

Generic Dijkstra: correctness and tractability

Ireneusz Szcześniak*, Bożena Woźna-Szcześniak†

*Częstochowa University of Technology, Department of Computer Science
ul. J. H. Dąbrowskiego 73, 42-201 Częstochowa, Poland

†Jan Długosz University in Częstochowa, Department of Mathematics and Computer Science
al. Armii Krajowej 13/15, 42-200 Częstochowa, Poland

Abstract—The recently-proposed generic Dijkstra algorithm finds shortest paths in networks with continuous and contiguous resources. The algorithm was proposed in the context of optical networks, but is applicable to networks with finite and discrete resources. The algorithm was published without a proof of correctness, and with a minor shortcoming. We provide that missing proof and offer a correction to the shortcoming. To prove the algorithm correct, we generalize the Bellman’s principle of optimality to algebraic structures with a partial ordering. By analyzing the size of the search space in the worst-case, we argue the stated problem is tractable. Thus we definitely answer a long-standing fundamental question of whether we can efficiently find a shortest path in a network with discrete resources subject to the continuity and contiguity constraints: yes, we can.

I. INTRODUCTION

The Dijkstra shortest-path algorithm finds shortest paths in a graph from a given source vertex to all other vertexes, provided the costs of edges are non-negative [1]. However, the algorithm cannot be used if the paths found have to meet the resource *continuity* and *contiguity* constraints.

Some networks have *finite* and *discrete* resources which have to be used by a connection along the path. For instance, in communication networks, frequency or time is such resource. The resource is divided into discrete units modeled by a set of integers $\Omega = [0, U)$, where U is a *finite* number of units offered. Each edge has a set of available units, which is a subset of Ω .

A path meets the resource continuity constraint if some given units are available on each of its edges. Therefore an algorithm should find not only a path, but also the set of units available along the path. When a connection is established along the found path, the units found are made unavailable along the path, thus changing the state of the network, and influencing the future path searches.

The continuity constraint follows from the characteristics of the modeled network. For instance, in optical networks, a unit is a wavelength or a frequency unit, which has to be used along the path in order to take advantage of the wavelength-division multiplexing. In networks capable of advance resource reservation, a unit is a time slot, which has to be used along the path in order to schedule transmission at a given time along the edges used.

The resource contiguity constraint requires that the units of a path be contiguous. Contiguous units can be modeled by a

half-closed integer interval. For example, the units of interval $[0, 3)$ are contiguous, and the units of set $\{0, 2\}$ are not.

The contiguity constraint also follows from the characteristics of the modeled network. For instance, in the elastic optical networks, the unit contiguity constraint follows from the physical and economical limitations of optical networks, where the frequency units of a connection should occupy a single frequency band. In networks capable of advance resource reservation, the (time slot) contiguity constraint follows from the assumption that the connection should last unintermitted.

The generic Dijkstra algorithm [2] finds efficiently shortest paths that meet the unit continuity and contiguity constraints. The generic Dijkstra algorithm is a generalization of the standard Dijkstra algorithm: while the standard Dijkstra requires the ordering (of the cost labels) to be total, the generic Dijkstra relaxes this requirement and allows for a partial ordering.

Our novel contribution is the generalization of the Bellman’s principle of optimality, the proof of correctness of the generic Dijkstra algorithm, and a correction to a minor shortcoming of the algorithm. By proposing the worst-case analysis of the search space, we argue that the problem is tractable, a crucial conclusion for network operations and management with the ever increasing requirements of densification, agility, performance and reliability.

II. RELATED WORKS

In [2], the routing problem was presented in the context of optical networks. In that article also the signal modulation constraint was considered, i.e., a path was dropped if it was unable to support a demand, while here we drop no paths. The generic Dijkstra algorithm was adapted to solve the dynamic routing problem with dedicated path protection [3]. The algorithm was simulatively demonstrated to efficiently find exact results [4], [2].

Bellman formulated the *principle of optimality* in [5], which is also known as the *dynamic programming principle*. For the shortest-path problem, this principle stipulates that shortest paths are made of shortest paths [6], i.e., shortest paths have the optimal substructure, and therefore form a shortest-path tree. This principle is not exactly a solution to a problem (which Bellman readily acknowledged in [6] that the principle does not offer a ready computational scheme), but rather a stipulation the solution should meet [7].

