

1 **MINIMUM SPANNING TREES IN INFINITE GRAPHS: THEORY**
2 **AND ALGORITHMS***

3 CHRISTOPHER THOMAS RYAN[†], ROBERT L. SMITH[‡], AND MARINA A. EPELMAN[§]

4 **Abstract.** We discuss finding minimum-cost spanning trees (MSTs) on connected graphs with
5 countably many nodes of finite degree. When edge costs are absolutely summable and an MST exists
6 (which is not guaranteed in general), we show that an algorithm that finds MSTs on finite subgraphs
7 (called *layers*) converges in objective value to the cost of an MST of the whole graph, as the sizes
8 of the layers grow to infinity. We call this the *layered greedy algorithm* since a greedy algorithm is
9 used to find MSTs on each finite layer. We stress that the overall algorithm is *not* greedy since edges
10 can enter and leave iterate spanning trees as larger layers are considered. However, in the setting
11 where the underlying graph has the *finite cycle* (FC) property (meaning, every edge is contained in
12 at most finitely many cycles) and distinct edge costs, we show that a unique MST T^* exists and
13 the layered greedy algorithm produces iterates that converge to T^* by eventually “locking in” edges
14 after finitely many iterations. Applications to network deployment are discussed.

15 **Key words.** minimum spanning trees, infinite graphs, infinite-dimensional optimization

16 **MSC codes.** 90C27, 90C35, 90C48

17 **1. Introduction.** The problem of finding minimum-cost spanning trees on finite
18 graphs is a classical combinatorial optimization problem with numerous applications
19 in practice [11, 13, 14, 20]. The problem is used as a subroutine or heuristic for
20 solving other graph optimization problems [3, 9, 27]. To our knowledge, an algorithmic
21 approach to the MST problem on *infinite* graphs has not been systematically pursued,
22 despite there being extensive literature on algorithms for infinite graphs in other
23 contexts (see, for instance, [2, 7, 10, 18]). Several references examine properties of
24 spanning trees in the limit of finite random graphs (see, for instance, [1, 4, 5, 16]),
25 but the focus of these papers is not on the questions of existence and performance of
26 algorithms, topics we emphasize here. The only paper we know of that deliberates on
27 producing an algorithm for finding MSTs in infinite graphs is [15] in the more general
28 context of infinite matroids (we discuss this paper in more detail below).

29 In a finite graph, an MST always exists and can be found by a greedy algorithm.
30 The MST problem on infinite graphs does not afford such luxuries. As we will show
31 through examples, an MST may not even exist in an infinite graph, and when it does,
32 it may not be reachable by a greedy algorithm.

33 In response to this, we develop an algorithm to tackle the MST problem (whenever
34 an MST exists) in any connected graph with countably many nodes of finite degree and
35 absolutely summable edge costs (see Assumption 2). This algorithm finds MSTs in a
36 growing sequence of finite subgraphs that, in the limit, converge in cost to that of an
37 MST of the original graph (see Theorem 3.3). We call this result — i.e., convergence
38 of the total cost of the edges of iterate trees to the total cost of the edges of an MST
39 — *convergence in objective value*.

*Submitted to the editors on 30 May 2023.

Funding: This work was funded by the Natural Sciences and Engineering Research Council of Canada Discovery Grant RGPIN-2020-06488.

[†]University of British Columbia, Sauder School of Business, (chris.ryan@sauder.ubc.ca, <https://www.sauder.ubc.ca/people/christopher-thomas-ryan>).

[‡]University of Michigan, Industrial and Operations Engineering, (rlsmith@umich.edu, <http://www-personal.umich.edu/~rlsmith/>).

[§]University of Michigan, Industrial and Operations Engineering, (mepelman@umich.edu, <http://www-personal.umich.edu/~mepelman/>).

40 The sequence of subgraphs considered by our algorithm are called *layers*, and so
 41 we call our algorithm the *layered greedy algorithm* since it applies a greedy algorithm
 42 repeatedly in a growing set of layers of the graph. To our knowledge, the concept of
 43 finite “layers” in infinite graphs was first introduced by Sharkey and Romeijn in [26],
 44 and has appeared in other papers on infinite dimensional graph theory, including [23].
 45 It is important to note that the layered greedy algorithm (as a whole) is not greedy
 46 since the MSTs found at each iteration must be computed “from scratch” and do not
 47 necessarily extend the previous MSTs from earlier layers in a greedy fashion. We also
 48 show, under an assumption akin to discounting of the edge costs in the graph, that
 49 finite termination of the infinite algorithm provides “good” solutions with bounded
 50 error in finite time.

51 The fact that the layered greedy algorithm guarantees convergence in objective
 52 value on a broad class of infinite graphs is the first important result in our paper.
 53 However, it naturally leads to three additional questions.

54 (Q1) We have convergence in objective value when an MST is known to exist
 55 in the original graph. How can we guarantee that an MST exists?

56 (Q2) Convergence in objective value is a nice feature, but we would also like
 57 convergence to an optimal solution. How can we ensure that the finite-
 58 sized iterates of the layered greedy algorithm converge to an MST of the
 59 original, infinite graph?

60 (Q3) Since the layered greedy algorithm is not greedy (but only locally greedy
 61 within layers), edges may come and go from iterate spanning trees as the
 62 algorithm proceeds. What are some sufficient conditions for an edge of
 63 the iterates to eventually “lock in” to an edge of the MST after finitely
 64 many iterations? Moreover, can these conditions be verified during the
 65 execution of the algorithm?

66 All three questions rely on careful consideration of the topological properties of
 67 the graph. Here, by “topological” we refer to questions of closedness, compactness,
 68 and convergence in the space of subgraphs of a given graph (and not the common
 69 alternative usage of referring to the arrangement of edges and nodes in a graph as its
 70 “topology”). Regarding (Q1), topology plays an important role because the existence
 71 of optimal solutions is naturally a topological consideration. As in other infinite-
 72 dimension optimization problems, we must argue that the feasible region has some
 73 notion of compactness and the objective function has some notion of continuity to
 74 conclude (using a Weierstrass-type result) that an optimal solution exists.

75 In (Q2), topology is implicit in the question itself because any notion of conver-
 76 gence requires a specification of topology. At first glance, (Q3) appears unrelated to
 77 topology, but the concept of “lock-in” has a topological flavor. Indeed, convergence
 78 via lock-in is precisely the notion of convergence in discrete topologies (see, for in-
 79 stance, Section 12 of [17]). This motivates our use of the *product discrete topology*,
 80 where convergence corresponds to lock-in edge by edge (here, the product is taken
 81 across edges).

82 With this topology, we can develop weak sufficient conditions for positive answers
 83 to the three questions we posed. In particular, we show (in Theorem 5.5) that the
 84 following condition:

85 (C1) *Finite Cycle (FC) property*: every edge of the graph is contained in at
 86 most finitely many cycles in the graph

87 is sufficient to guarantee that an MST always exists, answering (Q1). This follows
 88 by showing compactness of the set of spanning trees in the product discrete topology
 89 under the FC property (see Lemma 5.4).

90 If, additionally, the following holds:
 91 (C2) *Distinct Edge Costs*: no two edges have the same cost,
 92 then a *unique* MST always exists (Theorem 6.3). The uniqueness is useful in estab-
 93 lishing our answer to (Q2): the sequence of iterates of the layered greedy algorithm
 94 converges to the unique MST in the product discrete topology when (C1) and (C2)
 95 hold. Finally, because of the nature of convergence in the product discrete topology,
 96 (C1) and (C2) give conditions for lock-in of the edges, answering the first part of (Q3).
 97 At the end of Section 6, we provide verifiable conditions for discovery of these edges,
 98 that have lock-in in agreement with “early” edges of the MST, addressing the second
 99 part of the last question.

100 **Related work.** We add to the growing literature on algorithms for solving prob-
 101 lems on infinite graphs, including recent applications in deep learning (see the review
 102 article [28] for a summary of this work). Our work resembles a stream of work for
 103 solving network flow problems on infinite graphs [12, 18, 19, 21, 23, 26]. An important
 104 distinction between that line of work and what we pursue here is the definition of trees
 105 and connectedness. In the network flow literature, trees can be connected through a
 106 “node at infinity” that acts as a universal sink for flow generated at supply nodes in
 107 the graph. In contrast, our notion of connectivity is more classical, requiring nodes to
 108 be connected by a path of finitely many edges. Another important distinction is that
 109 network flow problems are typically studied as continuous optimization problems,
 110 allowing, for example, duality arguments and generalizations of the max-flow/min-
 111 cut theorem to infinite graphs [2]. By contrast, the MST problem is fundamentally
 112 discrete.

113 This paper is also related to the literature on infinite matroids (see, for instance,
 114 [8] and references therein). Here, the primary focus is on describing axiom systems
 115 for carefully defining the notion of infinite matroid to allow for a convenient matroid
 116 duality theory. As far as we know, little attention (other than [15]) has been given
 117 to the generality of greedy algorithms in infinite graphs. The workhorse of much of
 118 this infinite matroid theory is set-theoretic, using Zorn’s lemma to show the existence
 119 of maximal objects within infinite graphs. By contrast, our theory relies heavily on
 120 topological arguments, including Tychonoff’s Theorem, Weierstrass’s Theorem, and
 121 convergence proofs. Indeed, our layered greedy MST algorithm does not produce
 122 a “chain” of nested spanning trees that would be necessary to leverage Zorn-like
 123 arguments.

124 [15] proposes a “greedy” algorithm for finding bases in finitary infinite matroids
 125 (corresponding to MSTs in infinite graphs with nodes of finite degree). This algo-
 126 rithm, however, is shown to find an MST using transfinite induction. Infinite graph
 127 adaptations of greedy algorithms for finding MSTs in finite graphs may even fail to
 128 converge to trees or span the nodes of the graph. A greedy algorithm can be “in-
 129 definitely distracted” by an infinite subset of low-cost edges, never getting around to
 130 span other parts of the graph. Klee’s algorithm avoids this issue by continuing to
 131 analyze its execution after an infinite time is exhausted exploring a single subtree of
 132 a spanning tree. This is why the algorithm is called transfinite. The arguments in
 133 this paper do not use transfinite induction; in particular, we analyze the execution of
 134 algorithms without the use of ordinal numbers required in transfinite induction.

135 Of course, one can view the graph as a matroid [10] where the independent subsets
 136 of the edges of G are its forests. This view would allow to prove some of the results
 137 in this paper using matroid theory, but other results, such as solution convergence
 138 and early edge detection, are established exploiting the special properties we assume

139 about the graphical and cost structures. For this reason, we will use terminology
 140 and ideas familiar from studying finite graphs as much as possible, only delving into
 141 topics that are peculiar to infinite graphs when necessary and avoiding the even more
 142 general language of matroids altogether. We hope that this makes the paper more
 143 accessible to readers with little or no exposure to either matroids or infinite graph
 144 theory. It is an open direction for future research to examine the implications of our
 145 method for general infinite matroids.

146 The most closely related paper that we are aware of is [22], which considers a sim-
 147 ilar setup and concerns itself with related questions. The main contrast is that [22]
 148 studies *maximum*-cost spanning tree problems with positive costs and offers simple
 149 conditions where a greedy algorithm finds the maximum spanning tree. By contrast,
 150 this paper explores deficiencies with greedy thinking in the *minimum*-cost spanning
 151 tree setting with positive costs that do not arise when seeking a maximum-cost span-
 152 ning tree. Indeed, “very small” costs in the outer reaches of the graph can be readily
 153 handled when seeking a maximum spanning tree (they can effectively be ignored),
 154 but prove more delicate in minimization.

