

Routing Metrics Depending on Previous Edges: The M_n Taxonomy and its Corresponding Solutions

Amaury Van Bemten, Jochen W. Guck, Carmen Mas Machuca and Wolfgang Kellerer

Lehrstuhl für Kommunikationsnetze

Technical University of Munich, Germany

Email: {amaury.van-bemten, guck, cmas, wolfgang.kellerer}@tum.de

Abstract—The routing algorithms used by current operators aim at coping with the demanded QoS requirements while optimizing the use of their network resources. These algorithms rely on the *optimal substructure property* (OSP), which states that an optimal path contains other optimal paths within it. However, we show that QoS metrics such as queuing delay and buffer consumption do not satisfy this property, which implies that the used algorithms lose their optimality and/or completeness. This negatively impacts the operator economy by causing a waste of network resources and/or violating Service Level Agreements (SLAs). In this paper, we propose a new so-called *M_n taxonomy* defining new metric classes. An M_n metric corresponds to a metric which requires the knowledge of the n previously traversed edges to compute its value at a given edge. Based on this taxonomy, we present three solutions for solving routing problems with the newly defined classes of metrics. We show that state-of-the-art algorithms based on the OSP indeed lose their original optimality and/or completeness properties while our proposed solutions do not, at the price of an increased computation time.

Keywords—routing algorithm, routing metric, Dijkstra, optimal substructure property, taxonomy, graph transformation

I. INTRODUCTION

The new era of telecommunications considers new aspects such as virtual operators, large number of terminals and users and wide range of services with heterogeneous Quality of Service (QoS) requirements. These aspects force operators to deploy flexible networks able to guarantee all these requirements with an efficient use of their network resources. This paper shows that, with a strict modeling of classical QoS parameters such as delay and buffer consumption, link properties are not static and state-of-the-art algorithms lose their optimality and/or completeness properties¹. As a result, to prevent Service Level Agreements (SLAs) violations and ensure optimal use of their network resources, operators need new solutions.

One important challenge to ensure efficient resources usage is the optimal routing of data packets. Depending on the type of route that has to be found, different types of routing problems have been defined [1], [2]. For example, for a given graph, finding a path between two nodes and finding a tree from a node to a set of other nodes are two different routing problems. Among the set of possible solutions, it is usually also desirable to define a way for preferentially selecting one solution over another and for only selecting solutions satisfying a given set of requirements. For example, among the set of paths between two nodes, one might want to find the path with the greatest available bandwidth or simply to find any path

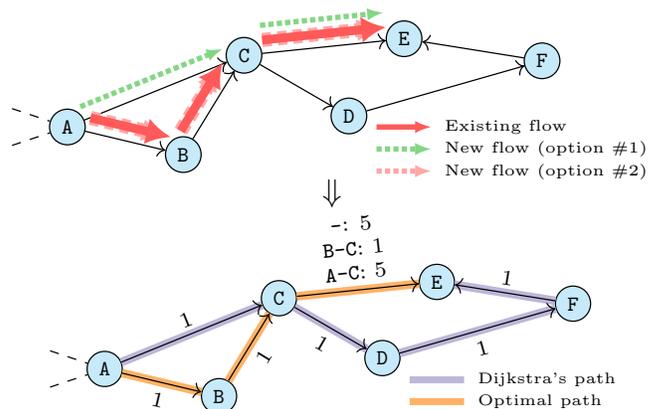


Fig. 1: Example scenario involving an M_1 metric. As elaborated in Sec. I-A, we observe that, in such a situation, Dijkstra finds a sub-optimal path.

with an available bandwidth greater than a given threshold. With this aim, so-called *metric values* are associated to each edge of the graph. For example, if propagation delay is defined as a metric, one might want to find the path between two nodes whose total propagation delay is minimal. Algorithms for solving routing problems using metrics both as constraints and/or as optimization objectives have been deeply studied [3].

A. Motivation: Violation of the Optimal Substructure Property

State-of-the-art routing algorithms rely on the fact that the metric values associated to the edges of the graph are *static* (i.e., constant) for a given routing request [1], [2]. However, this is not always the case.

Let us examine the example depicted in Fig. 1, which considers as metric the queuing delay experienced at the ingress of each link (referred to as delay in this section). Let us assume that the network carries a flow on path A-B-C-E (continuous red arrows in Fig. 1). Because there is only one flow in the network, no queuing occurs at any link and all the links have the same delay metric value (e.g., 1 in this example). A new flow has to be routed from A to E. In order to reach C, the flow can either be routed through A-C or A-B-C. If it is routed through A-C, the flow may generate queuing at the egress of C, i.e., at link C-E, because of potential collisions of packets of the two flows coming from the two different ingress links of C. This implies that the delay metric value of C-E should be greater than without collisions (e.g., 5 in this example). On the other hand, if the new flow is routed through A-B-C, it will follow the same path as the original flow, thereby generating no queuing delay at any link². Hence, all links keep the same original delay metric value (1 in our

¹These properties are defined in Sec. II-4.

