

A simplex method for uncapacitated pure-supply and pure-demand infinite network flow problems

Christopher Thomas Ryan* Robert L Smith[†] Marina Epelman[‡]

July 6, 2017

Abstract

We provide a simplex algorithm for a structured class of uncapacitated countably-infinite network flow problems. Previous efforts required explicit capacities on arcs with uniformity properties that facilitate duality arguments. By contrast, this paper takes a “primal” approach by devising a simplex method that provably converges to optimal value using arguments based on convergence of spanning trees and nonnegativity of reduced costs. This allows for removal of explicit capacity bounds. The method also converges, on a subsequence, to an extremal optimal solution. Our method is tailored to our problem setting — acyclic networks with nodes of only nonnegative supplies (or, alternatively, only demands). The necessary structure can be found in a variety of applied settings not amenable to existing methods including nonstationary infinite-horizon dynamic programming. A finite implementation of our simplex algorithm is provided for the infinite horizon dynamic lot sizing problem under linear costs.

Keywords: network simplex method, infinite networks, duality

AMS subject classifications: 90C35, 90C49, 05C63

1 Introduction

Network flow problems have a storied history of both theory and applications. Network problems have long driven the theory of linear optimization and seen countless uses in industrial settings (for a thorough survey see [2]). However, network flow problems over *infinite* graphs have seen only scant development over the last twenty years. This is despite their potential to drive the development of infinite-dimensional linear programming and find application in settings where infinitely large graphs are a natural modeling consequence, including infinite-horizon planning problems.

Our work draws on previous development in the infinite-dimensional linear optimization and infinite graph literatures. We work with directed graphs with countably-many nodes where each node has finite in- and out-degree. This is a special case of countably-infinite linear programming (CILP) studied by Smith and various co-authors [9, 11, 16, 17]. Ours is one of only two studies to handle degeneracy in CILPs, the other being the pioneering work of [21], which does so in the capacitated network flow setting. Our work also relates to the theory of infinite graphs [8]. To our

*Booth School of Business, University of Chicago, E-mail: chris.ryan@chicagobooth.edu

[†]Industrial and Operations Engineering, University of Michigan, E-mail: rlsmith@umich.edu

[‡]Industrial and Operations Engineering, University of Michigan, E-mail: mepelman@umich.edu

knowledge, there are no explorations of the simplex method on infinite graphs in that literature, where optimization is typically not the main focus (see [1, 15] for notable exceptions).

Previous work has examined the extremal structure of optimal solutions to infinite network flow problems [18] including an algorithm that applies to problems that obey certain regularity conditions on their infinite-dimensional duals [21]. This regularity is ensured by uniform capacity bounds on arcs ([21, Proposition 2.5]). Our goal is to devise a simplex method with desirable convergence properties that does not rely on duality arguments and removes the need for capacity restrictions. Somewhat paradoxically, and in contrast to the finite setting, uncapacitated problems can present even more issues than capacitated problems in the infinite setting. The lack of capacities can lead to unbounded flows on arcs that is another source of “infinity” that requires subtle handling. Hence, our analysis complements that of [21] in a new and challenging subclass of network flow problems.

A simplex method pivots between adjacent extreme points of the feasible region, never worsening in objective value. This method is well-understood in the finite-dimensional setting, but more care is needed in the infinite-dimensional setting. For general CILPs, it may be that extreme points cannot be characterized as basic feasible solutions [9], thus making it a challenge to algebraically define the pivot operation. Some existing algorithms explore nonadjacent extreme points and thus fail to be a classical simplex method [11]. Examples in the literature show that a sequence of feasible solutions that is strictly improving in objective value may nonetheless fail to converge to the optimal value [11]. However, simplex methods for special CILPs arising from Markov decision processes (MDP) have been introduced [10, 12].

The goal of this paper is to devise a simplex method for infinite network problems that has convergence guarantees without relying on uniform capacity bounds. Our proposed algorithm produces a sequence of monotone-improving adjacent extreme points that converges in value to the optimum, and converges to an optimal extreme point on a subsequence. Our method is “primal,” avoiding dual arguments based on *transversality* or *nondegeneracy* that are strong theoretical conditions that underpin simplex methods in the existing literature (see, for instance, [10–12, 21]).