155 **Organization of the paper.** The paper is organized as follows. In Section 2,
 156 we describe the general class of infinite graphs that we consider and define the MST
 157 problem in this class. Section 3 presents the layered greedy algorithm and analyzes
 158 its convergence in objective value. Section 4 formalizes the finite cycle property (C1).
 159 In Section 5, we establish the existence of MSTs under the FC property. In Section 6,
 160 we establish convergence of iterates of the layered greedy algorithm to an optimal
 161 spanning tree under the FC property and the additional condition (C2) of distinct
 162 edge costs. We also explore conditions that allow for discovery of early edges of the
 163 infinite MST and its implications for applications. Section 7 concludes the paper.

164 2. The minimum spanning tree problem.

165 **2.1. Basic definitions.** Let $G = (\mathcal{V}, \mathcal{E})$ be an undirected graph with node set
 166 $\mathcal{V} = \{1, 2, \dots\}$ and edge set \mathcal{E} . Let $c : \mathcal{E} \rightarrow \mathfrak{R}$ denote an edge-cost functional for
 167 G . We will sometimes use c_{ij} to denote the cost $c(\{i, j\})$ of edge $\{i, j\} \in \mathcal{E}$ when
 168 convenient.

169 The set $I(i)$ denotes the nodes that are adjacent to node i , that is, $I(i) := \{j \in$
 170 $\mathcal{V} \mid \{i, j\} \in \mathcal{E}\}$. The *degree* of node i is the cardinality of $I(i)$. A graph is *locally*
 171 *finite* if every node has finite degree. A *path* in G is a finite sequence of *distinct* nodes
 172 i_1, i_2, \dots, i_n , where $\{i_k, i_{k+1}\} \in \mathcal{E}$ for $k = 1, \dots, n - 1$. A *ray* is an infinite sequence
 173 of distinct nodes i_1, i_2, \dots , where $\{i_k, i_{k+1}\} \in \mathcal{E}$ for $k = 1, 2, \dots$. Two nodes i and j
 174 are *connected* in G if there exists a path starting with node i and ending with node j .
 175 The graph G is *connected* if all pairs of nodes i and j in G are connected. We make
 176 the following assumption throughout the paper:

177 ASSUMPTION 1. *The graph G is locally finite and connected.* \triangleleft

178 A *cycle* in G is a finite sequence of nodes $i_1, i_2, \dots, i_n, i_1$, where i_1, i_2, \dots, i_n is a
 179 path and $\{i_1, i_n\} \in \mathcal{E}$. A *bi-ray* consists of a node i and two distinct rays, that is, rays
 180 (i, i_1, i_2, \dots) and (i, j_1, j_2, \dots) , where all intermediate nodes i_k and j_ℓ are distinct.

181 Let H be a subgraph of G and let $\mathcal{V}(H)$ and $\mathcal{E}(H)$ denote the set of nodes and
 182 edges in H , respectively. In this paper, we only consider subgraphs with no isolated
 183 nodes, that is, for every node $i \in \mathcal{V}(H)$, there exists an edge $\{i, j\} \in \mathcal{E}(H)$ for some
 184 node $j \in \mathcal{V}(H)$. In light of this, we will typically refer to a subgraph H simply by its
 185 set $\mathcal{E}(H)$ of edges, since the set of nodes is implicit once the edges are defined. The cost
 186 function can be extended to be defined on the collection $\mathcal{P}(\mathcal{E})$ of subsets of edges of \mathcal{E}

187 (corresponding to subgraphs), where $c(H) := \sum_{e \in H} c(e)$ for any $H \in \mathcal{P}(\mathcal{E})$. (Without
 188 further assumptions, $c(H)$ might not be well-defined if the set $\mathcal{E}(H)$ is infinite. In the
 189 next section, we will make an assumption that the edge cost functional is absolutely
 190 summable — see Assumption 2 — to address this concern.)

191 A *forest* F of G is an acyclic subgraph of G , i.e., a subgraph of G without cycles.
 192 A connected forest is a *tree*. If a subgraph of G has node set \mathcal{V} , it is said to *span* G .
 193 A connected spanning forest is called a *spanning tree*.

194 We designate one of the nodes in G as its *root node* r . (The theory developed below
 195 is largely indifferent to which node in G is called the root node; only in subsections 3.3
 196 and 6.2 do we develop results that are tied to an *a priori* selection of r .) The first
 197 *layer* of nodes, denoted L_1 , consists of node r and all nodes that are adjacent to r ;
 198 that is, $L_1 := \{r\} \cup I(r)$. We define other layers recursively:

$$199 \quad L_{n+1} := L_n \cup \{i \in I(j) \text{ for some } j \in L_n\}, \quad n = 1, 2, \dots,$$

200 and sometimes refer to $\{r\}$ as layer 0 (cf. the definition of layers in [26]). Since G is
 201 locally finite and connected, each layer contains a finite number of nodes, every node
 202 is included in some layer, and once a node is in layer L_n , it is in every subsequent
 203 layer L_k for $k > n$. Let $G_n := (L_n, \mathcal{E}_n)$ for $n \geq 1$ denote the subgraph of G , where
 204 $\mathcal{E}_n := \{\{i, j\} \in \mathcal{E} \mid i, j \in L_n\}$ is the set of edges in the subgraph induced by the set of
 205 nodes L_n (we also use the term “layer n ” to refer to G_n).

206 **2.2. Formal statement of the minimum spanning tree problem.** Recall
 207 that the cost $c(T)$ of a spanning tree T of G is the sum of the costs of the edges of T ,
 208 i.e., $c(T) = \sum_{\{i,j\} \in \mathcal{E}(T)} c_{ij}$. Our problem is to find a minimum-cost spanning tree of
 209 G , i.e., solve

$$210 \quad (\text{P}) \quad c^* := \inf\{c(T) \mid T \text{ is a spanning tree of } G\}.$$

211 We call any optimal solution T^* of (P) a minimum spanning tree (MST). We say
 212 G *possesses an MST* if (P) has an optimal solution (that is, the infimum in (P) is
 213 attained).

214 **3. The layered greedy algorithm.** We now present the algorithm we analyze
 215 in this paper. The algorithm generates a sequence of spanning trees on finite restric-
 216 tions of the graph. We show that this sequence has nice convergence properties.

Algorithm 3.1 Layered greedy algorithm

- 1: **Input:** A locally finite and connected graph $G = (\mathcal{V}, \mathcal{E})$ with edge costs.
 - 2: **Initialize:** Set $n \leftarrow 1$ and T to be the empty subgraph of G with empty node set and empty edge set.
 - 3: **while** T is not a spanning tree **do**
 - 4: **Find MST on next layer:** Find an MST T^n on layer G_n using Prim’s algorithm (for completeness, we give a description of Prim’s algorithm below).
 - 5: Set $T \leftarrow T^n$ and $n \leftarrow n + 1$.
-

217 While most of the forthcoming analysis of the layered greedy algorithm is agnostic
 218 to the particular method used to find the MSTs on the layers in Step 4, Prim’s
 219 algorithm is instrumental in the early discovery of edges of an MST on G , which
 220 we discuss in subsection 6.2. It is one of the classical greedy algorithms for finding




FIG. 1. Graph for Example 1 illustrating that the layered greedy algorithm is not a greedy algorithm overall.

221 MSTs on *finite* graphs (see [3] for further details). In the usual statement of Prim's
 222 algorithm, the starting node that initializes the graph is arbitrary. We want ours to
 223 proceed from the root node r .

Algorithm 3.2 Prim's algorithm (for finding an MST on G_n)

- 1: **Input:** Graph $G_n = (L_n, \mathcal{E}_n)$ with edge costs.
 - 2: **Initialize:** Initialize a tree F to be the root node r .
 - 3: **while** F does not span G_n **do**
 - 4: **Append an edge:** Append to F the minimum-cost edge of \mathcal{E}_n emanating from F (that is, having one node in F and one outside of F), breaking ties arbitrarily.
-

224 It is important to note that while Prim's algorithm can be leveraged to find the
 225 tree iterates T^n on each of the finite graphs G_n , we may remove as well as add edges
 226 as we grow the layers G_n . The next example demonstrates this point.

227 **EXAMPLE 1.** Consider the ladder graph in Figure 1 with labeled nodes and edge
 228 costs written next to the edges. If node $r = 1$ is the root node, the nodes in layer 1 are
 229 nodes r , 2, and 3. The MST of graph G_1 consists of edges $\{r, 2\}$ and $\{r, 3\}$ for a cost
 230 of 11. The second layer has node set $\{r, 2, 3, 4, 5\}$. Now we can avoid the expensive
 231 edge $\{r, 3\}$ to construct the MST of G_2 consisting of the edges $\{r, 2\}$, $\{2, 4\}$, $\{3, 4\}$,
 232 and $\{3, 5\}$, for a total cost of 4. In other words, the cheapest edges in a given iteration
 233 (in this case, $\{r, 3\}$) may become too expensive by comparison as the subgraph grows,
 234 and get dropped in later iterations. \triangleleft

235 **3.1. Some preliminaries.** To analyze the performance of the layered greedy
 236 algorithm, we need a few preliminaries. First, we start with a classical result in
 237 infinite graph theory.

238 **PROPOSITION 3.1** (Proposition 8.1.1 in [10]). *Any locally finite and connected*
 239 *graph (Assumption 1) contains a spanning tree.*

240 Second, we need a mechanism for extending iterates of the layered greedy algo-
 241 rithm, which are not spanning trees of the entire graph G , into spanning trees.

242 **PROPOSITION 3.2.** *Suppose $T^n = (L_n, \mathcal{E}(T^n))$ is a spanning tree on the connected*
 243 *subgraph corresponding to the n -th layer graph $G_n = (L_n, \mathcal{E}_n)$. Then there exists a set*
 244 *of edges $\bar{\mathcal{E}} \subseteq \mathcal{E} \setminus \mathcal{E}_n$ such that $(\mathcal{V}, \mathcal{E}(T^n) \cup \bar{\mathcal{E}})$ is a spanning tree on G .*

245 *Proof.* Let \bar{G} be the graph obtained by removing from G nodes L_n and all edges
 246 incident to them (including both edges \mathcal{E}_n within the n -th layer and the edges con-

247 necting nodes in L_n to nodes in $L_{n+1} \setminus L_n$). Each connected component of \bar{G} satisfies
 248 Assumption 1, and therefore contains a spanning tree (Proposition 3.1). Moreover,
 249 each connected component of \bar{G} contains at least one node that belongs to $L_{n+1} \setminus L_n$
 250 — select one of these nodes in each connected component and select one of the edges
 251 connecting it to layer n . Then the union of T^n , the aforementioned spanning trees on
 252 the connected components of \bar{G} , and the selected edges that connect these connected
 253 components to L_n (and thus T^n) is a spanning tree on G . \square

254 Third, we must impose an additional assumption on the cost functional.

255 ASSUMPTION 2. *The edge cost functional $c : \mathcal{E} \rightarrow \mathfrak{R}$ is absolutely summable, i.e.,*
 256 $\sum_{e \in \mathcal{E}} |c(e)| < \infty$. \triangleleft

257 If we label the costs of the countably many edges in \mathcal{E} by c_ℓ for $\ell = 1, 2, \dots$, then
 258 Assumption 2 becomes $c = (c_1, c_2, \dots) \in \ell_1$ (where ℓ_1 is the vector space of absolutely
 259 summable sequences).