²We assume that the flows arrived in A via the same link and with the same total delay, thereby generating no queuing delay at the ingress of A-B.

case). We observe that the delay metric value associated to C-E depends on the previous edge traversed by the new flow. This is shown in the lower diagram of Fig. 1.

In such a setting, the least delay path from A to E is A-B-C-E, with a delay of 3. However, the Dijkstra algorithm [4], an optimal algorithm for the least-delay (i.e., shortest path) problem, finds the path A-C-D-F-E, with a delay of 4. Indeed, Dijkstra performs a breadth-first search and keeps only track of the best path to reach each node. While finding a path from A to E, Dijkstra will store A-C as the best path to reach C because it has a total delay of 1, which is lower than 2, the delay of A-B-C, the only other path to reach C. Then, Dijkstra will have two possibilities to reach the destination. Either using C-E with a delay of 5, or following C-D-F-E with a delay of 3. Since C-D-F-E has a lower delay, Dijkstra will choose it and return A-C-D-F-E, with a total delay of 4, as final solution. We observe that, in this situation, Dijkstra is not able to find the optimal path. We will show in Sec. III-B3 that this is due to the fact that the *optimal substructure property* (OSP), stating that an optimal path contains other optimal paths within it, is not satisfied.

Sec. III-C2 and Sec. III-C3 present other scenarios for which state-of-the-art algorithms based on the OSP (i.e., nearly all state-of-the-art routing algorithms) lose either their optimality or their completeness property, or both. Currently, operators solve this problem by allowing sub-optimality or by having a looser modeling of QoS parameters, thereby wasting network resources, or, in the worst case, by having a too optimistic modeling of QoS parameters, thereby potentially leading to SLAs violations. As a result, new optimal and complete algorithms for dealing with this type of metrics are needed.

B. Contributions: Mn Taxonomy and Corresponding Solutions

The contribution of this paper is twofold.

First, in Sec. III, we propose the *Mn taxonomy*, a new taxonomy defining new classes of routing metrics depending on the amount n of previous edges needed to compute their value at a given edge. We will see that this taxonomy allows to determine whether or not state-of-the-art algorithms based on the OSP lose their optimality and/or completeness properties (see Tab. I).

Second, for the new classes of metrics introduced by the taxonomy, new optimal and complete algorithms are needed. Hence, in Sec. IV, we present three solutions: *A*Prune*, *edge-based Dijkstra* (EBD) and a *graph transformation algorithm* (GTA). On the one hand, A*Prune [5], a state-of-the-art algorithm, and EBD, a newly proposed extension of the Dijkstra algorithm [4], are algorithms for the specific, respectively, *multi-constrained shortest path* (MCSP) and *shortest path* (SP) problems, which keep their optimality and completeness properties for the new defined classes of metrics. On the other hand, GTA can be used for any routing problem. Indeed, GTA is an extension that can be applied to any state-of-the-art algorithm for recovering its optimality and completeness properties.

In Sec. V, through evaluations, we show that state-of-the-art algorithms based on the OSP indeed lose their optimality

and/or completeness properties depending on the Mn classification of the considered metric. Further, we show that our proposed solutions are indeed complete and optimal for their respective problems. However, this comes at the price of an increased runtime.

II. ROUTING METRICS: DEFINITION

Depending on the type of route that has to be found (e.g., a simple path from a single source to a single destination, multiple paths from a single source to a single destination, or a tree from a single source to several destinations), there exist a wide range of different routing problems. There is usually more than one possible solution to a given routing problem. In order to prefer one solution over another, or to provide additional requirements regarding the solutions to be accepted, so-called *routing metrics* are introduced. Each metric defines a value, referred to as a *metric value*, for each edge of the subject graph. These values can be updated for each routing request. The metrics can then be used in three different ways for selecting one or several of the available solutions.

1) *Local Constraint Metrics*: First, the edges that can be used by the solutions can be restricted to those satisfying a given condition based on a metric value. We refer to such metrics as *local constraint metrics*. For example, in order to ensure enough bandwidth is available for a given video stream, one might want to use only links whose bandwidth is greater than the bit rate of the video stream.

2) *Global Constraint Metrics*: Secondly, in order to further restrict the set of solutions that can be returned, the values of a metric at each edge of a solution can be combined using a so-called *link combination operator* [6], e.g., the sum or the multiplication, whose result must satisfy a given constraint. We refer to such metrics as *global constraint metrics*. For example, in order to ensure that the packets of a critical unicast flow arrive on time, one might want to find a path for which the sum of the delays of each of its constituting links is lower than a given threshold.

3) *Global Optimization Metrics*: Finally, in order to rank solutions and search for the preferred one(s), the values of a metric at each edge of a solution can be combined using a given link combination operator whose result is used for ordering the solutions. We refer to such metrics as *global optimization metrics*. For example, in order to ensure that data from a unicast request is transferred as fast as possible, one might want to find the path for which the sum of the delays of each of its constituting links is minimal.

4) *Optimality and Completeness Properties*: An algorithm is said *complete* if it always finds a solution, if one exists, satisfying both the local and global constraints. In case of a single global optimization metric, an algorithm is said *optimal* if the solution it finds is always the optimal one.