In [8] the authors proposed an algorithm for finding a shortest path in a network with continuous and discrete resources. While we believe the algorithm is exact, the authors call it heuristic, albeit in the context of optimizing the overall network performance.

The status of this routing problem had been unclear: no proof of nondeterministic-polynomial (NP) completeness was proposed, no efficient algorithm with proven correctness has been published but heuristic algorithms were commonly used, most notably based on a K-shortest path algorithm or on the Dijkstra algorithm [9]. However, searching for a path that meets the unit continuity and contiguity constraints is an easy (in terms of complexity theory) problem, because the number of shortest paths is polynomially bounded [2].

The proofs of the correctness of the Dijkstra algorithm have been proposed in a number of mile-stone text books. Usually the proofs are inductive, and combined with contradiction. We find the most convincing the proof offered in [10], which does not rely on contradiction, and that strategy we use in the proof we propose.

The Martins algorithm is a general multilabeling algorithm for solving multicriteria optimization problems with real-valued criteria [11]. The algorithm could be used to solve the stated problem but with the exponential worst-case complexity. The generic Dijkstra algorithm is similar to the Martins algorithm in that it is also multilabeling, but different in that it is of a single criterion with partial ordering.

III. PRELIMINARIES AND TERMINOLOGY

A. Resource interval

We refer to a half-closed integer interval as a resource interval (RI). We describe a set of contiguous units with an RI, and refer to it with r . For instance, by $r = [10, 12)$ we mean units 10 and 11. Function $\min(r)$ returns the lower endpoint of r , and $\max(r)$ the upper endpoint. For example, $\min(r) = 10$, and $\max(r) = 12$.

We are interested in the inclusion relation between RIs, because the larger RI, the better, as it describes also the included RIs. For instance, $r_1 = [0, 1)$ is worse than $r_3 = [0, 2)$ (or r_3 is better than r_1), because r_3 properly includes r_1 , i.e., $r_3 \supset r_1$.

Relation \subset is a strict partial ordering as there can be RIs for which neither \subset nor the converse (i.e., \supset) holds. Two RIs are \subseteq -incomparable (or incomparable in short), denoted by \parallel , if one does not include the other, properly or not. For instance, $r_2 = [1, 3)$ is incomparable with r_1 , i.e., $r_1 \parallel r_2$, because neither \subset , \supset , nor equality holds.

The \subset -induced \parallel relation is intransitive. For instance, $r_1 \parallel r_2$ and $r_2 \parallel r_3$ does not imply $r_1 \parallel r_3$ because $r_1 \subset r_3$. Therefore, \parallel is not an equivalence relation, and so ordering \subset is not strict weak.

Standard sorting requires at least a strict weak ordering, and to that end we introduce a total ordering $<$, an extension of \subset , as defined in Table I, where a cell in its left part reports on the \subset relation, and in the right on $<$. The converse $r_i > r_j$ means $r_j < r_i$.

Relation $r_i \supseteq r_j$ has to imply $r_i \leq r_j$ (the unshaded cells in Table I) because r_i should be processed first as it offers a better solution (see Section V). The shaded cells in Table I report on one of two possibilities of defining $<$ for the incomparable RIs so that $<$ is transitive. The other possibility would have the relation flipped in both gray cells. The choice is arbitrary, and we prefer the RIs of smaller lower endpoints go first, i.e., $r_i < r_j$ if $\min(r_i) < \min(r_j)$. The ordering defined by (1) is lexicographic: the lower endpoints are compared with $<$, and the upper with $>$.

$$\begin{aligned} r_i < r_j &\iff \min(r_i) < \min(r_j) \text{ or} \\ &\min(r_i) = \min(r_j) \text{ and} \\ &\max(r_i) > \max(r_j) \end{aligned} \quad (1)$$

Lemma 1. *Relation $<$ for RIs is transitive.*

Proof. Relation $<$ is transitive if $r_i < r_j < r_k$ implies $r_i < r_k$. Relation $r_i < r_j$ holds in either of two cases: 1st, if $\min(r_i) < \min(r_j)$; and 2nd, if $\min(r_i) = \min(r_j)$ and $\max(r_i) > \max(r_j)$. If both $r_i < r_j$ and $r_j < r_k$ hold in the 1st case, then $\min(r_i) < \min(r_j) < \min(r_k)$ holds, and so $r_i < r_k$ does. If $r_i < r_j$ holds in the 1st case, and $r_j < r_k$ in the 2nd, then $\min(r_i) < \min(r_j) = \min(r_k)$ holds, and so $r_i < r_k$ does; similarly, if $r_i < r_j$ holds in the 2nd case, and $r_j < r_k$ in the 1st. If both $r_i < r_j$ and $r_j < r_k$ hold in the 2nd case, then $\min(r_i) = \min(r_j) = \min(r_k)$ and $\max(r_i) > \max(r_j) > \max(r_k)$ hold, and so $r_i < r_k$ does. \square

