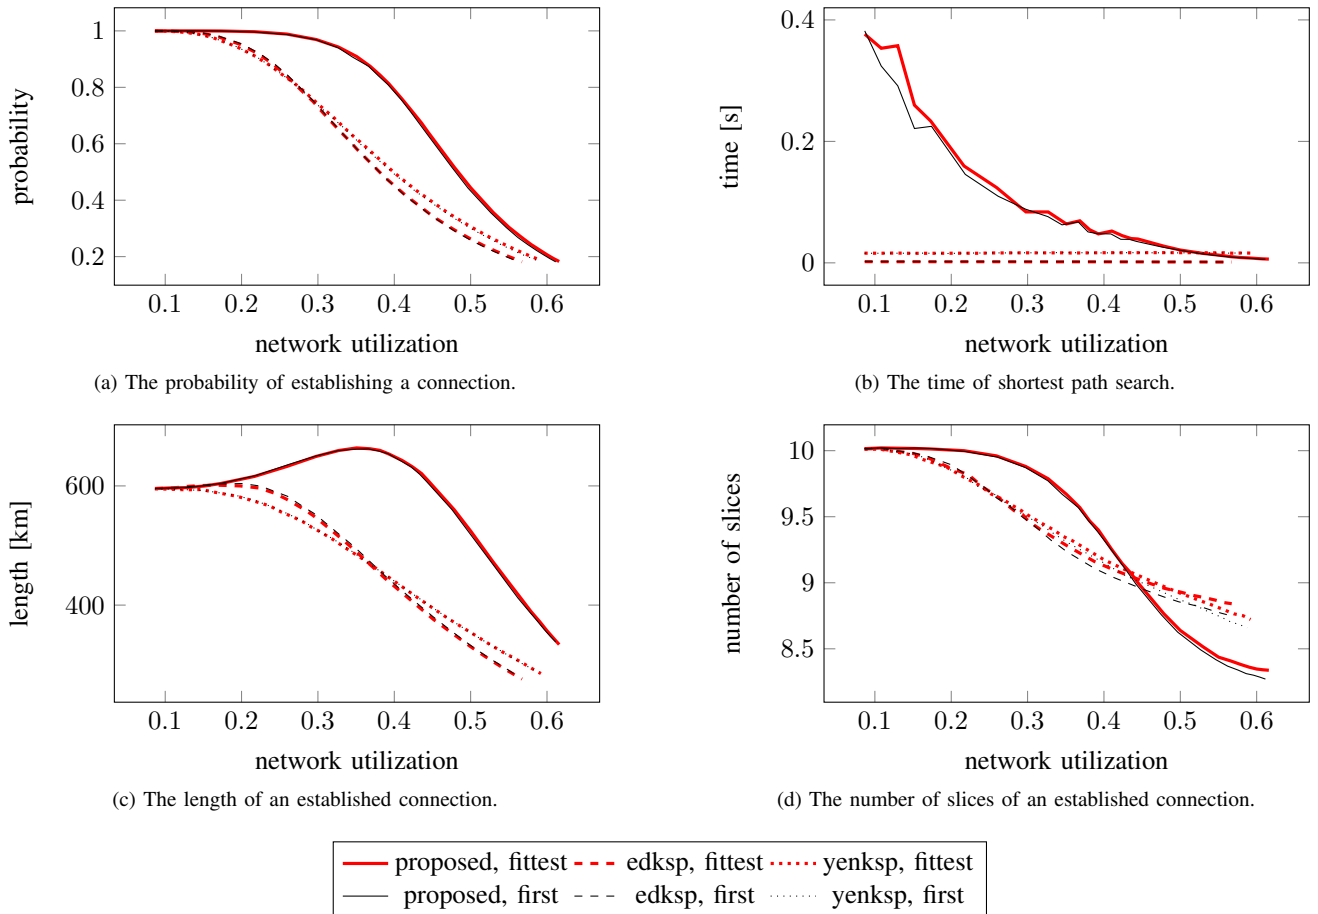


Fig. 3: Simulation results for Gabriel graphs.

VI. CONCLUSION

We proposed the adaptation and the constriction of the Dijkstra shortest path algorithm for finding shortest paths in elastic optical networks. The adaptation is a novel contribution, which takes into account the spectrum continuity and contiguity constraints.

The routing and spectrum assignment problem is known to be NP-complete, yet our algorithm has no difficulty finding a solution, since, we speculate, limiting the path length narrows the search space. Limiting the path length is not a problem, since it is needed in practice anyway.

Our extensive simulation studies show that the proposed algorithm outperforms two other routing types frequently used in research on elastic optical networks: routing along the edge-disjoint paths, and routing along the Yen K shortest paths. The studies also show that the algorithm can be used for routing in elastic optical networks of large sizes, thus making the algorithm practical.

Future work could concentrate on removing the path length constriction, and finding the stop condition for the algorithm, i.e., a condition to stop searching for a path, when it is known that no path can be found.

ACKNOWLEDGMENTS

This work was supported by the postdoctoral fellowship number DEC-2013/08/S/ST7/00576 from the Polish National Science Centre. The numerical results were obtained using PL-Grid, the Polish supercomputing infrastructure.

REFERENCES

- [1] I. Szcześniak. (2016, January) The implementation of the adapted and constrained Dijkstra for elastic optical networks. [Online]. Available: <http://www.irkos.org/abcd>
- [2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network flows: theory, algorithms, and applications*. Prentice Hall, 1993.
- [3] L. Velasco, A. Castro, M. Ruiz, and G. Junyent, "Solving routing and spectrum allocation related optimization problems: from off-line to in-operation flexgrid network planning," *Journal of Lightwave Technology*, vol. 32, no. 16, pp. 2780–2795, Aug 2014.
- [4] G. Zhang, M. De Leenheer, A. Morea, and B. Mukherjee, "A survey on OFDM-based elastic core optical networking," *IEEE Communications Surveys Tutorials*, vol. 15, no. 1, pp. 65–87, 2012.
- [5] X. Wan, L. Wang, N. Hua, H. Zhang, and X. Zheng, "Dynamic routing and spectrum assignment in flexible optical path networks," in *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2011*. Optical Society of America, 2011, p. JWA055.
- [6] E. Cetinkaya, M. Alenazi, Y. Cheng, A. Peck, and J. Sterbenz, "On the fitness of geographic graph generators for modelling physical level topologies," in *2013 5th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, Sept 2013, pp. 38–45.