III. THE Mn TAXONOMY

In this section, we present our novel routing metric taxonomy, the *Mn taxonomy*, classifying metrics into classes based on the amount of previous edges needed for computing their value at a given edge. We refer to a metric requiring the knowledge of the n previously traversed edges as an Mn metric.

l	M0	M1	... Mn ...	M_∞
Local constraint (Sec. II-1)	C O	C O	C O	C O
Global constraint (Sec. II-2)	C O	C O	C O	C O
Global optimization (Sec. II-3)	C O	C O	C O	C O

TABLE I: Impact of M_n metrics on the completeness (C) and optimality (O) of state-of-the-art algorithms based on the OSP.

A. M0 Metrics: No Additional Information Required

M0 metrics correspond to the traditional metrics considered in the state-of-the-art. The metric value associated to an edge depends only on the edge itself and requires no information on the other edges previously traversed. Examples of M0 metrics are the propagation delay and the total capacity of a link.

B. M1 Metrics: Values Depending on the Previous Edge

1) *Definition:* M1 metrics correspond to metrics whose value at a given edge depends on the previous edge used to reach the given edge.

2) *Motivation: Queuing Delay:* The example developed in Sec. I-A (and illustrated in Fig. 1) corresponds to an M1 metric. Indeed, depending on which ingress link to C is used, the queuing delay at the egress of C is different. As a result, the metric value of C-E depends on the previously traversed edge and the metric is an M1 metric.

3) *Impact as Global Optimization Metric:* The example developed in Sec. I-A (and illustrated in Fig. 1) corresponds to an M1 metric used as global optimization metric for a *shortest path* (SP) problem. We have seen that the Dijkstra algorithm loses its optimality. The reason for this is that Dijkstra relies on the *optimal substructure property* (OSP), which states that sub-paths of optimal paths are also optimal [7]. While the OSP is satisfied for M0 metrics, it is not necessarily satisfied anymore for M1 metrics. Other routing algorithms typically also rely on the OSP, either because they are based on Dijkstra itself, or because they are based on dynamic programming or greedy approaches which are themselves based on the OSP [7]. Consequently, other optimal SP algorithms such as Bellman-Ford [8]–[10] and A* [11] are also affected. Hence, when an M1 metric is used as optimization metric, state-of-the-art algorithms based on the OSP lose their optimality property. Note that, as the metric is only used for optimization, completeness is not impacted. This is summarized in Tab. I and will be confirmed in our evaluations (Sec. V-C).

4) *Impact as Global Constraint Metric:* Let us consider the M1 metric in Fig. 1 as global constraint metric (with a bound of 3.5) and further define the hop count (an M0 metric) as global optimization metric. This corresponds to a *constrained shortest path* (CSP) problem. CBF [12], an optimal CSP algorithm, is similar to Dijkstra. It performs a breadth-first search and keeps only track of the best path at each node. However, it discovers paths in order of the constraint metric and stops once the bound is reached. In our example, CBF would hence also find A-C as the best path to reach C. From this path, the destination E cannot be reached within the deadline. Hence, CBF will conclude that no path is available. However, A-B-C-E, with a delay of 3, is a valid solution. As a result, CBF is incomplete. It can easily be shown that, with a bound of 4.5, CBF would find A-C-D-F-E, with a delay of 4, which is sub-optimal (A-B-C-E, with a delay of 3, still being the optimal path).

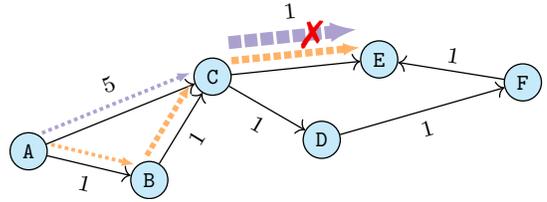


Fig. 2: Example scenario involving an M_∞ metric. As elaborated in Sec. III-C4, we observe that the new flow will be rejected or accepted at C-E depending on where it is coming from.

As for Dijkstra, this is due to the fact that the OSP is not satisfied anymore. Hence, other optimal CSP algorithms [13], which are all based on Dijkstra and hence on the OSP, are also affected. Consequently, when an M1 metric is used as a global constraint metric, state-of-the-art algorithms based on the OSP lose both their completeness and optimality properties. This is shown in Tab. I and will be confirmed in our evaluations (Sec. V-D).

5) *Impact as Local Constraint Metric:* It can also be easily shown that M1 local constraint metrics also lead to the sub-optimality and incompleteness of state-of-the-art algorithms based on the OSP (see Tab. I). Indeed, the OSP is also not necessarily satisfied.

C. M_∞ : Values Depending on the Complete Path

1) *Definition:* M_∞ metrics correspond to metrics whose value at a given edge depends on the complete path traversed to reach the current edge.

2) *Motivation: Buffer Management:* Let us consider a metric representing the buffer consumption of a flow. While traversing a given path, the burstiness of a flow is increased at each hop by an amount depending on the flow and on the hop characteristics [14], [15]. The buffer consumption of a flow at a given node depends, among other things, on the burstiness of this flow. Hence, the buffer consumption of a flow at a given hop depends on all the previously traversed links. Consequently, such a metric, which should be used for per-node strict buffer management [16], is an M_∞ metric.