Degeneracy is a key challenge here, since network flow problems tend to be highly degenerate, unlike the MDP settings studied in the literature. A convergent simplex algorithm must guarantee absence of cycling between feasible bases, an issue resolved by various methods in the finite-dimensional setting. We develop an anti-cycling pivoting rule that has no direct counterpart in the finite setting. Indeed, pure supply networks themselves have no direct finite counterpart since in finite networks, flow must originate and terminate at nodes within the graph. In the infinite setting, flow can always be sent “to infinity” rather than satisfy demand at any node. In this way, our study highlights several distinct features of the infinite network flow problems not shared by the finite case.

Another unexpected outcome of our approach is a proof of strong duality as a consequence of our simplex method. Although standard in the finite-dimensional setting (see for instance [5, Chapter 4]), using a simplex-like method to prove duality in the infinite setting is relatively uncommon (see [14] for an exception in the case of separated continuous linear programs) and, to our knowledge, not leveraged in CILPs.

Our development is specialized to acyclic pure supply (or pure demand) networks with geometrically decaying costs and uniformly bounded supplies. All nodes are either supply or transshipment nodes (pure demand is defined analogously). Although our setting is highly structured, it nonetheless captures a wide class of potential applications of infinite network flow problems. This includes infinite-horizon nonstationary dynamic programs and dynamic lot sizing problems with linear pro-

duction and holding costs. In the latter case, we are able to provide a finite implementation of our simplex method, suggesting a path for practical application of our ideas with further investigation. Stepping outside our setting presents numerous analytical challenges. A key feature of our method is that simplex iterates correspond to spanning in-trees rooted at infinity and thus have an appealing structure for analysis. The addition of even a single demand node can break the in-tree structure and disrupt our analysis. We believe there is scope for extending our primal approach to more general problems with additional insights. We leave this for future work.

The general simplex method we present for pure supply (and demand) problems is abstract in the sense that its implementation may require an infinite amount of data and computation in each iteration. This drawback is common to many infinite-dimensional optimization methods, with some notable exceptions [10–12]; the algorithm in [21] is also abstract, with a subclass of problems that allow finite iterations considered. We also derive a finite implementation of our approach when applied to the infinite horizon dynamic lot-sizing problem. This shows a path for future work in this direction.

The rest of the paper is organized as follows. [Section 2](#) introduces our problem and notation used throughout the paper. [Section 3](#) defines basic feasible flows, connects them to spanning trees and circulations, and discusses optimality conditions. [Section 4](#) develops our simplex algorithm using the notion of finite *layers* of nodes that control the set of entering variables (building on the concepts in [21]). Our proof of convergence relies on convergence of spanning trees and nonnegativity of reduced costs in a limiting basic feasible flow. [Section 5](#) explores an application to infinite-horizon nonstationary dynamic programming, and [Section 6](#) explores an extension to dynamic lot sizing. [Section 7](#) develops a strong duality result using the simplex method as the main tool of the proof. [Section 8](#) concludes the paper.

2 Pure-supply network flow problems

2.1 Graph structure

Let $G = (\mathcal{N}, \mathcal{A})$ be a directed graph with countably many nodes $\mathcal{N} = \{1, 2, \dots\}$ and arcs $\mathcal{A} \subseteq \mathcal{N} \times \mathcal{N}$. Let $I(i)$ denote the set of nodes that are tails of arcs entering node i : $I(i) := \{j \in \mathcal{N} : (j, i) \in \mathcal{A}\}$. Similarly, the set of nodes that are heads of arcs leaving i is $O(i) := \{j \in \mathcal{N} : (i, j) \in \mathcal{A}\}$. The *in-degree* and *out-degree* of node i in G are the cardinalities of $I(i)$ and $O(i)$, respectively. A graph is *locally finite* if every node has finite in- and out-degree.