260 **3.2. Convergence in objective value.** We are now ready to prove a main
 261 result of the paper.

262 THEOREM 3.3. *Suppose G is a locally finite and connected graph (Assumption 1)
 263 whose edge cost functional is absolutely summable (Assumption 2). If G possesses an
 264 MST of cost c^* then the layered greedy algorithm converges in objective value; that is,
 265 the sequence T^n of iterates satisfies $c(T^n) \rightarrow c^*$.*

266 *Proof.* Let T^* be an MST of a locally finite connected graph G , and let T_n^* denote
 267 the restriction of T^* to G_n . By construction, $c(T_n^*) \rightarrow c(T^*) = c^*$ as $n \rightarrow \infty$. Note
 268 that T_n^* is a forest on G_n , although not necessarily a spanning tree. It can be extended
 269 to a spanning tree on G_n with the addition of a finite number of edges (since G_n is a
 270 finite graph). Let \bar{T}_n^* be the cheapest such extension and define

$$271 \quad \epsilon'_n := c(\bar{T}_n^*) - c(T_n^*).$$

272 Since T^n is an MST on G_n , we have

$$273 \quad (3.1) \quad c(T^n) \leq c(\bar{T}_n^*) = c(T_n^*) + \epsilon'_n.$$

274 Since T^* is a spanning tree of G , for every pair of nodes i and j in G , there is a unique
 275 finite path P_{ij} connecting them in T^* . Moreover, path P_{ij} must be wholly contained
 276 in layer $G_{n_{ij}}$ for $n_{ij} = \max_{k \in P_{ij}} \ell(k)$, where $\ell(k)$ is the number of the smallest layer
 277 containing node k . Let

$$278 \quad (3.2) \quad m(n) := \max\{m \mid n_{ij} \leq n \text{ for all } i, j \in G_m\}.$$

279 In other words, given n , $m(n)$ is the number of the largest layer such that all pairs of
 280 nodes in this layer are connected in T^* by paths wholly contained in G_n .

281 Note that none of the edges added to T_n^* to construct \bar{T}_n^* are in $G_{m(n)}$, since
 282 every pair of nodes in $G_{m(n)}$ is already connected by a path in T_n^* . Hence, $\epsilon'_n =$
 283 $c(\bar{T}_n^*) - c(T_n^*) \leq \epsilon_{m(n)}$, where

$$284 \quad (3.3) \quad \epsilon_{m(n)} := \sum_{e \in \mathcal{E} \setminus \mathcal{E}_{m(n)}} |c(e)|$$

285 is the sum of the absolute values of costs of edges outside of layer $G_{m(n)}$.

286 Observe that $m(n) \rightarrow \infty$ as $n \rightarrow \infty$, which follows from the finiteness of the path
 287 P_{ij} between any two nodes i and j and local finiteness and connectedness of G . Hence
 288 $\epsilon_{m(n)} \rightarrow 0$ as $n \rightarrow \infty$ since $c \in \ell_1$ by Assumption 2.

289 By Proposition 3.2, T^n can be extended to span G . Let S^n denote one such
 290 extended spanning tree, with additional edges $\mathcal{E}(S^n) \setminus \mathcal{E}(T^n) \subseteq G \setminus G_n$, and let

$$291 \quad \Delta_n := c(S^n) - c(T^n).$$

292 Observe that, by construction, $\Delta_n \rightarrow 0$ as $n \rightarrow \infty$. Now,

$$293 \quad (3.4) \quad c^* \leq c(S^n) = c(T^n) + \Delta_n \leq c(T_n^*) + \epsilon'_n + \Delta_n,$$

294 where the first inequality holds since c^* is the cost of an MST on G and the second
 295 inequality holds by (3.1). Since, as $n \rightarrow \infty$, $\Delta_n \rightarrow 0$, $\epsilon'_n \leq \epsilon_{m(n)} \rightarrow 0$, and $c(T_n^*) \rightarrow$
 296 $c(T^*) = c^*$, (3.4) implies $c(S^n) \rightarrow c^*$ and $c(T^n) \rightarrow c^*$ as $n \rightarrow \infty$, establishing the
 297 result. \square

298 **3.3. Error bound after finite termination.** We are also interested in the
 299 question of how fast the costs $c(T^n)$ of the iterates T^n approach the optimal value c^* .
 300 To provide a partial answer, we need the following additional assumption (which is
 301 only made in this subsection and not in the rest of the paper).

302 ASSUMPTION 3. *The graph $G = (\mathcal{V}, \mathcal{E})$ and the cost function $c : \mathcal{E} \rightarrow \Re$ satisfy*
 303 *the following: (i) there exist $\beta \in (0, 1)$ and $\gamma \in (0, +\infty)$ such that for every edge*
 304 *$\{i, j\} \in \mathcal{E}$, $0 \leq c_{ij} \leq \gamma \beta^{\min\{\ell(i), \ell(j)\}}$, where $\ell(i)$ is the number of the smallest layer*
 305 *containing node i , and (ii) there exists a uniform bound M on the cardinality of node*
 306 *degrees in G , with $M < 1/\beta$. \triangleleft*

307 Note that part (i) of Assumption 3 is tied to a particular selection of root node
 308 r and corresponding specification of the layers. It is, however, easy to show that if
 309 Assumption 3 is satisfied for a particular choice of r , it will also be satisfied for any
 310 other choice if the value of γ is adjusted accordingly. The iterates T^n will also change
 311 with a different choice of r , so it is natural that the bound in Proposition 3.4 depends
 312 on a particular choice of the root node.

313 Under Assumption 3, we can prove the following.

314 PROPOSITION 3.4. *Let S^n denote the extensions to spanning trees (via Propo-*
 315 *sition 3.2) of the iterates T^n produced by the layered greedy algorithm. Under the*
 316 *assumptions of Theorem 3.3 and Assumption 3, the errors in cost satisfy the follow-*
 317 *ing bound:*

$$318 \quad (3.5) \quad 0 \leq c(S^n) - c(T^*) \leq \frac{M\gamma}{(1-\delta)}(\delta^n + \delta^{m(n)}),$$

319 where $m(n)$ is defined in (3.2) and $\delta = M\beta < 1$.

320 *Proof.* This proof refers to several bounds established in the course of the proof
 321 of Theorem 3.3. We can bound

$$322 \quad (3.6) \quad 0 \leq c(S^n) - c(T^*) \leq c(T_n^*) - c(T^*) + \epsilon'_n + \Delta_n \leq \epsilon'_n + \Delta_n,$$

323 where the first inequality follows by optimality of T^* , the second inequality reproduces
 324 (3.4), and the last inequality follows because T_n^* is a subgraph of T^* , and the edge
 325 costs are nonnegative by Assumption 3(i).

326 Recall that, by definition, L_n is the set of all nodes that are at most n edges
 327 “away” from the root node r , i.e., for every node in L_n , there exists a path between
 328 that node and r that is at most n edges long.

329 Let $\epsilon_n := \sum_{e \in \mathcal{E} \setminus \mathcal{E}_n} c(e)$ be the sum of the costs of all edges in $\mathcal{E} \setminus \mathcal{E}_n$.¹ From
 330 Assumption 3(ii), the number of edges joining layer n to layer $n+1$ is bounded above
 331 by M^{n+1} . This follows by induction on the layer number, noticing that the maximum
 332 number of nodes in L_n is M times the number of nodes in L_{n-1} . Moreover, the cost
 333 of each edge joining layer n to $n+1$ is bounded above by $\gamma\beta^n$, by Assumption 3(i).
 334 Combining these observations, we establish

$$\begin{aligned}
 335 \quad \Delta_n \leq \epsilon_n &= \sum_{e \in \mathcal{E} \setminus \mathcal{E}_n} c(e) \leq \sum_{m=n}^{\infty} M^{m+1} \gamma \beta^m \\
 336 \quad &= M\gamma (M\beta)^n \sum_{m=0}^{\infty} (M\beta)^m = M\gamma \delta^n \sum_{m=0}^{\infty} \delta^m = M\gamma (\delta^n / (1 - \delta)).
 \end{aligned}$$

337 As part of the proof of Theorem 3.3, we showed that ϵ'_n can be bounded above
 338 by $\epsilon_{m(n)}$, and so

$$339 \quad \epsilon'_n \leq \epsilon_{m(n)} \leq M\gamma (\delta^{m(n)} / (1 - \delta)).$$

340 Substituting these bounds into (3.6), we derive

$$341 \quad 0 \leq c(S^n) - c(T^*) = M\gamma (\delta^n / (1 - \delta)) + M\gamma (\delta^{m(n)} / (1 - \delta)) = \frac{M\gamma}{(1 - \delta)} (\delta^n + \delta^{m(n)}),$$

342 as required. \square

343 From the proof of Theorem 3.3, we know $m(n) \rightarrow \infty$ as $n \rightarrow \infty$ and so the error
 344 bound in (3.5) converges to 0 as n grows. Of course, there remains the question of
 345 assessing the rate at which the sequence $m(n)$ grows with n to further analyze the
 346 convergence rate of the algorithm. The growth rate of $m(n)$ depends on the structure
 347 of the graph, and different MSTs can give rise to different functions $m(n)$.

348 Let $L(m)$ be the maximum number of edges over all paths P_{ij} in the tree T^*
 349 connecting nodes i and j in layer G_m . Note that $L(m) < \infty$ since G_m is finite. For
 350 $i, j \in G_m$, we have $n_{ij} = \max_{k \in P_{ij}} \ell(k) \leq m + L(m)$. Moreover, $\{m \mid n_{ij} \leq n \text{ for all}$
 351 $i, j \in G_m\} \supseteq \{m \mid m + L(m) \leq n\}$. Hence, $m(n) = \max\{m \mid n_{ij} \leq n \text{ for all } i, j \in$
 352 $G_m\} \geq \max\{m \mid m + L(m) \leq n\} = \max\{m \mid L(m) \leq n - m\}$. Now, note that $L(x)$ is
 353 increasing in positive real numbers x so that $\max\{x \mid L(x) \leq n - x\}$ is attained at a
 354 unique positive real solution $x(n)$ to the equation $L(x) = n - x$. Thus $m(n) = \lfloor x(n) \rfloor$;
 355 that is, $m(n)$ is the largest integer less than or equal to $x(n)$. This concrete formula
 356 can be used to assess the growth of the bound in (3.5), if one has an understanding of
 357 the function $L(m)$ and its connection to the structure of an optimal tree T^* in specific
 358 applications.

359 **REMARK 1.** *Since we employ Prim’s Algorithm to find an MST in layer G_n , the*
 360 *computational time in iteration n of the layered greedy algorithm is $O(|L_n|^2)$. This,*
 361 *together with (3.5), yields a bound on the computational time to find a spanning tree*
 362 *achieving a cost error within a pre-specified error from optimal. \triangleleft*

¹We introduced similar notation in equation (3.3) in the proof of Theorem 3.3; here, it is no longer necessary to take absolute values since the costs are assumed to be nonnegative.




FIG. 2. A graph where FC holds (see Example 2).