3) *Motivation: Routing through Service Function Chains:* Let us consider a metric representing the total bandwidth consumption of a flow at each link. While routing through *service function chains* (SFCs), loops can be introduced. Hence, the total bandwidth consumption of a flow at a given link depends on how many times the flow already visited the given link. As a result, all the previously visited links have to be known and this corresponds to an M_∞ metric.

4) *Impact as Local Constraint Metric:* Let us consider the example developed in Sec. III-C2 (illustrated in Fig. 2) and the problem of routing a flow from A to E (Fig. 2) through the least-hop path (M0 metric). We consider a local constraint metric rejecting flows consuming too much buffer space. In Fig. 2, each link is labeled with its queuing delay (generated by already embedded flows). Note that this queuing delay is not used as a metric. As per network calculus [14], [15], the burstiness of the flow, i.e., its buffer consumption, will increase approximately proportionally to the delay it experiences along its path. Fig. 2 shows the two options for routing the new flow. The width of the arrows represent the buffer consumption of the flow. While routing, Dijkstra will save the path A-C, with

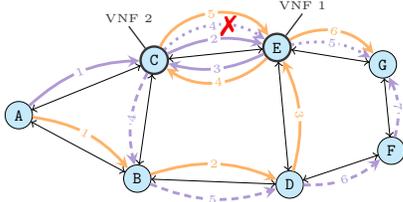


Fig. 3: Example scenario involving an M_∞ metric. As elaborated in Sec. III-C4, the bandwidth consumption of the new flow at a link depends on whether or not it visited this link before. Hence, acceptance of the new flow at a link depends on all the previously visited links.

a hop count of 1, as the best path towards C. However, because the queuing delay experienced at A-C is 5, the burstiness of the flow greatly increased and it cannot be accepted at link C-E. As a result, Dijkstra would either find no solution (thereby losing completeness) or find A-C-D-F-E, with a hop count of 4, if the latter has enough buffer space available. However, A-B-C-E, with a hop count of 3, has a low delay, thereby only slightly increasing the burstiness of the flow and hence allowing it to use C-E. In this case, Dijkstra is hence sub-optimal. This is again due to the fact that the OSP is not satisfied.

Let us further consider the example developed in Sec. III-C3 (illustrated in Fig. 3) and the problem of routing a flow on the least-hop path (M0 metric) from A to G (Fig. 3) visiting the two *virtual network functions* (VNFs) E and F in the specified order. The flow consumes 700 Mbps and each link has 1 Gbps available bandwidth. Hence, the flow can only visit each link once. A least-hop algorithm would first reach the first VNF, i.e., E, through the least-hop path, i.e., A-C-E. Then, it would visit the second VNF, i.e., C, through the least-hop path from E, i.e., E-C. Finally, it would try to reach the destination with the least-hop path from C, i.e., C-E-G. However, this would imply visiting a second time C-E, which is not allowed. If the algorithm does not notice that the access to C-E is refused (e.g., if it cached the result of the access control when it first traversed the link), it will return an invalid path and hence be incomplete. If the algorithm notices that the access to C-E is refused, it would follow the second least-hop path, i.e., C-B-D-F-E, thereby finding A-C-E-C-B-D-F-E, with a hop count of 7, as final solution. However, A-B-D-E-C-E-G is a valid path with a hop count of 6. Hence, the algorithm loses its optimality.

As a result, when an M_∞ metric is used as a local constraint metric, state-of-the-art algorithms based on the OSP lose both their completeness and optimality properties (Tab. I).

5) Impact as Global Optimization or Constraint Metric:

As an M1 metric is also an M_∞ metric, it can easily be shown that the impact as global optimization and constraint metric is the same for both M1 and M_∞ metrics (Tab. I).

D. M_n : Values Depending on the n Previous Edges

As a generalization, we further introduce the class of M_n metrics. M_n metrics correspond to metrics whose value at a given edge depend on the n previous edges used to reach the current edge.

Obviously, for the same reasons as for M_∞ metrics, the impact of state-of-the-art algorithms based on the OSP is the same as for M1 metrics. This is shown in Tab. I.

IV. SOLUTIONS FOR THE M_n TAXONOMY

In this section, we present three opportunities (Sec. IV-A, IV-B and IV-C) to optimally and completely solve problems with M_n metrics, $n > 0$.

A. Existing Solution: A*Prune

A*Prune [5] is a complete and optimal state-of-the-art algorithm able to solve the *shortest path* (SP) and *multi-constrained shortest path* (MCSP) problems. Although similar to Dijkstra, A*Prune does not rely on the OSP but is only faster when it is satisfied. Hence, it keeps its optimality and completeness properties for both M_n and M_∞ metrics. The reason for this is that it does not keep track of only one best path to reach each node. Instead, all feasible paths are kept in memory and the best ones are extended first. Path extension is stopped only once the next path to extend has an optimization metric value higher than the current best path for the destination.