B. Network model

The network is modeled by a weighted, directed multigraph $G = (V, E)$, where $V = \{v_i\}$ is a set of vertexes, and $E = \{e_i\}$ is a set of edges. For edge e_i , function $\text{cost}(e_i)$ gives its cost, and function $\text{AU}(e_i)$ gives its set of available units (i.e., a set of integers), which do not have to be contiguous. Function $\text{target}(e_i)$ gives the target vertex of edge e_i . Function $\text{I}(v_i)$ gives the set of incoming edges for v_i .

Path $p = (e_i)$ is a sequence of edges e_i where neighboring edges meet at the same vertex. A path has a cost, and an RI. The RI is available on every edge of the path.

The algebraic structure of the cost should have two operators defined. First, the relation $<$ operator should be transitive. Second, the \oplus operator with the identity 0 should calculate the combined cost based on its two operands. For instance, integer and real numbers with the $<$ and $+$ operators meet these requirements.

The cost of a path is the cost of its edges combined with \oplus as given by (2). We assume that appending edge e to path p cannot decrease the cost of the path, i.e., $\text{cost}(p) \oplus \text{cost}(e) \geq \text{cost}(p)$. If \oplus is the addition operator, then $\text{cost}(e) \geq 0$, the requirement of the Dijkstra algorithm.

$$\text{cost}(p) = \bigoplus_i \text{cost}(e_i) \quad (2)$$

TABLE I: Relations between RIs r_i and r_j .

	$\max(r_i) < \max(r_j)$		$\max(r_i) = \max(r_j)$		$\max(r_i) > \max(r_j)$	
$\min(r_i) < \min(r_j)$	$r_i \parallel r_j$	$r_i < r_j$	$r_i \supset r_j$	$r_i < r_j$	$r_i \supset r_j$	$r_i < r_j$
$\min(r_i) = \min(r_j)$	$r_i \subset r_j$	$r_i > r_j$	$r_i = r_j$		$r_i \supset r_j$	$r_i < r_j$
$\min(r_i) > \min(r_j)$	$r_i \subset r_j$	$r_i > r_j$	$r_i \subset r_j$	$r_i > r_j$	$r_i \parallel r_j$	$r_i > r_j$

C. Label

While we search for paths, we find it more efficient (because of the dynamic programming principle) to describe a path with a *label* defined as a pair of cost and an RI. A path is a feasible solution, so a label is feasible too. We denote labels with l .

A vertex can be reached by different paths of the same cost and RIs, i.e., equivalent paths of equal labels. We are interested only in one of these equivalent paths chosen arbitrarily, and so we do not allow a vertex to have equal labels. Different vertexes can have equal labels, but these labels would represent different paths.

The label has the cost and the RI of its path. The cost of label l is denoted by $\text{cost}(l)$, and the RI by $\text{RI}(l)$. For instance, label $l = (0, \Omega)$ is of $\text{cost}(l) = 0$, and $\text{RI}(l) = \Omega$. Function $\text{edge}(l)$ gives the edge of the label, i.e., the edge that was appended last to produce the label.

A set C of candidate labels l' is produced when to a path described with label l an edge e is appended, which is denoted by $l \oplus e$ and defined by (3). We say that l' is derived from l . Labels l' have equal cost that depends on the cost of l and e , i.e., $\text{cost}(l') = \text{cost}(l) \oplus \text{cost}(e)$. However, their RIs differ because $\text{RI}(l) \cap \text{AU}(e)$ can fan out to a set of RIs. The RI of l' is in $\text{RI}(l)$ and in $\text{AU}(e)$ to meet the continuity constraint.

$$C = l \oplus e = \{l' : \text{cost}(l') = \text{cost}(l) \oplus \text{cost}(e) \text{ and } \text{RI}(l') \in \text{RI}(l) \cap \text{AU}(e)\} \quad (3)$$

Labels are compared for two reasons. First, with \prec for relaxation to determine whether a label should be kept or dropped. Second, with $<$ for sorting to determine which label should be retrieved from the priority queue first. Table II shows relations between two labels l_i and l_j depending on their costs and RIs, where a cell in its left part reports on the \prec relation, and in the right on $<$.