A *finite (undirected) path* in G is a finite sequence of distinct nodes i_1, i_2, \dots, i_n , where $(i_k, i_{k+1}) \in \mathcal{A}$ or $(i_{k+1}, i_k) \in \mathcal{A}$ for $k = 1, \dots, n - 1$. A *path to infinity* is a sequence of distinct nodes i_1, i_2, \dots where $(i_k, i_{k+1}) \in \mathcal{A}$ or $(i_{k+1}, i_k) \in \mathcal{A}$ for $k = 1, 2, \dots$. We typically use P_{ij} to denote a finite path from node i to node j , and $P_{i\infty}$ to denote a path from node i to infinity. Two nodes i and j are *finitely connected* in G if there exists a finite path P_{ij} . Two nodes i and j are *connected at infinity* if G contains two paths to infinity, $P_{i\infty}$ and $P_{j\infty}$, that share no common nodes. Nodes i and j are *connected* if they are either finitely connected or connected at infinity. The graph G is *finitely connected* if all nodes i and j in G are finitely connected.

A *finite 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 path and either $(i_1, i_n) \in \mathcal{A}$ or $(i_n, i_1) \in \mathcal{A}$. An *infinite cycle*, also called a *cycle at infinity*, consists of two paths to infinity from some node i , (i, i_1, i_2, \dots) and (i, j_1, j_2, \dots) , where all intermediate nodes i_k and j_ℓ are distinct.

We also need directed versions of these definitions. A *finite directed path* from i_1 to i_n , denoted $P_{i_1 i_n}^{\rightarrow}$, is a finite path i_1, \dots, i_n where $(i_k, i_{k+1}) \in \mathcal{A}$ for all $k = 1, 2, \dots, n-1$. A *directed path from node i to infinity*, denoted $P_{i\infty}^{\rightarrow}$, is a path to infinity $P_{i\infty}$ where each arc in the path is directed away from node i . A *directed path from infinity to node i* , denoted $P_{i\infty}^{\leftarrow}$, is a path to infinity $P_{i\infty}$ where each arc in the path is directed towards node i . A *directed finite cycle* is a finite cycle that consists of a finite directed path i_1, \dots, i_n and the arc $(i_n, i_1) \in \mathcal{A}$. A *directed cycle at infinity* is a cycle at infinity where both paths to infinity from a given node i are directed, one from infinity to i , and the other from i to infinity. A graph is *acyclic* if it contains no finite or infinite directed cycles.

Assumption 1. The graph G is: (i) locally finite, (ii) finitely connected, and (iii) acyclic.¹

Following [21], we define *layers* in G as follows. Let r be an arbitrary node in G . The first layer of nodes, denoted L_1 , consists of node r and all nodes that are adjacent to r ; that is, $L_1 := \{r\} \cup \{i : (i, r) \in \mathcal{A} \text{ or } (r, i) \in \mathcal{A}\}$. We define other layers in the graph recursively as follows:

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

Since G is locally finite and finitely connected, each layer contains a finite number of nodes, every node is included in some layer, and once a node is in layer L_n it is in every subsequent layer L_k for $k > n$. Let $G_n = (L_n, \mathcal{A}_n)$ denote the subgraph of G induced by the layer L_n , where $\mathcal{A}_n := \{(i, j) \in \mathcal{A} : i, j \in L_n\}$ is the set of arcs in the subgraph induced by the layer of nodes L_n .

2.2 Node supplies and arc costs

We associate node and arc data with G to specify a network. Each node i has *supply* b_i and is a *supply node* if $b_i > 0$, a *demand node* if $b_i < 0$, and a *transshipment node* if $b_i = 0$. Each arc (i, j) has *cost* c_{ij} . The tuple $N := (\mathcal{N}, \mathcal{A}, b, c)$ denotes the *infinite network* with node set \mathcal{N} , arc set \mathcal{A} , supplies $b = (b_i : i \in \mathcal{N})$, and arc costs $c = (c_{ij} : (i, j) \in \mathcal{A})$.