363 **4. The finite cycle property.** In Theorem 3.3, we showed that the layered
 364 greedy algorithm satisfies convergence in objective value (under Assumptions 1 and 2)
 365 whenever the graph possesses an MST. This naturally leads to the question of what
 366 graphs possess MSTs (question (Q1) in the introduction). In this section, we describe
 367 an elegant sufficient condition (and prove it suffices for existence in the next section).

368 We say that a graph satisfies the *finite cycle* (FC) property if every edge is con-
 369 tained in at most finitely many cycles of G . The graph in Figure 1 fails the FC
 370 property because the edge $\{1, 2\}$ is in infinitely many cycles in the graph. The next
 371 example satisfies the FC property.

372 **EXAMPLE 2.** Consider the graph in Figure 2. Observe that every edge lies in a
 373 unique cycle in the graph, and thus satisfies the FC property. \triangleleft

374 We capture the FC property in the following assumption, and refer to this as-
 375 sumption whenever the FC property is invoked later in the paper:

376 **ASSUMPTION 4.** The graph G satisfies the FC property.

377 Before moving on to studying the implications of the FC property for the MST
 378 problem, we take a brief detour to discuss the simple sufficient condition of absence
 379 of bi-rays for a graph to satisfy the FC property. Because the proof of this result will
 380 take us off the main path of our development, we put it in an appendix. The reader
 381 should be aware, however, that the proof relies on the contents of Section 5.

382 **PROPOSITION 4.1.** If G contains no bi-rays, then G satisfies the FC property.

383 *Proof.* See Appendix A. \square

384 Clearly, the converse of Proposition 4.1 is not true. Consider again the graph in
 385 Figure 2. The bottom path connecting all of the “triangle” pieces is a bi-ray, but the
 386 graph nonetheless satisfies the FC property.

387 **5. Existence of a minimum spanning tree.** Proposition 3.1 shows that a
 388 spanning tree always exists, but this does not ensure that an optimal solution to the
 389 MST problem (P) exists. Consider the following example.

390 **EXAMPLE 3.** Consider the one-way-infinite ladder graph in Figure 3, with top
 391 and bottom rays of 0-cost edges connected by infinitely many rungs with decreasing
 392 costs. The most expensive spanning tree has cost 1, consisting of the left-most rung
 393 of cost 1 connecting the top and bottom rays. A spanning tree of cost $1/4$ is drawn in
 394 non-dashed edges in the figure. One can similarly construct spanning trees of cost $1/8$,
 395 $1/16$, etc. Thus, a sequence of spanning trees whose costs converge to 0 can be found
 396 in the graph. However, no spanning tree has cost 0 since all edges have nonnegative
 397 cost and the 0-cost edges do not form a connected graph. Therefore, a minimum-cost




FIG. 3. A graph with no minimum spanning tree (see Example 3).

398 *spanning tree does not exist.* ◁

399 To establish existence, we will use Weierstrass’s standard optimization result (see,
 400 for instance, Theorem 2.35 in [6]) that minimizing a continuous function over a compact set always yields a minimizer. The challenge here is to develop the appropriate
 401 notion of topology to define continuity and compactness.
 402

403 **5.1. The product discrete topology.** Our desire to apply Weierstrass’s Theorem to (P) motivates the following notion of convergence.²
 404

405 DEFINITION 5.1. A sequence of subgraphs S^k of graph G converges to a subgraph
 406 S in G in the product discrete topology if there is a positive integer K_e for each edge
 407 $e \in \mathcal{E}$ such that for all $k \geq K_e$, $e \in S^k$ if and only if $e \in S$. We call this the lock-in
 408 property of edges of the sequence of subgraphs to the edges of the limiting subgraph.
 409 ◁

410 We can understand the use of the terminology “product” and “discrete” better in
 411 light of the following construction. For each edge $e \in \mathcal{E}$, define a set $B_e := \{0, 1\}$ and
 412 endow that set with the discrete metric $d_e(x, y) = 0$ if $x = y$ and 1 if $x \neq y$. That is,
 413 $d_e(0, 1) = d_e(1, 0) = 1$ and $d_e(0, 0) = d_e(1, 1) = 0$. Then, clearly, B_e is a metric space
 414 under metric d_e . There is a bijection between $\mathcal{P}(\mathcal{E})$ and the product $\prod_{e \in \mathcal{E}} B_e$, where
 415 $\mathcal{P}(\mathcal{E})$ is the power set of \mathcal{E} . Indeed, any subset H of \mathcal{E} corresponds to an element χ_H
 416 of $\prod_{e \in \mathcal{E}} B_e$ where $\chi_H(e) = 1$ if $e \in H$ and 0 otherwise (and vice versa). We call χ_H
 417 the characteristic function of the subset of edges H .

418 The product $\prod_{e \in \mathcal{E}} B_e$ can be endowed with the product topology τ of the discrete
 419 topologies on B_e for every $e \in \mathcal{E}$. By Theorem 3.36 in [6], the topology τ is metrizable.
 420 The significance of this for our purposes is that it suffices to consider subsequences
 421 (as opposed to nets) to establish topological properties involving τ . In particular, a
 422 set B in $\prod_{e \in \mathcal{E}} B_e$ is closed if every convergent (in τ) sequence χ_k of elements in B
 423 has a limit $\chi \in B$. Here, convergence in τ means that for every e , there exists a K_e
 424 such that $\chi_k(e) = \chi(e)$ for $k \geq K_e$. Moreover, compactness of a set in B is equivalent
 425 to sequential compactness (see Theorem 3.28 in [6]).

426 Returning to the product discrete topology on $\mathcal{P}(\mathcal{E})$, it can be seen as correspond-
 427 ing to the topology τ on $\prod_{e \in \mathcal{E}} B_e$ under the bijection $H \leftrightarrow \chi_H$. More precisely, a
 428 subset H of $\mathcal{P}(\mathcal{E})$ is open in the product discrete topology if and only if the subset
 429 $\{\chi_h \mid h \in H\}$ of $\prod_{e \in \mathcal{E}} B_e$ is open in τ . This notion defines a *product discrete topology*
 430 on the collection of all subgraphs on G , as defined in Definition 5.1. In particular, if
 431 S^k converges to S in the product discrete topology then, for any finite subset of \mathcal{E} ,
 432 the S^k ’s agree with S on this set of edges for sufficiently large k .

²Others use different notions of convergence, mostly based on the fact that they study random graphs and so are interested in probabilistic notions of convergence. See, for instance, [4].

433 **5.2. Cost continuity in the product discrete topology.** Having set our
 434 topology, we now want to establish the continuity and compactness needed for Weier-
 435 strass's Theorem. We start with establishing continuity of the objective function.

436 LEMMA 5.2. *Suppose the edge cost functional $c : \mathcal{E} \rightarrow \mathfrak{R}$ is absolutely summable*
 437 *(Assumption 2). Then $c(\cdot)$ is continuous in the product discrete topology.*

438 *Proof.* To establish continuity of $c(\cdot)$, it suffices to show that if a sequence H^k of
 439 elements of $\mathcal{P}(\mathcal{E})$ converges to H in the product discrete topology, then $c(H^k) \rightarrow c(H)$
 440 in the usual topology on the reals. That is, for an arbitrary $\epsilon > 0$, we want to show that
 441 there exists a K_ϵ such that $|c(H^k) - c(H)| < \epsilon$ for all $k \geq K_\epsilon$. Under Assumption 2,
 442 there exists a subset E of \mathcal{E} such that $E' = \mathcal{E} \setminus E$ is finite and $\sum_{e \in E} |c(e)| < \epsilon/2$.
 443 Since E' is a finite subset of \mathcal{E} , there exists a K_ϵ such that H^k agrees with H on all
 444 edges in E' for $k \geq K_\epsilon$ by the lock-in property. That is, for all $k \geq K_\epsilon$ we have

$$\begin{aligned}
 445 \quad |c(H^k) - c(H)| &= \left| \sum_{e \in H^k \cap E} c(e) + \sum_{e \in H^k \cap E'} c(e) - \sum_{e \in H \cap E} c(e) - \sum_{e \in H \cap E'} c(e) \right| \\
 446 \quad &= \left| \sum_{e \in H^k \cap E} c(e) - \sum_{e \in H \cap E} c(e) \right| \\
 447 \quad &\leq 2 \sum_{e \in E} |c(e)| < \epsilon.
 \end{aligned}$$

448 This establishes the result. \square

449 **5.3. Compactness in the product discrete topology.** The final ingredient
 450 in our existence proof is establishing the compactness of the set of spanning trees.
 451 The FC property is crucial to this argument. First, we state a preliminary lemma to
 452 establish the compactness of a superset.

453 LEMMA 5.3. *Let G be a locally finite and connected graph (Assumption 1). The*
 454 *space of all subgraphs of G is compact in the product discrete topology τ .*

455 *Proof.* Immediate from Tychonoff's theorem (Theorem 2.61 in [6]). \square

456 LEMMA 5.4. *Let G be a locally finite and connected graph (Assumption 1) that*
 457 *satisfies the FC property (Assumption 4). Then, the set of all spanning trees is com-*
 458 *compact in the product discrete topology.*

459 *Proof.* In light of Lemma 5.3, it suffices to show that the set of all spanning trees
 460 is closed in the product discrete topology.

461 Let $S^k, k = 1, 2, \dots$, be a sequence of spanning trees in G that converges in the
 462 product discrete topology to a subgraph S of G . It then suffices to show that S is,
 463 itself, a spanning tree. This is achieved in three parts: (i) show S is spanning, (ii)
 464 show S is acyclic, and (iii) show S is connected.

465 To establish (i), observe that if a node i is disconnected from S then each of the
 466 edges incident to i can only lie in finitely many of the iterates S^k . Then this means
 467 that node i is isolated in S^k for n sufficiently large, a contradiction of the fact that
 468 all S^k are connected.

469 To establish (ii), suppose that S contains a cycle C . Then, since C contains finitely
 470 many edges, the lock-in property of convergence in the product discrete topology
 471 implies that C is in each S^k for k sufficiently large. This contradicts the fact that
 472 each S^k is acyclic.

473 We now establish (iii). We will show that there is a path from i to j in S for any
 474 pair of nodes i and j . By connectedness of each S^k , there are paths P^k connecting i

475 and j in S^k for all k . Consider an arbitrary “reference” path P_{ij} in G connecting i
 476 and j . Path P_{ij} contains finitely many edges, and by the FC property, each edge is in
 477 at most finitely many cycles in G . Let us collect all these cycles into a finite collection
 478 of cycles $\tilde{\mathcal{C}}$, and let $\mathcal{C} := \{C \setminus P_{ij} \mid C \in \tilde{\mathcal{C}}\}$. That is, for every cycle $C \in \tilde{\mathcal{C}}$, the subset
 479 of edges of C that are not in the reference path P_{ij} is an element of \mathcal{C} . Again by the
 480 FC property, \mathcal{C} is a finite collection of subsets of edges in G .