1) *The \prec relation:* Label l_i is better than l_j (denoted by $l_i \prec l_j$) in two cases: first, if it offers at a lower cost a better or equal RI; second, if it offers at a lower or equal cost a better RI. Better or equal relation is denoted by \preceq . The converse $l_i \succ l_j$ means $l_j \prec l_i$. A better label is kept, and a worse discarded. Relation \prec is used to define label efficiency.

Definition 2 (Label efficiency). *Label l is efficient if, for a given vertex, there does not exist label l' such that $l' \prec l$.*

Proposition 3. *Relation $l \preceq l'$ holds for l' derived from l .*

Proof. Relation $l \preceq l'$ holds for two reasons. First, $\text{cost}(l) \leq \text{cost}(l')$ by an assumption of the network model that the cost of a path cannot decrease when an edge is appended. Second, $\text{RI}(l) \supseteq \text{RI}(l')$ by the continuity constraint. In Table II, the related cells are unshaded. \square

Relation \prec is a strict partial ordering as there can be labels for which neither \prec nor \succ holds. Two labels are \preceq -incomparable (or incomparable in short), denoted by \parallel , if one is not better than or equal to the other, which has two reasons: first, because their RIs are incomparable; second, because one of the labels offers a better RI at a larger cost than the other label. For instance, if $l_1 = (0, [0, 1])$ and $l_2 = (1, [0, 2])$, then $l_1 \parallel l_2$, because neither \prec , \succ , nor equality holds.

The \prec -induced label incomparability is not an equivalence relation as it is intransitive. For example, for labels of equal cost and incomparable RIs (shaded dark in Table II), the intransitivity of label incomparability follows from the intransitivity of RI incomparability: e.g., given $l_1 = (0, [0, 1])$, $l_2 = (0, [1, 3])$, and $l_3 = (0, [0, 2])$, relations $l_1 \parallel l_2$ and $l_2 \parallel l_3$ hold, but $l_1 \parallel l_3$ does not as $l_1 \prec l_3$ does.

For a priority queue, standard sorting cannot use \prec because of two reasons related to the label incomparability. First, \prec is not a strict weak ordering, because label incomparability is not an equivalence relation. Second, sorting must have ordering established between all nonequal labels, but \prec leaves some labels incomparable (shaded in Table II).

2) *The $<$ relation:* For sorting, (4) defines a total ordering $<$ so that either $<$ or $>$ holds for all nonequal labels. Ordering $<$ extends \prec , i.e., $< \supset \prec$, so that \prec implies $<$ because a better label should be processed first (see Section V).

$$l_i < l_j \iff \begin{aligned} &\text{cost}(l_i) < \text{cost}(l_j) \text{ or} \\ &\text{cost}(l_i) = \text{cost}(l_j) \text{ and} \\ &\text{RI}(l_i) < \text{RI}(l_j) \end{aligned} \quad (4)$$

The ordering is lexicographic by cost first, and by RI next. For labels with different costs (including the incomparable labels, shaded light in Table II), we compare costs only: $l_i < l_j$ if $\text{cost}(l_i) < \text{cost}(l_j)$, because the greedy strategy minimizes cost.

For labels with equal costs, we compare RIs: $l_i < l_j$ if $\text{cost}(l_i) = \text{cost}(l_j)$ and $\text{RI}(l_i) < \text{RI}(l_j)$. For the comparable labels (unshaded in the middle row of Table II), the choice is not arbitrary (i.e., we cannot flip the relation for RIs), because \preceq must imply \leq . However, for the incomparable labels (shaded dark in Table II), the choice is arbitrary: we could flip the relation for RIs (and $<$ for labels would still be transitive), but we prefer to keep (4) simple.

Lemma 4. *Relation $<$ for labels is transitive.*

Proof. Analogous to the proof of Lemma 1, provided $<$ is transitive for the values compared: for the cost it is by assumption, and for the RI by Lemma 1. \square

TABLE II: Relations between labels l_i and l_j .