Assumption 2. The node and arc data for the network $N := (\mathcal{N}, \mathcal{A}, b, c)$ satisfy: (i) $b_i \geq 0$ for all $i \in \mathcal{N}$, (ii) b_i is integer for all $i \in \mathcal{N}$, (iii) $b \in \ell_\infty(\mathcal{N})$, i.e., there exists a uniform upper bound \bar{b} on all node supplies, and (iv) $c \in \ell_1(\mathcal{A})$.

The *countably-infinite network flow* (CINF) problem (P) on infinite network N is to find a real nonnegative flow vector x that minimizes the cost $Z(x) := \sum_{(i,j) \in \mathcal{A}} c_{ij} x_{ij}$ and satisfies *flow balance* at every node:

$$Z^* = \inf_x Z(x) := \sum_{(i,j) \in \mathcal{A}} c_{ij} x_{ij} \tag{1a}$$

$$(P) \quad \text{s.t.} \quad \sum_{j \in O(i)} x_{ij} - \sum_{j \in I(i)} x_{ji} = b_i \text{ for } i \in \mathcal{N} \tag{1b}$$

$$x_{ij} \geq 0 \text{ for } (i, j) \in \mathcal{A}, \tag{1c}$$

where constraints (1b) are the flow balance constraints. A vector x that satisfies constraints (1b)–(1c) is a *feasible flow*. An *extremal flow* is a feasible flow that cannot be expressed as a strict convex combination of two other feasible flows.

¹Once stated, assumptions hold without further reference. Exceptions are explicitly noted.

In general, $Z(x)$ may be undefined for some feasible flow x , or Z^* may be infinite. We now make assumptions that guarantee neither of these abnormalities occur.

We begin with some pre-processing. Identify all transshipment nodes with out-degree zero. Since no feasible flow will send positive flow along arcs into such nodes, they can be removed along with all of their adjacent arcs without loss of generality. Apply this rule recursively until no such nodes remain. Moreover, without loss of *feasibility*, each supply node has out-degree at least one — otherwise flow balance is violated. We have thus established the following result.

Proposition 2.1. In every feasible instance of (P), each node has out-degree at least one without loss of generality.

We call node i a *predecessor* of node j (in G) if there exists a directed path P_{ij}^{\rightarrow} .

Lemma 2.2. Every node has finitely-many predecessors.

Proof. Suppose $i \in \mathcal{N}$ has infinitely many predecessors. Since the graph is locally finite, this means there exists a directed path from infinity to node i . Since all nodes have out-degree at least one (via [Proposition 2.1](#)) and there are no finite cycles in the graph, there must also be an infinite directed path from i to infinity. This implies G has a directed infinite cycle, contradicting [Assumption 1\(iii\)](#). \square

Assumption 3. $G = (\mathcal{N}, \mathcal{A})$ has finitely many nodes with in-degree 0.

[Lemma 2.2](#) and [Assumption 3](#) together allow us to specify an ordering of the nodes and arcs. We define *stages* of nodes as follows. Stage 0 is the set of all nodes with in-degree 0. This set is finite by [Assumption 3](#). Stage 1 consists of all nodes with in-degree 0 in the modified graph that results from removing all Stage 0 nodes and their adjacent arcs. This set is again finite since each node with in-degree 0 in the modified graph must be adjacent to one of the finitely-many removed nodes. Repeat this procedure to construct the remaining stages. Since the graph is acyclic, each node is contained in exactly one stage. Each stage is finite and the only possible incoming arcs into nodes in Stage k must have tails at nodes in earlier stages. Let S_k denote the set of nodes in Stage k . We label nodes with the natural numbers \mathbb{N} so that all nodes in S_k have smaller labels than the nodes in S_{k+1} , for all k . This means $i < j$ for every $(i, j) \in \mathcal{A}$. Arcs can also be labeled with the natural numbers so that arcs with tails in lower-numbered stages have smaller labels than arcs with tails in higher-numbered stages. Let $s(i)$ denote the stage of node i . Clearly, $s(i) \leq i$, and $s(i) \leq s(j)$ if $i < j$.