481 Observe that each P^k arises by taking some edges from P_{ij} and some subsets of
 482 edges from \mathcal{C} (in the degenerate cases, P^k can exactly equal P_{ij} or just be composed
 483 of subsets of edges taken from \mathcal{C}). Thus, there are only finitely many possibilities
 484 for the structure of P^n since \mathcal{C} is a finite collection and P_{ij} has finitely many edges.
 485 According to the pigeonhole principle, infinitely many of the P^k are thus equal and
 486 so a subsequence of them converges in the product discrete topology to a path P that
 487 connects i and j . Since we have assumed that the S^k converge to S in the product
 488 discrete topology, this implies that P is in S and so i and j are connected in S . This
 489 implies that S is connected. \square

490 **THEOREM 5.5.** *Consider the minimum-cost spanning tree problem (P) and sup-*
 491 *pose G is a locally finite and connected graph (Assumption 1) with the FC property*
 492 *(Assumption 4) and with costs that are absolutely convergent (Assumption 2). Then,*
 493 *an MST (i.e., an optimal solution to (P)) exists.*

494 *Proof.* Note that (i) the objective function of (P) is continuous in the product
 495 discrete topology by Lemma 5.2, and (ii) the feasible region is compact in the product
 496 discrete topology by Lemma 5.4. The result then follows by Weierstrass’s theorem
 497 (Theorem 2.35 in [6]). \square

498 The above result implies that if the graph G has the FC property, then the
 499 layered greedy algorithm can be used to find a sequence of trees in G that converges
 500 to optimality in objective value (combining Theorems 3.3 and 5.5).

501 **6. Solution convergence.** In the previous section, we showed that if a graph
 502 is locally finite, connected, and satisfies the FC property with absolutely summable
 503 costs (Assuptions 1, 2, and 4) then the layered greedy algorithm always achieves
 504 convergence in objective value. However, this does not imply that the iterates of the
 505 graph converge to an MST. Consider the following example.

506 **EXAMPLE 4.** *Consider the graph in Figure 4, which satisfies Assuptions 1, 2, and*
 507 *4. If we apply the layered greedy algorithm, there is a tie between the two identical-cost*
 508 *vertical edges within each four-node cycle contained in the layer. Suppose for T^n with*
 509 *n odd, the algorithm chooses the “left” edges (shown as the dotted (purple) edges in*
 510 *Figure 4), and for T^n with n even, the algorithm chooses the “right” edges (shown as*
 511 *the dashed (green) edges in Figure 4). Then the sequence T^n does not converge in the*
 512 *product discrete topology at all, let alone to an MST. Thus, the iterates of the layered*
 513 *greedy algorithm can fail to converge. \triangleleft*

514

515 **REMARK 2.** *An astute reader may recognize that this example does have a conver-*
 516 *gent sequence of iterates if Prim’s algorithm uses a consistent tie-breaking rule across*
 517 *iterations. Indeed, we conjecture that it is possible to show more general convergence*
 518 *results than what we attempt below if we are careful in how Prim’s algorithm breaks*
 519 *ties. There is certainly precedent for this type of analysis in the literature: consider,*
 520 *for instance, [25, 24]. Pursuing this is a promising direction for future work, but is*
 521 *tangential to our focus here.*




FIG. 4. Graph for Example 4 illustrating that the layered greedy algorithm fails to find an optimal MST even when one exists.

522 **6.1. Solution convergence when there is a unique MST.** One sufficient
 523 condition to avoid pathological behavior illustrated in Example 4 is having a unique
 524 MST in the graph.

525 **THEOREM 6.1.** *Suppose G is a locally finite and connected graph (Assumption 1)*
 526 *that satisfies the FC property (Assumption 4) and whose edge cost functional is abso-*
 527 *lutely summable (Assumption 2). If G possesses a unique MST T^* then the extensions*
 528 *S^n (as defined in Proposition 3.2) of the iterates T^n of the layered greedy algorithm*
 529 *converge to T^* in the product discrete topology.*

530 *Proof.* Let T^n be the n -th iterate of the layered greedy algorithm. By Proposi-
 531 tion 3.2 each iterate can be extended to a spanning tree S^n of G . Suppose, by way of
 532 contradiction, that the sequence S^n does not converge to T^* in the product discrete
 533 topology. By the compactness of the set of spanning trees (Lemma 5.4), a subsequence
 534 S^{n_k} , $k = 1, 2, \dots$, converges to a spanning tree T' where $T' \neq T^*$. By convergence
 535 in objective value (Theorem 3.3) and continuity (Lemma 5.2), we conclude that T' is
 536 also an MST. Since T^* is the unique MST, this is a contradiction. \square

537 The following simple assumption is sufficient to ensure that a graph has at most
 538 one MST:

539 **ASSUMPTION 5.** *The graph G has distinct edge costs; that is, for every two dis-*
 540 *tinct edges $\{i, j\}$ and $\{i', j'\}$ we have $c_{ij} \neq c_{i'j'}$. \triangleleft*

541 To prove uniqueness under Assumption 5, we need the following generalization of
 542 a well-known condition in finite graphs (see, for instance, Theorem 13.1 in [3]). The
 543 proof follows identical logic to the finite case, and is thus omitted.

544 **PROPOSITION 6.2 (Cut optimality condition).** *If T^* is an MST of a locally finite*
 545 *and connected (Assumption 1) graph G then for all $\{i, j\} \in T^*$, $c_{ij} \leq c_{k\ell}$ for any edge*
 546 *$\{k, \ell\}$ crossing the cut formed by deleting edge $\{i, j\}$ from T^* .*

547 **THEOREM 6.3.** *Let G be a locally finite and connected graph (Assumption 1) with*
 548 *distinct edge costs (Assumption 5). If an MST exists for G then this MST is unique.*

549 *Proof.* To show uniqueness, suppose S and T are two distinct MSTs (at least
 550 one is guaranteed to exist by assumption), and let $\{i, j\} \in S \setminus T$. Furthermore, let
 551 $\{k, \ell\} \in T$ be in the cut created in G by removing $\{i, j\}$ from S . Since S and T
 552 are both MSTs, they both satisfy the cut optimality condition (Proposition 6.2), i.e.,
 553 $c_{ij} \leq c_{k\ell}$ and $c_{k\ell} \leq c_{ij}$, implying that $c_{ij} = c_{k\ell}$. This is a contradiction, establishing
 554 that $S = T$. \square

555 This result (via Theorem 6.1) shows that when we apply the layered greedy al-

556 gorithm to a locally finite, connected graph with the FC property and absolutely
 557 summable *distinct* edge costs, then the algorithm's iterates converge to an MST, i.e.,
 558 it provides an affirmative answer to question (Q2). Moreover, for each edge, we get
 559 lock-in after finitely many iterations via convergence in the product discrete topology.

560 **6.2. Discovery of early edges of an MST.** Of course, we would like a stronger
 561 convergence result than Theorem 6.1 in the following sense. Convergence in product
 562 discrete topology tells us that every edge *eventually* locks into an edge of an MST
 563 of G , but it would be better if we had a verifiable sufficient condition for when an
 564 edge has locked in. As we will see, the layered view of the graph and the nature
 565 of Prim's algorithm allow us to provide some partial results in this area. We note
 566 that the conditions we develop are tied to a particular selection of root node r and
 567 corresponding specification of the layers. These sufficient condition for early lock in
 568 are such that an edge e of the MST might satisfy them for some choices of r , but not
 569 for others.

570 In what follows, we adopt Assumption 5 that the graph has distinct edge costs.
 571 By Theorem 13.1 in [3], which is the finite-graph version of Theorem 6.3, this implies
 572 that for every n , T^n is the unique MST of the graph G_n and moreover, there will be
 573 no tie-breaking in Step 4 of Prim's algorithm.

574 With this assumption, we can make the following simple, yet powerful, obser-
 575 vation. Since in each iteration of the layered greedy algorithm the iterate T^n is
 576 constructed via Prim's algorithm, and because Prim's algorithm always starts with
 577 the root node and grows the tree T^n from there, the uniqueness in the choice of T^n
 578 greatly restricts the possibility of deviation in the "early" edges among the iterates
 579 T^n . The next result formalizes this idea.

580 Let e_k^n be the k -th edge added by Prim's Algorithm applied to G_n initialized with
 581 the root node r , where $k = 1, 2, \dots, |L_n| - 1$. We add a little more interpretation
 582 here for clarity. We are executing the layered greedy algorithm and are on its n -th
 583 iteration; that is, we are constructing T^n on the graph G_n of layer n . In Step 4 of the
 584 layered greedy algorithm, there is a call to Prim's algorithm to construct T^n . The
 585 subscript k in e_k^n refers to the k -th iteration of Prim's algorithm *within* Step 4 of the
 586 layered greedy algorithm.

587 Let $k_n^* := \max\{k \mid e_k^n \in \mathcal{E}_{n-1}, \ell = 1, 2, \dots, k \text{ and } 1 \leq k \leq |L_n| - 1\}$, i.e., the last
 588 iteration of Prim's algorithm applied to G_n before an edge that is *not* contained in
 589 G_{n-1} is selected. Since Prim's algorithm is initialized with the root node, $1 \leq k_n^* <$
 590 $|L_n| - 1$ for $n > 1$ (we let $k_1^* = 0$). Furthermore, let

$$591 \quad (6.1) \quad F_n^* = \{e_\ell^n, \ell = 1, 2, \dots, k_n^* + 1\}.$$

592 In other words, F_n^* is the set of edges added by Prim's algorithm applied to G_n up to
 593 and including the first edge that connects a node in L_{n-1} and a node in $L_n \setminus L_{n-1}$,
 594 namely $e_{k_n^*+1}^n \in \mathcal{E}_n$.

595 **PROPOSITION 6.4.** *Suppose G is a locally finite and connected graph (Assump-*
 596 *tion 1) with distinct edge costs (Assumption 5). Then $F_n^* \subseteq T^m$, for $m \geq n$ and*
 597 *$n = 1, 2, \dots$, where F_n^* is defined in (6.1).*

598 *Proof.* Consider an arbitrary $n \geq 1$ and arbitrary $m \geq n$. For $n = 1$, the result
 599 is trivially true, since in this case F_n^* will include only the cheapest edge incident to
 600 the root node, and this edge will be added as the first iterate of each application of
 601 Prim's algorithm. Consider now $n > 1$ and $m \geq n$. We will show that $e_\ell^m = e_\ell^n$
 602 for all $\ell = 1, 2, \dots, k_n^* + 1$, which implies that $F_n^* \subseteq T^m$. We will prove this by

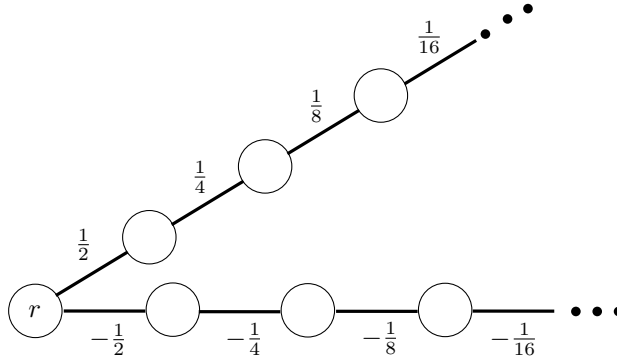


FIG. 5. A graph with some minimum spanning tree edges that do not satisfy the sufficient condition in Corollary 6.5 (see Example 5).