	$\text{RI}(l_i) \subset \text{RI}(l_j)$		$\text{RI}(l_i) = \text{RI}(l_j)$		$\text{RI}(l_i) \supset \text{RI}(l_j)$		$\text{RI}(l_i) \parallel \text{RI}(l_j)$	
$\text{cost}(l_i) < \text{cost}(l_j)$	$l_i \parallel l_j$	$l_i < l_j$	$l_i < l_j$	$l_i < l_j$	$l_i < l_j$	$l_i < l_j$	$l_i \parallel l_j$	$l_i < l_j$
$\text{cost}(l_i) = \text{cost}(l_j)$	$l_i \succ l_j$	$l_i > l_j$	$l_i = l_j$		$l_i < l_j$	$l_i < l_j$	$l_i \parallel l_j$	$l_i < l_j$ if $\text{RI}(l_i) < \text{RI}(l_j)$ $l_i > l_j$ if $\text{RI}(l_i) > \text{RI}(l_j)$
$\text{cost}(l_i) > \text{cost}(l_j)$	$l_i \succ l_j$	$l_i > l_j$	$l_i \succ l_j$	$l_i > l_j$	$l_i \parallel l_j$	$l_i > l_j$	$l_i \parallel l_j$	$l_i > l_j$

IV. GENERIC PRINCIPLE OF OPTIMALITY

The principle of optimality was formulated for the shortest-path problem in [6] by (5) where: the shortest paths are found to target vertex N (from all other vertexes), there are no parallel edges, f_i is the cost of reaching v_N from v_i , and t_{ij} is the cost of the edge from v_i to v_j (of ∞ if the edge does not exist). Over the years, for various shortest-path problems, the principle has been reformulated.

$$\begin{aligned} f_i &= \min_j \{t_{ij} + f_j\} & \text{if } i \neq N \\ f_N &= 0 \end{aligned} \quad (5)$$

A. Reformulation

If we search for shortest paths from the source vertex s (to all other vertexes), and there are parallel edges, the principle of optimality can be formulated by (6) where \oplus can be any operator provided the equations hold for an optimal solution (i.e., of the lowest cost). Operator \oplus is usually $+$, but can also be, e.g., the multiplication operator if we search for, e.g., a path of the lowest probability of failure (i.e., highest availability). The relaxation of the Dijkstra algorithm uses this formulation with the addition operator.

$$\begin{aligned} f_s &= 0 \\ f_i &= \min_{e \in I(v_i)} \{f_{\text{source}(e)} \oplus \text{cost}(e)\} & \text{if } i \neq s \end{aligned} \quad (6)$$

Formulations (5) and (6) require a total ordering between labels (e.g., the ordering between real numbers), because they allow a vertex to have one label only. If the ordering between labels is partial, then a vertex can have a set of incomparable labels, and so the principle of optimality should be generalized.

B. Generalization

We propose the generalization given by (7), which we call the *generic principle of optimality*, and which could also be called the generic dynamic programming principle. The generalization describes a solution that we call the *efficient-path tree*, a parallel to the shortest-path tree described by the principle of optimality. C_e is the set of candidate labels we get after traversing edge e . A union of candidate labels for all edges that lead to v_i is the set of candidate labels for v_i . Then from the union we take a minimum to get the set of efficient labels P_i .

$$\begin{aligned} P_s &= \{(0, \Omega)\} \\ P_i &= \min \left\{ \bigcup_{e \in I(v_i)} C_e \right\} & \text{if } i \neq s \end{aligned} \quad (7)$$

The minimum of a set C of labels is a set of incomparable labels, as given by (8), i.e., worse labels are discarded.

$$\min C = \{l \in C : \forall l' \neq l \in C \ l \parallel l'\} \quad (8)$$

Candidate labels C_e depend on the efficient labels $P_{\text{source}(e)}$ of the source vertex of e , and e itself, as given by (9), which (depending on the definition of \oplus) could be or not constrained.

$$C_e = P_{\text{source}(e)} \oplus e \quad (9)$$

C. Constriction

We could include in C_e every label $l' \in l \oplus e$ for every label $l \in P_{\text{source}(e)}$ but that could produce a label with an empty RI. While such a label could be considered feasible for some problems and would meet the requirements of the generic principle of optimality, it would not describe a path that meets the resource continuity and contiguity constraints. Therefore we allow labels with nonempty RIs, as defined by (10). If empty RIs were allowed, the efficient-path tree would include the shortest-path tree.

$$P_{\text{source}(e)} \oplus e = \{l' \in l \oplus e : l \in P_{\text{source}(e)} \text{ and } \text{RI}(l') \neq \emptyset\} \quad (10)$$

V. A SHORTCOMING OF THE GENERIC DIJKSTRA

The generic Dijkstra algorithm sorts the labels in the priority queue in the ascending order of their cost only [2], and that is a shortcoming, but only when adding an edge to a path does not increase the cost of the path.