In the example of Fig. 1, as Dijkstra, A*Prune will first find the path A-B-C-E with a total metric value of 7. However, the next path to extend, namely A-C, has a total metric value of 4, which is lower than the current best path to the destination. Hence, A*Prune will further extend A-C and thereby find A-C-E, which is optimal.

Unfortunately, this optimality for any type of M_n and M_∞ metric comes at the price of a poor scalability behavior. This will be shown in our evaluations (Sec. V).

B. New Solution: Edge-based Dijkstra (EBD)

In the particular case of the SP problem with an M1 optimization metric, Dijkstra can be slightly adapted. Instead of keeping track of the best path towards each node, our proposed adaptation keeps track of the best path towards each edge. We refer to this algorithm as *edge-based Dijkstra* (EBD).

In the example of Fig. 1, instead of keeping track of the best path towards node C, EBD will keep track of the best path towards A-C (which is A-C itself) and towards B-C (which is A-B-C). Then, when extending these paths to obtain the best path towards C-E, both paths will be considered and the path A-C-E, missed by the normal Dijkstra, will be found. Once EBD stops, the final solution then corresponds to the best path among those stored at all the ingress links to the destination node.

Because the amount of edges in a graph is usually higher than the amount of nodes, EBD keeps track and extends more paths than the traditional Dijkstra and the usage of EBD for M1 metrics hence results in a runtime increase compared to a normal Dijkstra run. This will be confirmed in Sec. V-C.

C. New Solution: Graph Transformation Algorithm (GTA)

In this section, we propose a *graph transformation algorithm* (GTA) which transforms a graph with M_n metrics to an equivalent graph with M0 metrics such that any routing problem with M_n metrics can be solved with a state-of-the-art algorithm for the given routing problem with traditional M0 metrics.

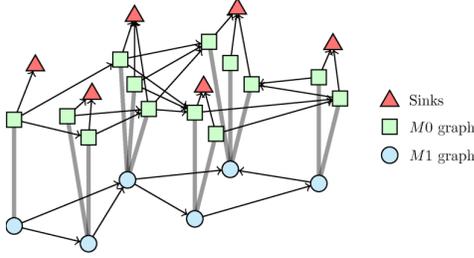


Fig. 4: Illustration of the GTA procedure for M1 metrics. From the original graph with M1 metrics (the M1 graph), the algorithm creates a new graph (the M0 graph) with M0 metrics on which any state-of-the-art algorithm can run to solve the original M1 problem.

1) *Reasoning*: For M1 metrics, edges can have as many different metric values as ingress links (plus one for the “null” ingress link when the flow starts at the given edge). This is shown in Fig. 1. The idea of GTA is to duplicate links as many times as they have different metric values such that each new edge has a static M0 metric value. Let us refer to the original graph with M_n metrics as the M_n graph and to the transformed graph with only M0 metrics as the M_0 graph. Each edge in the M0 graph then corresponds to (i) an edge of the original M_n graph, and (ii) a set of n previous edges. In this way, edges in the M0 graph have only one metric value (i.e., the metric is now an M0 metric) and state-of-the-art algorithms based on the OSP can operate properly.

2) *Algorithm Description for M1 Metrics*: The GTA algorithm for M1 metrics is illustrated in Fig. 4. The blue circle nodes correspond to the original M1 graph. From this graph, each node is copied and then duplicated as many times as it has ingress links. Each M0 node (green square nodes in Fig. 4) then corresponds to an original M1 node and to one ingress link of this node (including the “null” ingress link). Then, M0 edges are obtained by creating edges towards the created M0 nodes from all the M0 nodes corresponding to the source of the original edge to which the destination M0 node corresponds. Each M0 edge then corresponds to (i) an edge of the M1 graph, and (ii) an ingress link to this edge. Hence, each M0 edge can be assigned a static M0 metric value corresponding to the metric value of the original edge when the given ingress edge is used to reach it. We refer to this procedure as $GTA()$.

3) *Request and Result Transformation*: In order for the original algorithm to run on the transformed graph, the original request has to be mapped. First, the source nodes now correspond to their M0 equivalent which have no ingress link. Secondly, the destination nodes have now several M0 equivalents. To overcome this problem, so-called *sink nodes* have to be created. All the M0 nodes corresponding to the same original M1 node have to be connected to the same sink node with edges whose metric value does not change the metric value of the overall solution (that is, e.g., 0 for an additive metric). The destination(s) of the original request then become(s) the corresponding sink node(s). We refer to the procedure of creating the sink nodes as $ADD_SINKS()$. Once the algorithm found a solution on the M0 graph, the solution on the M1 graph can be recovered by taking all the M1 equivalents of the elements of the M0 solution returned by the algorithm.

4) *Algorithm Description for M_n Metrics*: In order to transform a graph with M_n metrics, the $GTA()$ procedure described in Sec. IV-C2 simply has to be applied n times and followed by the addition of sink nodes (Fig. 5). Indeed, as

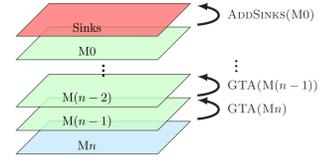


Fig. 5: Illustration of the GTA procedure for M_n metrics. The procedure described in Fig. 4 simply has to be applied n times, provided that the sinks are only added at the end.