603 mathematical induction on ℓ . The claim is clearly true for $\ell = 1$ since the minimum-
 604 cost edge emanating from node r is the same for all graphs G_m with $m \geq 1$. Adopt
 605 the inductive hypothesis that $e_\ell^m = e_\ell^n$ for all $\ell = 1, 2, \dots, k$ for some $k \leq k_n^*$. Then
 606 Prim's Algorithm, before its $k + 1$ -st iteration, has created trees identical to $F_k :=$
 607 $\{e_\ell^n, \ell = 1, 2, \dots, k\} \subseteq G_{n-1}$ when applied to graphs G_n and G_m for $m \geq n$. Then the
 608 $k + 1$ -st iteration of Prim's algorithm for both graphs finds the same minimum-cost
 609 edge e_{k+1}^n out of F_k since all edges emanating from F_k in G_m are in \mathcal{E}_n for all $m \geq n$,
 610 thus restoring the inductive hypothesis. \square

611 **REMARK 3.** *The distinct edge costs assumption (Assumption 5) is important to*
 612 *the above result as it ensures that different calls to Prim's algorithm do not need to*
 613 *make tie-breaking decisions and potentially select different edges on earlier layers of*
 614 *the graph.* \triangleleft

615 If the graph possesses an MST T^* , we can further demonstrate that all edges of
 616 F_n^* are guaranteed to be in the set \mathcal{E}^* of edges of T^* .

617 **COROLLARY 6.5.** *Suppose G is a locally finite and connected graph (Assump-*
 618 *tion 1) with distinct edge costs (Assumption 5) and (a unique) MST $T^* = \{\mathcal{V}, \mathcal{E}^*\}$*
 619 *exists. Then $F_n^* \subseteq \mathcal{E}^*$, $n = 1, 2, \dots$, where F_n^* is defined in (6.1).*

620 *Proof.* Let $T^*(n)$ be the smallest (connected, finite) subtree of T^* that contains
 621 all nodes of layer n , and let $G^*(n)$ be the subgraph of G spanned by $T^*(n)$. It is
 622 easy to show (e.g., by contradiction) that $T^*(n)$ is an MST of $G^*(n)$; moreover, it is a
 623 unique MST due to Assumption 5. Applying Prim's algorithm to $G^*(n)$ starting with
 624 the root node, we will generate F_n^* on the way to generating $T^*(n)$, since $G_n \subseteq G^*(n)$.
 625 Hence $F_n^* \subseteq T^*(n) \subseteq \mathcal{E}^*$. \square

626 Corollary 6.5 provides a basic sufficient condition for an edge e to lie in an MST
 627 under appropriate assumptions: if $e \in F_n^*$ for some n , then e is an edge of an MST.
 628 This condition can be readily verified by running Prim's algorithm until it first reaches
 629 outside the smallest layer that contains e and checking whether e has been added to
 630 T^n by this point. Therefore, we have a partial answer to question (Q3).

631 It is important to stress that this condition is only sufficient. If an edge e does not
 632 lie in F_n^* for any n , this does not mean that e is not an edge of any MST. A simple
 633 example illustrates this point.




FIG. 6. An example that satisfies condition (6.2) in Corollary 6.6.

634 EXAMPLE 5. Consider the graph in Figure 5 and let the node in the bottom left
 635 corner be the root node. Clearly, this graph satisfies the assumptions of Corollary 6.5,
 636 and its single minimum spanning tree is the whole graph itself. In the n -th iteration
 637 of the layered greedy algorithm, Prim’s algorithm selects every available negative-cost
 638 edge before selecting any positive-cost edge. This implies that edge count K_n^* is reached
 639 before a single positive-cost edge is reached. This implies that the positive-cost edges
 640 do not lie in F_n^* , even though they are part of the minimum spanning tree. This
 641 implies that the sufficient condition in Corollary 6.5 cannot identify the positive-cost
 642 edges of this graph as belonging to the minimum spanning tree. \triangleleft

643 In the next set of results, we build on Proposition 6.4 and Corollary 6.5 to identify
 644 scenarios where we can tell that an entire iterate T^n of the layered greedy algorithm
 645 lies in T^* .

646 COROLLARY 6.6. Suppose G is a locally finite and connected graph (Assump-
 647 tion 1) with distinct edge costs (Assumption 5) and (a unique) MST $T^* = \{\mathcal{V}, \mathcal{E}^*\}$
 648 exists. Suppose

649 (6.2)
$$\min_{e \in \mathcal{K}_{\bar{n}}} c(e) > \max_{e \in \mathcal{E}_{\bar{n}}} c(e)$$

650 for some $\bar{n} > 1$, where $\mathcal{K}_{\bar{n}} := \{\{i, j\} : i \in L_{\bar{n}} \text{ and } j \in L_{\bar{n}+1} \setminus L_{\bar{n}}\}$. Then all edges of
 651 layered greedy iterate $T^{\bar{n}}$ lie in every subsequent iterate T^n , $n \geq \bar{n}$, and therefore, $T^{\bar{n}}$
 652 is contained in T^* .

653 *Proof.* Observe that (6.2) ensures that all edges of $T^{\bar{n}}$ lie in $F_{\bar{n}+1}^*$, since this
 654 condition implies that, when Prim’s algorithm is applied to layer $\bar{n} + 1$ and beyond,
 655 all nodes *within* layer \bar{n} get spanned before any node outside of this layer is reached.
 656 The rest of the argument follows by Proposition 6.4 and Corollary 6.5. \square

657 It is straightforward to see that condition (6.2) fails in the graph in Figure 5. The
 658 next example provides a case where condition (6.2) holds.

659 EXAMPLE 6. To illustrate condition (6.2), consider the graph in Figure 6 that is
 660 adapted from Figure 16.7 in [3]. We can see that condition (6.2) holds for $\bar{n} = 2$ since
 661 $\min\{45, 50, 60\} > \max\{35, 40, 25, 10, 20, 15, 30\}$. Thus, the layered greedy algorithm
 662 locks in the edges of T^2 starting with iteration 3. In this case, these edges have costs
 663 35, 10, 20, and 15, and they are guaranteed to be in T^* independently of the structure
 664 and costs of G after layer 3 (aside from ensuring that assumptions of Corollary 6.6
 665 hold). \triangleleft

666 Condition (6.2) can be interpreted as follows: the edges in $\mathcal{K}_{\bar{n}}$ create a “mountain




FIG. 7. Graph for Example 7 illustrating the notation defined in Corollary 6.7. Here, $n_1 = 1$ and $n_2 = 3$. The edges in K_{n_2} are dashed. The edges in $\mathcal{E}(n_1, n_2)$ are in bold. It is easy to see that (6.3) is satisfied for n_1 and n_2 .

667 range” or a “ridge” of costs, while all edges within the subgraph $G_{\bar{n}}$ form a cost
 668 “valley”; as a result, all the nodes in the valley should be spanned before the MST
 669 ventures across the ridge.

670 Note that in a graph with positive edge costs, this condition cannot hold for all n ,
 671 or even for an infinite subsequence of n , and satisfy the other assumptions imposed on
 672 our graphs. Indeed, for (6.2) to hold on an infinite subsequence n_k , $k = 1, 2, \dots$, we
 673 must have a subsequence of edges with costs that are increasing. But this condition
 674 violates Assumption 2, which may be needed to establish existence of an MST, since
 675 it requires the sequence of edge costs to converge to 0 for them to be summable.

676 Luckily, we can provide a modification of condition (6.2) that can hold on a sub-
 677 sequence of layers without contradicting Assumption 2 while providing a workable
 678 approach to identifying early edges in T^* . The new condition is discussed in Corol-
 679 lary 6.7 and illustrated in Figure 7.

680 COROLLARY 6.7. Suppose G is a locally finite and connected graph (Assump-
 681 tion 1) with distinct edge costs (Assumption 5), and (a unique) MST $T^* = \{\mathcal{V}, \mathcal{E}^*\}$
 682 exists. Suppose further that there is an increasing sequence n_k , $k = 1, 2, \dots$, with
 683 $n_1 > 1$, that satisfies the following conditions:

$$684 \quad (6.3) \quad \min_{e \in K_{n_1}} c(e) > \max_{e \in \mathcal{E}_{n_1}} c(e), \text{ and } \min_{e \in K_{n_k}} c(e) > \max_{e \in \mathcal{E}(n_{k-1}, n_k)} c(e) \text{ for } k > 1,$$

685 where $\mathcal{E}(n, m) = \mathcal{E}_m \setminus (\mathcal{E}_n \cup K_n)$ for $n < m$, i.e., it is the set of all edges of G with both
 686 endpoints in layer m , but outside layer n (thus extending notation $\mathcal{E}_m = \mathcal{E}(0, m)$).
 687 Furthermore, assume that whenever the set $L_{n_k} \setminus L_{n_{k-1}}$ contains more than one
 688 node, this node set is connected in the graph induced by $\mathcal{E}(n_{k-1}, n_k)$. Then, for all
 689 $k = 1, 2, \dots$, all edges of layered greedy iterate T^{n_k} lie in every subsequent iterate T^n ,
 690 $n \geq n_k$, and therefore, T^{n_k} is contained in T^* .

691 EXAMPLE 7. Consider the graph in Figure 7. We have $n_1 = 1$, since

$$692 \quad \max \left\{ \frac{1}{2}, 1 \right\} < \min \left\{ 1 + \frac{1}{2}, 1 + \frac{1}{4}, 1 + \frac{1}{8} \right\},$$

693 and T^1 consists of the two edges emanating from the root node. It is easy to see that
 694 Prim’s algorithm applied to any G_n with $n \geq 1$ in this example will begin by adding
 695 these two edges, which are therefore locked in.

696 Furthermore, $n_2 = 3$ satisfies (6.3), since the most expensive of the bold edges has
 697 cost $\frac{1}{4}$, and the cheapest of the dashed edges has cost $\frac{1}{4} + \frac{1}{64}$. T^3 consists of edges with
 698 costs $\frac{1}{2}$, 1 , $1 + \frac{1}{8}$, $\frac{1}{128}$, $\frac{1}{256}$, $\frac{1}{512}$, $\frac{1}{32}$, and $\frac{1}{16}$ (listed here in the order they are added
 699 by Prim's algorithm). The application of Prim's algorithm to construct T^4 will also
 700 begin by adding these edges.

701 Notice, however, that T^2 contains the edge with cost $\frac{1}{8}$, which is not included in
 702 the subsequent iterates, illustrating that the result in Corollary 6.7 is only guaranteed
 703 to hold on the specified subsequence.

704 *Proof of Corollary 6.7.* We will prove, by induction on k , that $T^{n_k} \subset F_{n_{k+1}}^*$ for
 705 $k = 1, 2, \dots$. For $k = 1$, condition (6.3) coincides with (6.2), and this conclusion
 706 follows by Corollary 6.6. For $k > 1$, let us adopt the inductive hypothesis that
 707 $T^{n_{(k-1)}} \subset F_{n_{(k-1)+1}}^*$, and show that $T^{n_k} \subset F_{n_{k+1}}^*$.

708 If the set $L_{n_k} \setminus L_{n_{(k-1)}}$ consists of a single node (say, v), the claim is trivially true,
 709 since then $n_k = n_{(k-1)} + 1$, T^{n_k} consists of $T^{n_{(k-1)}}$ combined with the cheapest edge
 710 connecting $L_{n_{(k-1)}}$ with v , and $F_{n_{k+1}}^*$ consist of T^{n_k} combined with the cheapest edge
 711 connecting v with a node in $L_{n_{k+1}} \setminus L_{n_k}$. We will therefore consider the case when
 712 $L_{n_k} \setminus L_{n_{(k-1)}}$ contains multiple nodes.