A. A failing example

Figure 1 shows an example where the efficient path from vertex s to t goes through vertex u , is of cost 1 and the RI of $[0, 2)$. However, generic Dijkstra, depending on the implementation, could mistakenly find efficient a path of cost 1 and the RI of $[0, 1)$ that has edge e_1 only.

When visiting vertex s , the algorithm relaxes edges e_1 and e_2 . Next, label $(1, [0, 1))$ could be retrieved from the priority queue first, thus erroneously finding that label (for the path of edge e_1 only) efficient, while it is label $(1, [0, 2))$ (for the path of edges e_2 and e_3) that is efficient.

B. A correction of the shortcoming

Labels in the priority queue should be sorted using the \leq relation. Labels for the same vertex are $<$ -comparable, but in the queue there can be equal labels for different vertexes, and so the \leq relation has to be used. For any labels l_i and l_j in the priority queue, label l_i will be retrieved before l_j if $l_i \leq l_j$. Equal labels are retrieved in arbitrary order. At the top of the

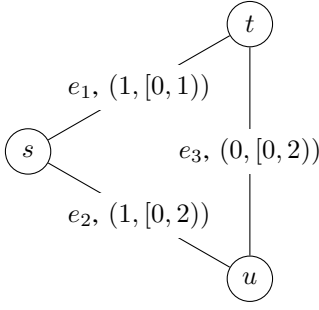


Fig. 1: A sample failing example.

queue there is the label that is \leq -comparable with every other label in the queue.

With the correction, the failing example is solved correctly. When visiting vertex s , the algorithm relaxes edges e_1 and e_2 . In the queue there are two labels: $l_1 = (1, [0, 1))$, and $l_2 = (1, [0, 2))$. Even though $l_2 \prec l_1$, label l_1 is kept as it is for a different vertex. Label l_2 is retrieved from the queue first, because $l_2 < l_1$, and so vertex u is visited next, and label $l_3 = (1, [0, 2))$ is inserted into the queue. Label l_3 is retrieved next, since $l_3 < l_1$, thus label l_3 is found efficient for vertex t .

VI. PROBLEM STATEMENT

Prove the correctness of the generic Dijkstra algorithm that solves the following problem.

Given:

- weighted, directed multigraph $G = (V, E)$, where $V = \{v_i\}$ is a set of vertexes, and $E = \{e_i\}$ is a set of edges,
- cost function $\text{cost}(e_i)$, which gives the cost of edge e_i ,
- available units function $\text{AU}(e_i)$, which gives the available units of edge e_i ,
- a transitive relation $<$ for cost,
- appending an edge to a path does not decrease the cost of the path,
- an efficient label meets the generic principle of optimality,
- the set of all units Ω on every edge,
- the source vertex s .

Find:

- an efficient-path tree rooted at s that meets the resource continuity and contiguity constraints.

In the efficient-path tree, each vertex has a minimal and complete set of efficient labels. Therefore we search for a minimal and complete set of efficient labels. Minimal, because from the set of equal labels for a vertex (that describe equivalent paths) we choose one label arbitrarily. Complete, because we produce all efficient labels, i.e., there does not exist an efficient label that we could add to the set.

A. The algorithm

Algorithm 1 shows the generic Dijkstra algorithm, and Algorithm 2 shows the relaxation procedure.

The labels that are processed by the algorithm are called known, the other labels are unknown, and as the search progresses, more labels become known. From an efficient label, relaxation derives a candidate label that becomes tentative if no better or equal label is known. A tentative label is incomparable with other known labels, and can be efficient. A tentative label is discarded if a better label is found later, or becomes permanent when it is found efficient. Known labels for v_i are in the set of permanent labels P_i and tentative labels T_i . All known labels are in $P = \{P_i\}$ and $T = \{T_i\}$. T is organized as a priority queue with the \leq ordering.

Algorithm 1 Generic Dijkstra

In: graph G , source vertex s

Out: an efficient-path tree

Here we concentrate on permanent labels.