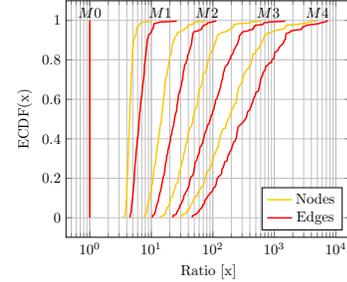


Fig. 6: Evolution of the amount of nodes and edges for the Internet Topology Zoo [17] topologies for different amount of executions of GTA (including the sinks creation).

$GTA()$ duplicates edges for each ingress link, applying it n times will duplicate edges for each possible set of n ingress links and hence lead to an M0 graph.

5) *Cost of the Transformation*: The size of the M0 graph increases with n . Fig. 6 shows the evolution of the amount of nodes and edges for all the topologies from the Internet Topology Zoo [17]. We can observe that the amount of nodes and edges increases by up to one order of magnitude for each application of the $GTA()$ procedure. That is, while GTA allows to optimally solve any problem with M_n metrics using algorithms for M0 metrics, this comes at the price of a huge increase in the graph size. An insight on the runtime impact will be given in Sec. V.

V. EVALUATION

The goal of the evaluation is to observe the impact of M1 and M_∞ metrics on the optimality and completeness of state-of-the-art algorithms based on the OSP and to show that our proposed solutions are correct. In particular, we show the influence of M1 and M_∞ metrics on one SP and one CSP algorithm both with and without GTA . Algorithms are compared to A*Prune, which provides a benchmark for both completeness and optimality.

A. Setup

1) *Topologies*: We use the topologies from the Internet Topology Zoo [17] which are connected and have more than 10 nodes. Further, because A*Prune poorly scales both in terms of memory consumption and runtime [13], we filter out topologies with more than 100 nodes or 200 edges.

2) *Requests*: For a given topology, source and destination nodes for requests are randomly selected from the whole set of nodes of the topology.

3) *Metrics*: We define one M0, one M1 and one M_∞ metric. The metric values are random values between 1 and 2. For the M0 metric, the values are defined for each edge. For the M1 metric, the values are defined for each combination of edge and previous edge. For the M_∞ metric, the values are defined for each path.

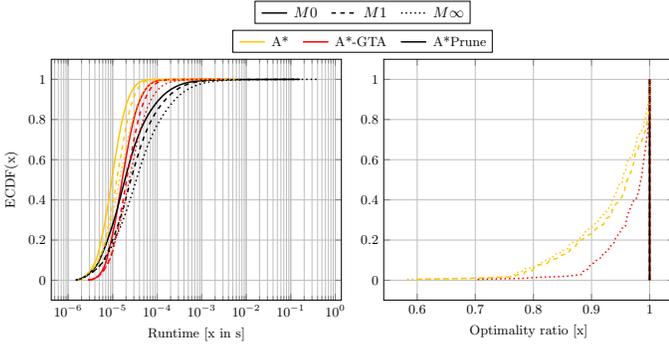


Fig. 7: Runtime and optimality ratio of A*, A* with one GTA transformation (A*-GTA) and A*Prune for the SP problem with M0, M1 and M ∞ optimization metrics.

B. Measurement Environment

In both SP and CSP scenarios, each algorithm is ran 20.000 times for each topology and metric type. Prior to these 20.000 runs, 1.000 warm-up runs are used to prevent the Java Hotspot optimizer from influencing the runtime measurements. The evaluation ran on an Intel Core i7-4790 CPU @ 3.60GHz.

C. Shortest Path: Optimality Influence

1) *Setup*: We observe the runtime and optimality of A* [11] (an optimal SP algorithm) and A* with GTA applied once (referred to as A*-GTA) for M0, M1 and M ∞ metrics using A*Prune as a benchmark. The guess values for both A*Prune and A* correspond to the hop count.

2) *Results: Runtime*: The left plot of Fig. 7 shows the empirical cumulative distribution function (ECDF) of the observed running times during the simulation³. We can observe that the GTA transformation for A* leads to an increase in runtime of around half an order of magnitude. Even though big topologies were not used, the figure also illustrates the poor scalability of A*Prune. Indeed, while A*Prune is sometimes significantly faster than A*-GTA (around 40% of the cases for the M0 metric, around 30% of the cases for the M1 metric and around 20% of the cases for the M ∞ metric) its runtime becomes very high for bigger topologies. Note that the runtime increase for each algorithm for the M1 and M ∞ metrics is mostly due to the increased complexity of the metric values computation.

3) *Results: Optimality Ratio*: The right plot of Fig. 7 shows the ECDF of the *optimality ratio* observed for each algorithm and metric types. The optimality ratio is defined per topology as the percentage of requests that the algorithm was able to solve optimally. As expected, both A* and A*-GTA are always optimal for M0 metrics and show a sub-optimal behavior for M ∞ metrics. However, while A* presents a sub-optimal behavior for M1 metrics, A*-GTA does not. The sub-optimality of A* for M1 metrics corresponds to the behavior described in Sec. I-A and Fig. 1. This confirms that GTA allows state-of-the-art algorithms based on the OSP to keep their original properties for M1 metrics. For M ∞ metrics, though GTA does not guarantee optimality for A*, we can observe that it improves its optimality ratio. All algorithms were complete, confirming that an M1 or M ∞ global optimization metric does not impact the completeness of algorithms (see Tab. I).