713 As before, let e_m^n be the edge added by the m -th iteration of Prim's algorithm
 714 applied to G_n . To prove our claim, we need to show that, for $m = 1, \dots, |L_{n_k}| - 1$,

715 (6.4)
$$e_m^{n_k+1} = e_m^{n_k}.$$

716 By the inductive hypothesis, (6.4) is true for all $m \leq |L_{n_{(k-1)}}| - 1$ (while both Prim's
 717 algorithms are constructing $T^{n_{(k-1)}}$) and for $m = |L_{n_{(k-1)}}|$ (when they both add the
 718 cheapest edge from $\mathcal{K}_{n_{(k-1)}}$ to reach $L_{n_{(k-1)+1}}$, thus completing $F_{n_{(k-1)+1}}^*$).

719 We now construct an induction on ℓ where we suppose (6.4) is true for all $m \leq \ell$,
 720 where $|L_{n_{(k-1)}}| \leq \ell < |L_{n_k}| - 1$, and consider the sets of edges each algorithm chooses
 721 from in iteration $\ell + 1$. During the first ℓ iterations, the algorithms have spanned,
 722 using the same edges, all of $L_{n_{(k-1)}}$ and a strict subset \mathcal{V}_ℓ of $L_{n_k} \setminus L_{n_{(k-1)}}$. Let
 723 $\mathcal{V}' = (L_{n_k} \setminus L_{n_{(k-1)}}) \setminus \mathcal{V}_\ell$ — these are precisely the nodes of L_{n_k} that have not yet
 724 been spanned.

725 We now prove the inductive step in iteration $\ell + 1$. In that iteration, Prim's
 726 algorithm applied to G_{n_k} is comparing the costs of edges in $\mathcal{K}_{n_{(k-1)}}$ incident to nodes
 727 in \mathcal{V}' and edges connecting nodes in \mathcal{V}_ℓ to nodes in \mathcal{V}' , while the algorithm applied
 728 to $G_{n_{k+1}}$ is comparing the costs of all the aforementioned edges as well as any edges
 729 in \mathcal{K}_{n_k} incident to nodes in \mathcal{V}_ℓ . Due to the assumption that node set $L_{n_k} \setminus L_{n_{(k-1)}}$ is
 730 connected in $\mathcal{E}(n_{(k-1)}, n_k)$, at least one of the edges from this edge set is considered
 731 in the cost comparison by both algorithms, and by (6.3), it will be cheaper than any
 732 edge in \mathcal{K}_{n_k} . Therefore, Prim's algorithm applied to $G_{n_{k+1}}$ will not choose an edge
 733 from \mathcal{K}_{n_k} until all nodes in L_{n_k} have been spanned, i.e., until it constructs the MST
 734 T^{n_k} . This establishes (6.4) for $\ell + 1$ and completes our induction on ℓ , which in turn
 735 closes the outer induction on k .

736 The rest of the argument follows by Proposition 6.4 and Corollary 6.5. \square

737 This last corollary shows that the MST T^* can be constructed by building the
 738 smaller finite trees T^{n_k} where later iterations do not add or remove edges from the
 739 layer of G spanned by the T^{n_k} uncovered so far.

740 It is worth noting that assumptions of Corollary 6.7 and Assumption 2 can be met
 741 simultaneously in graphs with positive costs. Roughly speaking, condition (6.3) only
 742 requires that, occasionally, costs of edges connecting to a new layer form a "ridge," but

743 only relative to the costs of edges in the previous valley. However, the heights of the
 744 subsequent ridges K_{n_k} can get smaller as long as the subsequent valleys $\mathcal{E}(n_{(k-1)}, n_k)$
 745 also get more shallow.

746 Corollaries 6.6 and 6.7 provide additional partial answers to question (Q3).

747 **6.3. An application: High-speed information channels.** In this subsection,
 748 we illustrate how the results in this section can be used to solve a minimum
 749 spanning tree problem on an infinite graph that arises from an application. The infinite
 750 graph models an underlying indefinite but large finite graph whose nodes we
 751 expect to ultimately connect via a spanning tree of telecommunication links.

752 Suppose in particular that a telecommunications company is building high-speed
 753 information channels (e.g., via laying fiber-optic cables) to connect a large number
 754 of locations to a single service provider at minimum cost. The collection of these
 755 locations is modeled as countably infinite since the goal is to connect discrete locations
 756 over a long but uncertain life of the project. For more discussion of using infinite
 757 graphs to study infinite-horizon optimization problems see [21]. The cost of an edge
 758 $\{i, j\}$ is the cost of building an information channel that directly connects location i
 759 and location j .

760 We view the layers of the graph as nodes reached by edges over time. The first
 761 layer consists of locations that can be connected to the root node (the service provider
 762 location) in a certain interval of time, say, 1 year. The second layer consists of locations
 763 that can be connected to the root node (via a node in layer 1) in two time periods,
 764 say, 2 years. Under this time interpretation of layering, it follows that each node has
 765 finite degree, since in finite time a location can only be connected to finitely many
 766 other locations. This supports Assumption 1. As for Assumption 2, it is natural to
 767 assume that future costs are discounted by a discount factor that assures summable
 768 costs. These two assumptions then assure that the layered greedy algorithm will find
 769 a sequence of spanning-tree iterates that converge in value to optimality.

770 The nature of the layered greedy algorithm, however, is that the edges in the
 771 tree iterates will shift around, as we saw in Example 1. For an application like
 772 laying fiber-optic cable, such “shifting around” can lead to very expensive reworking
 773 requiring removal of previously added edges. We would prefer to be able to apply
 774 a rolling horizon approach to this problem. In particular, we would like to be able
 775 to finalize our decisions of which potential edges within a few initial layers will and
 776 will not be built based on whether they are included in T_{n_1} for some small n_1 (and
 777 proceed to lay cable along the chosen edges during the first few years of construction);
 778 then finalize the decisions regarding the edges in the next few layers based on T_{n_2} for
 779 some $n_2 > n_1$, etc., without sacrificing optimality of the overall spanning tree that is
 780 being constructed.

781 If we assume more about the underlying graph, we can get stronger convergence
 782 results. These assumptions are in fact quite natural in our setting. The condition
 783 of distinct edge costs (Assumption 5) is easy to guarantee since it is unlikely that
 784 two projects to connect two different pairs of locations have exactly the same costs.
 785 The recursive ridges and valleys condition (6.3) is natural in this application, with
 786 “valleys” and “mountain ranges” representing either the actual topography of the
 787 area or the difference in difficulty and costs of laying cable with and without pre-
 788 existing underground conduits. We may assume the costs are summable if we take
 789 time discounting into consideration, so even though “far off” mountains may be high,
 790 their costs will be sufficiently discounted. Finally, the connectedness assumption of
 791 Corollary 6.7 is natural if the population of the valleys is dense enough to allow it to




FIG. 8. A graph with an MST that fails the FC property.

792 be connected by cheap local infrastructure. Accordingly, we can apply the result of
 793 Corollary 6.7 ensuring that we can construct the MST recursively in finite subtrees
 794 whose edges become stable at finite intervals (the associated sequence $\{n_k\}$) without
 795 edges entering or leaving the MST.

796 **7. Conclusion.** In this paper, we gave an algorithm that yields convergence
 797 in objective value for a broad class of infinite graphs (locally finite and connected,
 798 with absolutely summable edge costs) that works as long as an MST is known to
 799 exist (Theorem 3.3). We offer the combination of the FC property on the graph
 800 and absolute summability of the costs as a sufficient condition for existence, but
 801 acknowledge that these are not necessary conditions. Indeed, consider the graph in
 802 Figure 8. It satisfies the properties of absolutely summable and distinct edge costs
 803 but fails the FC property. Nonetheless, an MST exists, as indicated in dashed (green)
 804 edges. An interesting open question is whether there is a meaningful characterization
 805 of when an MST exists in a locally finite and connected graph that is weaker than
 806 the FC property, or substantially different from it.

807 In this paper, we also showed convergence of the layered greedy iterates in the
 808 scenario where there exists a unique MST (Theorem 6.1). Unlike in many other
 809 optimization problems, where the uniqueness of the optimal solution is hard to verify,
 810 this problem has the simple sufficient condition of unique edge costs. We also showed
 811 in Example 4 that if there is more than one MST then the iterates of the layered greedy
 812 algorithm may fail to converge to an MST. The convergence issue arose because of an
 813 “unfortunate” selection of edges of equal cost as the algorithm proceeds. We believe
 814 that this “selection” issue could potentially be resolved, using an approach similar in
 815 spirit to [25]. We will leave this for future work.

816 Finally, we explored a verifiable sufficient conditions that allow us to confirm
 817 whether an iterate of the layered greedy algorithm has “locked in,” i.e., verify that all
 818 its edges will be contained in all of the future iterates (and thus the optimal MST if
 819 it exists).

820 **Appendix A. Appendix: Proof of Proposition 4.1.**

821 We start with the following preliminary lemma.

822 **LEMMA A.1.** *If a locally finite and connected graph G contains no bi-rays, then*
 823 *every pair of rays must have infinitely many nodes in common.*

824 *Proof.* Let $r^1 := (i_1, i_2, \dots)$ and $r^2 := (j_1, j_2, \dots)$ be two rays in the graph, and
 825 suppose they have at most finitely many nodes in common. If they have no nodes
 826 in common, then a bi-ray is produced by connecting nodes i_1 and j_1 and, if needed,
 827 deleting a finite number of repeated edges or cycles that may be created by the added

828 path.

829 Otherwise, we can construct a bi-ray as follows. Suppose $k = j_n$ is the last node
 830 r^1 and r^2 have in common, where we determine the “last” node according to the
 831 ordering or nodes in r^2 . Furthermore, let m be such that $i_m = j_n = k$. Then the rays
 832 $\bar{r}^1 := (k = i_m, i_{m+1}, \dots)$ and $\bar{r}^2 := (k = j_n, j_{n+1}, \dots)$ are distinct except for node k .
 833 Indeed, although the ray \bar{r}^1 may contain nodes j_ℓ for $\ell < m$, they are now excluded
 834 from the ray \bar{r}^2 . The union of rays \bar{r}^1 and \bar{r}^2 is a bi-ray, a contradiction.

835 LEMMA A.2. *The collection of all paths and rays in a locally finite and connected*
 836 *graph that contains no bi-rays is compact in the product discrete topology.*

837 *Proof.* Observe that a subgraph is a path or a ray if and only if it is a connected
 838 and acyclic subgraph where each node has degree at most two. (Bi-rays also have these
 839 properties, but we are assuming that our graph has no bi-rays.) Let P^k , $k = 1, 2, \dots$,
 840 be a sequence of paths and rays that converges in the product discrete topology to
 841 some subgraph P . We claim that P has no cycles, is connected, and each node in P
 842 has degree at most 2, i.e., P is either a path or a ray.

843 The proof that P is acyclic follows the same logic as claim (ii) in Lemma 5.4 using
 844 the lock-in property of convergence.

845 Next, suppose P has a node of degree 3 or greater. Again, by lock-in, this implies
 846 that infinitely many of the P^k also have a node of degree 3 or greater, contradicting
 847 the fact they are paths or rays.

848 Finally, we establish by contradiction that P is connected. Suppose there are two
 849 nodes $i, j \in P$ that are not connected in P . Since these two nodes are in P , P contains
 850 at least one edge incident to i and at least one edge incident to j . This means that,
 851 for sufficiently large k , each P^k contains those edges and thus contains both nodes i
 852 and j ; we can pass to a subsequence to make this claim for all k . Let P_{ij}^k be the path
 853 that connects i and j in P^k .