$T_s = \{(0, \Omega)\}$ // The initial label.

while T is not empty **do**

$l = \text{pop}(T)$

$e = \text{edge}(l)$

$v = \text{target}(e)$

 // Add l to the set of permanent labels for vertex v .

$P_v = P_v \cup \{l\}$

for each out edge e' of v in G **do**

$\text{relax}(e', l)$

return P

VII. CORRECTNESS

We propose an inductive proof for the generic Dijkstra algorithm that also holds for the standard Dijkstra algorithm since a total ordering (required by the Dijkstra algorithm) is an extension of a partial ordering (required by the generic Dijkstra algorithm). A total ordering is more restrictive as it does not allow for incomparability, and that simplifies searching.

A. Intuition

Generic Dijkstra algorithm is correct for two reasons. First, from among the tentative labels (that are a link away from efficient labels), the priority queue provides at the top an efficient label l , because labels are sorted with \leq : the relation $l \leq l'$ holds for any other label l' in the queue, and so no label l' is better than l , i.e., $l \leq l' \iff l \preceq l' \text{ or } l \parallel l' \iff l' \not\prec l$. Second, relaxation replenishes the priority queue with labels l' derived from l , and so, by Proposition 3, they cannot be better than l .

Relaxation is not necessary for the algorithm correctness. Instead of relaxation, it would suffice to ensure a candidate label l' for vertex v' does not represent a loop (a path that returns to one of its vertexes) before inserting it into the queue, and is not worse than any permanent label for vertex v' before making it permanent. Such procedure, however, would be inefficient: a candidate label that is worse than any known label would be inserted, retrieved, and inevitably discarded.

Relaxation improves efficiency. By two steps. The first inserts into the queue a candidate label l' only if it is promising

Algorithm 2 relax

In: edge e' , label l *Here we concentrate on tentative labels.*

 $c' = \text{cost}(l) \oplus \text{cost}(e')$ $v' = \text{target}(e')$ **for each** RI $I' \neq \emptyset$ in $\text{RI}(l) \cap \text{AU}(e')$ **do** $l' = (c', I')$ // Can the candidate label l' become tentative?**if** $\nexists l_{v'} \in P_{v'} : l_{v'} \preceq l'$ **then****if** $\nexists l_{v'} \in T_{v'} : l_{v'} \preceq l'$ **then**// Discard tentative labels $l_{v'}$ such that $l' \preceq l_{v'}$,// leave in $T_{v'}$ only labels incomparable with l' . $T_{v'} = T_{v'} - \{l_{v'} \in T_{v'} : l' \preceq l_{v'}\}$ // Add l' to the set of tentative labels for vertex v' . $T_{v'} = T_{v'} \cup \{l'\}$

at the time of relaxation, i.e., incomparable with the known labels for vertex v' . The second discards the tentative labels for vertex v' which turn out to be worse than the candidate label l' to be inserted into the priority queue. Even though we look for worse labels, we use \preceq because its definition is simpler than that of \prec and because the equality would never hold, as checked by the first step.

B. Proof

Theorem 5. *The algorithm terminates with a complete set of efficient labels.*

Proof. We prove by induction. The induction step corresponds to an iteration of the main loop. The induction hypotheses are:

- 1) P has efficient labels derived from efficient labels,
- 2) T has incomparable labels derived from efficient labels.

Basis. $|P| = 1$. In the first iteration, the initial label $(0, \Omega)$ for vertex s is retrieved from T_s and added to P_s . This label is the root of the efficient-path tree, and is the only label that has not been derived.

The initial label is efficient because no better label could exist. Such a better label would describe a path from s to s either of cost lower than 0 (and that cannot be since, by the problem assumption, adding an edge cannot decrease cost), or of units that properly include Ω (and that cannot be either since, by the continuity constraint, traversing an edge cannot add units).

Inductive step. In an inductive step, efficient label l from the top of T is moved to P . Label l is for vertex v .

To prove that label l is efficient, we show that any other label l' for vertex v could not be better, i.e., $l' \prec l$ cannot hold. Label l' must be unknown because otherwise l would not be in T by the second hypothesis.

Label l' must be derived from a known label l_u for some preceding vertex u . Label l_u must be tentative because otherwise (i.e., l_u is permanent) l' would be known (tentative) by the induction hypotheses. The algorithm would never produce such l' derived from tentative l_u , but that would prove the algorithm incorrect if l' turns out to be efficient.