³Note that, for GTA, the graph creation is not taken into account. Indeed, it can be done once for all the requests.

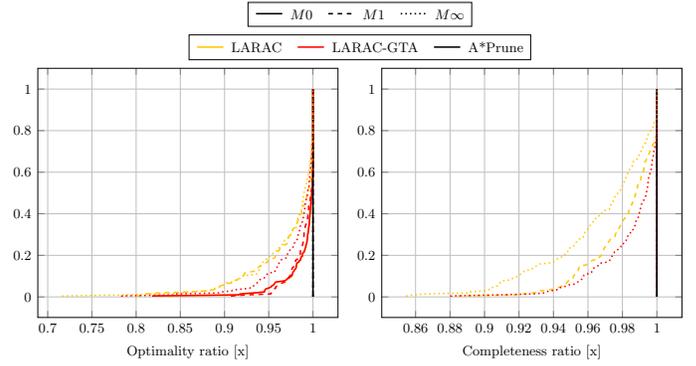


Fig. 8: Optimality and completeness ratios of LARAC, LARAC with one GTA transformation (LARAC-GTA) and A*Prune for the CSP problem with M0, M1 and M ∞ global constraint metrics and an M0 optimization metric.

4) *Results: Conclusions*: M1 and M ∞ global optimization metrics indeed lead to the sub-optimality of A*. While A*Prune provides optimality for any type of metric, it presents a problematic scalability behavior. For its part, GTA allows A* to optimally solve problems for M1 metrics and improves its optimality ratio for M ∞ metric, at the price of a reasonable increased runtime. Note that EBD led to an identical optimality behavior and a similar runtime behavior as A*-GTA.

D. Constrained Shortest Path: Completeness Influence

1) *Setup*: Because CBF, the optimal CSP algorithm, presents an exponential runtime behavior [12], we use LARAC [18]–[21], a sub-optimal but fast and complete CSP algorithm. We then observe the optimality and completeness of LARAC and LARAC with GTA applied once (referred to as LARAC-GTA) for M0, M1 and M ∞ metrics, using A*Prune as benchmark. Note that, because of the poor scalability of A*Prune and the higher runtime required by a CSP search compared to an SP search, we further reduced the topologies to those with less than 50 nodes and 100 edges.

2) *Metrics and Constraint Bounds*: For the optimization metric, we use an M0 metric. The constraint metric upper bound is randomly distributed among all the possible values (between the minimum and the maximum values).

3) *Results: Runtime*: Because of space constraints and because the runtime impact of GTA appeared to be the same as for the SP problem, we omit the runtime values for the CSP problem.

4) *Results: Optimality Ratio*: The left plot of Fig. 8 shows the ECDF of the optimality observed for each algorithm. LARAC and LARAC-GTA present exactly the same behavior for the M0 metric (the yellow curve being hidden in Fig. 8). As expected, LARAC is not optimal for the M0 metric. For LARAC, the M1 metric reduces its optimality ratio. However, the M ∞ metric does not appear to further reduce this optimality ratio. For LARAC-GTA, the M1 metric does not have a big influence on its optimality. However, the M ∞ metric appears to further reduce this optimality ratio. It is interesting to notice that this behavior is different for LARAC and LARAC-GTA.

5) *Results: Completeness Ratio*: The right plot of Fig. 8 shows the ECDF of the *completeness ratio* observed for each algorithm. The completeness ratio is defined per topology as the percentage of requests for which the algorithm was able to find a solution. The same conclusions as for the optimality ratio of the SP simulations can be drawn. Indeed, GTA allows

LARAC to be complete for the M_1 metric and improves on its completeness ratio for the M_∞ metric.

6) *Results: Conclusions:* The M_1 and M_∞ global constraint metrics indeed lead to the incompleteness of LARAC. GTA allows LARAC to completely solve problems for the M_1 metric and improves its completeness ratio for the M_∞ metric.

VI. APPLICABILITY OF THE SOLUTIONS

A*Prune and EBD can only be used for the SP/MCSP and SP problems, respectively. On the other hand, GTA can be applied to any routing algorithm. That is, GTA can be used for any routing problem (unicast, multicast, multipath, etc.) and amount of optimization and constraint metrics as long as a state-of-the-art algorithm for the corresponding problem with M_0 metrics exists. However, we identify that multipath routing requires small adaptations. Indeed, after a GTA transformation, disjointness on the transformed graph does not guarantee disjointness of the corresponding elements on the original graph and multipath routing algorithms might hence return disjoint paths on the transformed graph which are not disjoint on the original graph. To circumvent this problem, the GTA transformation can be adapted by adding an intermediate *sink edge* to which all the transformed edges corresponding to an identical original edge connect. In this way, if the algorithm finds two paths that use different edges corresponding to the same original edge, it will have to use the same sink edge and will hence conclude that these paths are not disjoint. This additional procedure however again comes at the price of an increased amount of nodes and edges in the transformed graph.