854 Let $i_1^k \in I(i)$ be such that $\{i, i_1^k\} \in P_{ij}^k$. By the pigeonhole principle, one of these
 855 edges locks in, so that for some $i_1 \in I(i)$, $\{i, i_1\} \in P_{ij}^k$ for k sufficiently large, and
 856 thus $\{i, i_1\} \in P$. Note that $i_1 \neq j$ by our assumption. Let us pass to a subsequence
 857 so that $\{i, i_1\} \in P_{ij}^k$ for all k .

858 We continue following each of the paths P_{ij}^k from i_1 towards j . Consider nodes
 859 $i_2^k \in I(i_1)$ such that $i_2^k \neq i$ and $\{i_1, i_2^k\} \in P_{ij}^k$. Following the same logic, one of these
 860 edges, denoted $\{i_1, i_2\}$, is contained in all paths P_{ij}^k for sufficiently large k , and thus
 861 is contained in P . Note that $i_2 \neq j$ and, since P is acyclic, $i_2 \neq i$.

862 We will repeat the above process iteratively. At each step, we will continue fol-
 863 lowing the paths P_{ij}^k towards j from the most-recently identified node i_m , and adding
 864 a node i_{m+1} such that the edge sequence $(\{i, i_1\}, \{i_1, i_2\}, \dots, \{i_{m-1}, i_m\}, \{i_m, i_{m+1}\})$
 865 is in P_{ij}^k for all (sufficiently large) k , and thus is in P . Since i_{m+1} is different from i_m
 866 by construction, and from every other identified node since P_{ij}^k is acyclic, this process
 867 will create a ray $R_i = (\{i, i_1\}, \{i_1, i_2\}, \dots) \subset P$ that does not include node j .

868 Using the same process starting from j , we can create a ray

$$869 R_j = (\{j, j_1\}, \{j_1, j_2\}, \dots) \subset P$$

870 that does not include node i . Moreover, this ray has no nodes in common with R_i ,
 871 since otherwise there is a path connecting i and j in P . This, however, contradicts
 872 Lemma A.1 in a graph with no bi-rays, thus establishing that P is connected. \square

873 *Proof of Proposition 4.1.* Suppose G is a locally finite and connected graph with
 874 no bi-rays. By way of contradiction, suppose there exists an edge $\{i, j\} \in \mathcal{E}$ that is

875 contained in infinitely many cycles. Deleting the edge from those cycles, we conclude
 876 that there are infinitely many distinct paths P_{ij}^n , $n = 1, 2, \dots$, connecting i and j .

877 Observe that there must be an infinite subsequence $P_{ij}^{n_k}$, $k = 1, 2, \dots$, such that
 878 $P_{ij}^{n_{k+1}}$ contains strictly more edges than $P_{ij}^{n_k}$, for all k . Suppose otherwise, that
 879 there is a maximum number N of edges in all paths between nodes i and j . By
 880 local finiteness, there are finitely many potential paths of length N leaving node i .
 881 However, we have supposed there are infinitely many paths of length N leaving node
 882 i and reaching node j . Hence, such a sequence $P_{ij}^{n_k}$, $k = 1, 2, \dots$, exists.

883 Let N_k , $k = 1, 2, \dots$, denote the increasing sequence of cardinalities of the edge
 884 sets of paths $P_{ij}^{n_k}$, and let m_k be the $\lfloor N_k/2 \rfloor$ -th node in the path $P_{ij}^{n_k}$. Break each
 885 $P_{ij}^{n_k}$ into two subpaths, $P_i^{n_k}$ and $P_j^{n_k}$, where $P_i^{n_k}$ connects node i and node m_k , and
 886 $P_j^{n_k}$ connects node j and node m_k ; i.e., $P_i^{n_k}$ and $P_j^{n_k}$ have only node m_k in common.
 887 Passing to subsequences if necessary and using Lemma A.2, sequences $P_i^{n_k}$ and $P_j^{n_k}$
 888 each have a limit P_i and P_j , respectively, that are either paths or rays. Moreover, by
 889 the construction of $P_i^{n_k}$ and $P_j^{n_k}$, they cannot converge to limits with finitely many
 890 nodes, and so P_i and P_j must be rays.

891 Our contradiction comes from the properties of rays P_i and P_j . We argue that P_i
 892 and P_j have at most one node in common. Suppose otherwise that P_i and P_j have at
 893 least two nodes in common, say, u and v . Then P_i contains a finite path p_i between
 894 u and v and P_j contains a finite path p_j between u and v . There are two cases to
 895 consider. The first is where p_i and p_j share an edge. In this case, by the lock-in
 896 property, $P_i^{n_k}$ and $P_j^{n_k}$ both contain that edge for large enough k , contradicting the
 897 fact that $P_i^{n_k}$ and $P_j^{n_k}$ do not have any edges in common by construction.

898 On the other hand, if p_i and p_j do not share edges, then their union contains a
 899 cycle C in $P_i \cup P_j$. Recall that $P_{ij}^{n_k}$ is equal to the union of $P_i^{n_k}$ and $P_j^{n_k}$, and since
 900 $P_i^{n_k}$ converges to P_i and $P_j^{n_k}$ converges to P_j , we must have that $P_{ij}^{n_k}$ converges to
 901 $P_i \cup P_j$. This implies that infinitely many elements in the sequence $P_{ij}^{n_k}$ contain the
 902 cycle C by the lock-in property. This contradicts the fact that each $P_{ij}^{n_k}$ is a path.

903 This establishes that the rays P_i and P_j intersect in at most one node. On the
 904 other hand, since P_i and P_j are rays in a graph with no bi-rays, by Lemma A.1 they
 905 must have infinitely many nodes in common. We have arrived at a contradiction, and
 906 thus every edge of G is contained in at most finitely many cycles. \square

907

REFERENCES

- 908 [1] L. ADDARIO-BERRY, N. BROUTIN, C. GOLDSCHMIDT, AND G. MIERMONT, *The scaling limit of*
 909 *the minimum spanning tree of the complete graph*, The Annals of Probability, 45 (2017),
 910 pp. 3075–3144.
 911 [2] R. AHARONI, E. BERGER, A. GEORGAKOPOULOS, A. PERLSTEIN, AND P. SPRÜSSEL, *The max-*
 912 *flow min-cut theorem for countable networks*, Journal of Combinatorial Theory Series B,
 913 101 (2011), pp. 1–17.
 914 [3] R. K. AHUJA, T. L. MAGNANTI, AND J. B. ORLIN, *Network Flows: Theory, Algorithms, and*
 915 *Applications*, Prentice Hall, 1993.
 916 [4] D. ALDOUS AND J. M. STEELE, *Asymptotics for Euclidean minimal spanning trees on random*
 917 *points*, Probability Theory and Related Fields, 92 (1992), pp. 247–258.
 918 [5] K. S. ALEXANDER, *Percolation and minimal spanning forests in infinite graphs*, The Annals of
 919 Probability, (1995), pp. 87–104.
 920 [6] C. D. ALIPRANTIS AND K. C. BORDER, *Infinite Dimensional Analysis: A Hitchhiker’s Guide*,
 921 Springer, 3rd ed., 2006.
 922 [7] E. J. ANDERSON AND A. B. PHILPOTT, *A continuous-time network simplex algorithm*, Networks,
 923 19 (1989), pp. 395–425.
 924 [8] H. BRUHN, R. DIESTEL, M. KRIESELL, R. PENDAVINGH, AND P. WOLLAN, *Axioms for infinite*
 925 *matroids*, Advances in Mathematics, 239 (2013), pp. 18–46.

- 926 [9] N. CHRISTOFIDES, *Worst-case analysis of a new heuristic for the travelling salesman problem*,
927 tech. report, Carnegie-Mellon University Technical Report, 1976.
- 928 [10] R. DIESTEL, *Graph Theory*, Springer, 4th ed., 2010.
- 929 [11] M. A. DJAUHARI AND S. L. GAN, *Optimality problem of network topology in stock market*
930 *analysis*, Physica A: Statistical Mechanics and Its Applications, 419 (2015), pp. 108–114.
- 931 [12] A. GHATE, *Duality in countably infinite monotropic programs*, SIAM Opt., 27 (2017), pp. 2010–
932 2033.
- 933 [13] R. L. GRAHAM AND P. HELL, *On the history of the minimum spanning tree problem*, Annals
934 of the History of Computing, 7 (1985), pp. 43–57.
- 935 [14] D. GRANOT AND G. HUBERMAN, *Minimum cost spanning tree games*, Mathematical Program-
936 ming, 21 (1981), pp. 1–18.
- 937 [15] V. KLEE, *The greedy algorithm for finitary and cofinitary matroids*, in Combinatorics: Pro-
938 ceedings of Symposia in Pure Mathematics, T. S. Motzkin, ed., 1971, pp. 137–152.
- 939 [16] R. LYONS, Y. PERES, AND O. SCHRAMM, *Minimal spanning forests*, The Annals of Probability,
940 34 (2006), pp. 1665–1692.
- 941 [17] J. MUNKRES, *Topology*, Prentice Hall, 2000.
- 942 [18] S. NOUROLLAHI AND A. GHATE, *Duality in convex minimum cost flow problems on infinite*
943 *networks and hypernetworks*, Networks, 70 (2017), pp. 98–115.
- 944 [19] S. NOUROLLAHI AND A. GHATE, *Inverse optimization in minimum cost flow problems on count-*
945 *ably infinite networks*, Networks, 73 (2019), pp. 292–305.
- 946 [20] A. PAUL, D. FREUND, A. FERBER, D. B. SHMOYS, AND D. P. WILLIAMSON, *Budgeted prize-*
947 *collecting traveling salesman and minimum spanning tree problems*, Mathematics of Op-
948 erations Research, 45 (2019), pp. 576–590.
- 949 [21] H. E. ROMELJN, D. SHARMA, AND R. L. SMITH, *Extreme point characterizations for infinite*
950 *network flow problems*, Networks, 48 (2006), pp. 209–22.
- 951 [22] C. T. RYAN AND R. L. SMITH, *A greedy algorithm for finding maximum spanning trees in*
952 *infinite graphs*, Operations Research Letters, 50 (2022), pp. 655–659.
- 953 [23] C. T. RYAN, R. L. SMITH, AND M. A. EPELMAN, *A simplex method for uncapacitated pure-*
954 *supply infinite network flow problems*, SIAM Journal on Optimization, 28 (2018), pp. 2022–
955 2048.
- 956 [24] S. M. RYAN, J. C. BEAN, AND R. L. SMITH, *A tie-breaking rule for discrete infinite horizon*
957 *optimization*, Operations Research, 40 (1992), pp. S117–S126.
- 958 [25] I. E. SCHOCHETMAN AND R. L. SMITH, *Convergence of selections with applications in optimiza-*
959 *tion*, Journal of Mathematical Analysis and Applications, 155 (1991), pp. 278–292.
- 960 [26] T. C. SHARKEY AND H. E. ROMELJN, *A simplex algorithm for minimum-cost network-flow*
961 *problems in infinite networks*, Networks, 52 (2008), pp. 14–31.
- 962 [27] K. J. SUPOWIT, D. A. PLAISTED, AND E. M. REINGOLD, *Heuristics for weighted perfect match-*
963 *ing*, in ACM STOC Symposium on Theory of Computing, 1980, pp. 398–419.
- 964 [28] S. ZHANG, H. TONG, J. XU, AND R. MACIEJEWSKI, *Graph convolutional networks: A compre-*
965 *hensive review*, Computational Social Networks, 6 (2019), p. 11.