Relation $l_u \leq l'$ holds by Proposition 3. Relation $l \leq l_u$ holds, since it was label l that was retrieved from the top of the priority queue. Relation $l \leq l'$ holds, since $l \leq l_u \leq l'$, and since \leq is transitive by Lemma 4. Therefore $l' \prec l$ cannot hold (because $l' < l$ does not), and so l is efficient.

The efficient label l added to P preserves the first hypothesis. Relaxation preserves the second hypothesis using the equations of the generic principle of optimality: incomparable labels l' derived from l are produced, and labels worse than l' are discarded.

Termination. The algorithm terminates when the priority queue gets empty, and then set P is complete. The algorithm terminates as set P is finite: P is replenished by T , and the number of labels in T is polynomially bounded (see Section VIII). □

VIII. TRACTABILITY

We argue the problem is tractable because the size of the search space is polynomially bounded. The worst-case number L of labels to process, which is the number of incomparable labels for all vertexes, is $L = S|V|$, where S is the number of incomparable labels a vertex can have.

Number S depends on $|\Omega|$ only, and is given by (11). The maximal set of incomparable labels has labels whose cost increases along with the size of their RIs. The set has $|\Omega|$ subsets: the first has $|\Omega|$ labels with RIs of a single unit and the lowest cost; the second has $|\Omega| - 1$ labels with RIs of two units and a higher cost, \dots ; and the last has a single label with the RI of $|\Omega|$ units and the highest cost. The set therefore has $1 + 2 + \dots + |\Omega| = (|\Omega| + 1)|\Omega|/2$ incomparable labels.

Therefore, $O(L) = O(|\Omega|^2|V|)$, a polynomial bound.

$$S = \frac{(|\Omega| + 1)|\Omega|}{2} \quad (11)$$

IX. CONCLUSIONS

Routing of a single connection is one of the most important tasks of network operations and management. We have shown that in networks with discrete resources under the contiguity and continuity constraints, that task can be efficiently and exactly performed with the generic Dijkstra algorithm.

As future work, the average and worst-case complexities of time and memory could be analytically evaluated to allow for comparison with other algorithms, e.g., the filtered-graphs algorithm. The simulative results demonstrate that the generic Dijkstra algorithm is efficient. Whether the algorithm is the most efficient in comparison with other algorithms, perhaps even optimal, deserves further research.

X. ACKNOWLEDGMENT

We dedicate this work to Alexander Stepanov for his decades-long inspiration, and his contributions to generic programming and C++. This work was funded by the Polish Ministry of Education and Science under grant number 020/RID/2018/19.

REFERENCES

- [1] Edsger Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.
- [2] Ireneusz Szcześniak, Andrzej Jajszczyk, and Bożena Woźna-Szcześniak. Generic Dijkstra for optical networks. *IEEE/OSA Journal of Optical Communications and Networking*, 11(11):568–577, November 2019.
- [3] Ireneusz Szcześniak, Ireneusz Olszewski, and Bożena Woźna-Szcześniak. Towards an efficient and exact algorithm for dynamic dedicated path protection. *Entropy*, 23(9), 2021.
- [4] Piotr Jurkiewicz, Edyta Biernacka, Jerzy Domżał, and Robert Wójcik. Empirical time complexity of generic Dijkstra algorithm. In *2021 IFIP/IEEE International Symposium on Integrated Network Management*, pages 594–598, 2021.
- [5] Richard Bellman. On the theory of dynamic programming. *Proceedings of the National Academy of Sciences*, 38(8):716–719, 1952.
- [6] Richard Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16:87–90, 1958.
- [7] M. Sniedovich. Dijkstra’s algorithm revisited: the dynamic programming connexion. *Control and Cybernetics*, 35(3):599–620, 2006.
- [8] K.-C. Lee and V.O.K. Li. Routing and switching in a wavelength convertible optical network. In *IEEE INFOCOM ’93 The Conference on Computer Communications, Proceedings*, volume 2, pages 578–585, 1993.
- [9] Ireneusz Olszewski. Improved dynamic routing algorithms in elastic optical networks. *Photonic Network Communications*, 34(3):323–333, Dec 2017.
- [10] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice Hall, 1993.
- [11] Ernesto Queirós Vieira Martins. On a multicriteria shortest path problem. *European Journal of Operational Research*, 16(2):236 – 245, 1984.