VII. CONCLUSIONS

State-of-the-art routing algorithms are optimal and complete when using metrics that satisfy the *optimal substructure property* (OSP). However, we have shown that relevant QoS metrics such as delay or buffer consumption do not satisfy this property. Hence, the algorithms lose their optimality and/or completeness. This causes operators to violate Service Level Agreements (SLAs), hence incurring penalties, as well as to inefficiently use their network resources.

In order to still guarantee optimal and complete results, we first proposed a new M_n metric taxonomy for classifying routing metrics based on the amount n of previously traversed edges needed to compute their value at a given edge. Based on this taxonomy, we presented solutions guaranteeing optimality and completeness. First, we presented A*Prune [5], a state-of-the-art algorithm that can deal with any type of M_n and M_∞ metric for solving the shortest path (SP) and multi-constrained shortest path (MCSP) problems. Second, we proposed *edge-based Dijkstra* (EBD), a newly proposed modification of Dijkstra for solving SP problems with M_1 metrics. Finally, because A*Prune and EBD can only be used for particular problems and metric types, we proposed a *graph transformation algorithm* (GTA) that allows any state-of-the-art algorithm for any routing problem (e.g., unicast, multicast, multi-constrained, etc.) to solve problems with M_n metrics. While A*Prune is the only opportunity for optimally solving a problem with M_∞ metric, we have shown that it presents a poor scalability behavior. Besides, on the example of the A* [11] and LARAC [18]–[21] state-of-the-art algorithms, we

have shown that GTA indeed recovers their properties for M_n metrics, at the cost of an increased running time.

The effective impact of the completeness and optimality behavior of state-of-the-art algorithms and of the proposed algorithms on the network resources usage of operators is an interesting future research direction and is left for future work.

ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 671648 (VirtuWind).

REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network flows: theory, algorithms, and applications*. Prentice Hall, 1993.
- [2] J. L. Gross and J. Yellen, *Graph theory and its applications*. CRC press, 2005.
- [3] P. Paul and S. V. Raghavan, "Survey of QoS routing," in *Proceedings of the 15th International Conference on Computer Communication (ICCC)*, 2002, pp. 50–75.
- [4] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [5] G. Liu and K. Ramakrishnan, "A*Prune: an algorithm for finding k shortest paths subject to multiple constraints," in *Proc. IEEE INFOCOM*, vol. 2, 2001, pp. 743–749.
- [6] R. Baumann, S. Heimlicher, M. Strasser, and A. Weibel, "A survey on routing metrics," *TIK report*, vol. 262, pp. 1–53, 2007.
- [7] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. MIT Press, 2009.
- [8] R. Bellman, "On a routing problem," *Quarterly of Applied Mathematics*, vol. 16, no. 1, pp. 87–90, Apr. 1958.
- [9] L. R. Ford Jr, "Network flow theory," DTIC Document, Tech. Rep., 1956.
- [10] E. F. Moore, *The Shortest Path Through a Maze*. Bell Telephone System, 1959.
- [11] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [12] R. Widyono, "The design and evaluation of routing algorithms for real-time channels," International Computer Science Institute Berkeley, Tech. Rep. TR-94-024, 1994.
- [13] J. W. Guck, A. Van Bemten, M. Reisslein, and W. Kellerer, "Unicast QoS routing algorithms for SDN: A comprehensive survey and performance evaluation," *IEEE Communications Surveys & Tutorials*, 2018.
- [14] J.-Y. Le Boudec and P. Thiran, *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*. Springer Verlag, 2001.
- [15] A. Van Bemten and W. Kellerer, "Network calculus: A comprehensive guide," *Technical University of Munich, Chair of Communication Networks, Technical Report No. 201603*, October 2016.
- [16] J. W. Guck, A. Van Bemten, and W. Kellerer, "Detserv: Network models for real-time QoS provisioning in SDN-based industrial environments," *IEEE Transactions on Network and Service Management*, vol. 14, no. 4, pp. 1003–1017, Dec 2017.
- [17] S. Knight *et al.*, "The internet topology zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 9, pp. 1765–1775, October 2011.
- [18] Y. P. Aneja and K. Nair, "The constrained shortest path problem," *Naval Research Logistics Quarterly*, vol. 25, no. 3, pp. 549–555, Sep. 1978.
- [19] G. Y. Handler and I. Zang, "A dual algorithm for the constrained shortest path problem," *Networks*, vol. 10, no. 4, pp. 293–309, 1980.
- [20] D. Blokh and G. Gutin, "An approximate algorithm for combinatorial optimization problems with two parameters," *Australasian Journal of Combinatorics*, vol. 14, pp. 157–164, 1996.
- [21] A. Jüttner, B. Szviatovski, I. Mécs, and Z. Rajkó, "Lagrange relaxation based method for the QoS routing problem," in *Proc. IEEE INFOCOM*, vol. 2, 2001, pp. 859–868.