Assumption 4. Network $N := (\mathcal{N}, \mathcal{A}, b, c)$ satisfies the following: (i) there exist $\beta \in (0, 1)$ and $\gamma \in (0, +\infty)$ such that for every $(i, j) \in \mathcal{A}$, $|c_{ij}| \leq \gamma\beta^{s(i)}$, where β can be interpreted as a discount factor, and (ii) there exists a sub-exponential function $g(k)$ such that the cardinality of Stage k is bounded above by $g(k)$: $|S_k| \leq g(k)$ for all k . In particular, we require $\sum_{k=0}^{\infty} \beta^k \sum_{j=0}^k g(j) < \infty$. Any bounding polynomial $g(k)$ suffices.

Remark 2.3. Before proceeding, we point out a potential point of confusion between stages and layers. Both are finite subsets of nodes. However, given an infinite network satisfying [Assumptions 1](#) through [3](#), there is a unique staging of the nodes. By contrast, there are many possible layerings — indeed, any choice of a node as the root node r produces a different layering.

Definition 2.4. A network $N := (\mathcal{N}, \mathcal{A}, b, c)$ is called a *pure-supply network* if it satisfies [Assumptions 1](#) through [4](#).

In the remainder of the paper we assume the network underlying problem (P) is a pure supply network. Consequently, we call (P) a pure-supply problem.

2.3 Properties of pure-supply problems

We now explore basic properties of pure supply problems.

Lemma 2.5. Every feasible flow has finite objective value.

Proof. Let x be a feasible flow. We have

$$\sum_{k=0}^{\infty} \sum_{i \in S_k} \sum_{j: (i,j) \in \mathcal{A}} |c_{ij}| x_{ij} \leq \bar{b} \gamma \sum_{k=0}^{\infty} \beta^k \sum_{j=0}^k g(j) < \infty,$$

where the first inequality holds since $x \geq 0$ and all flow on arcs with tails in Stage k must come from supplies at nodes in the first k stages (since the graph is acyclic) and hence can be bounded by $\bar{b} \sum_{j=0}^k |S_j| \leq \bar{b} \sum_{j=0}^k g(j)$, and the cost of corresponding arcs is bounded above by $\gamma \beta^k$. The second inequality follows from [Assumption 4\(ii\)](#). Thus, the sum $\sum_{k=0}^{\infty} \sum_{i \in S_k} c_{ij} x_{ij}$ converges absolutely, and the cost of flow x can be broken up by stage; that is,

$$\sum_{(i,j) \in \mathcal{A}} c_{ij} x_{ij} = \sum_{k=0}^{\infty} \sum_{i \in S_k} \sum_{j: (i,j) \in \mathcal{A}} c_{ij} x_{ij},$$

and is thus finite (see, e.g., [19, Theorem 3.55]). □

Next, we discuss important topological properties of the feasible region and objective function. Let $A(x) = \{(i, j) \in \mathcal{A} : x_{ij} > 0\}$ denote the set of *active arcs* of feasible flow x .

Lemma 2.6. Suppose x is a feasible flow with active arc $(i, j) \in A(x)$. Then there exists a directed path from i to infinity in the graph $G(x) = \{\mathcal{N}, A(x)\}$.

Proof. Every node in the network is either a supply or a transshipment node. Since $x_{ij} > 0$, node j has some incoming flow, and so by flow balance, there must be flow leaving node j along at least one arc $(j, j_1) \in A(x)$. Repeat this argument for node j_1 , etc. Since there are no directed cycles in the graph, this generates a directed path from i to infinity in $G(x)$. □

Lemma 2.7. The set of all feasible flows is compact in the product topology.

Proof. We apply Tychonoff's theorem. The first step is to show that the flow on each arc is uniformly bounded across all feasible flows. Constraint (1c) gives a lower bound of 0, so it remains to show that each arc (i, j) has an implied upper bound u_{ij} .

Suppose otherwise that there exists an arc (i, j) and a sequence of feasible flows x^n , $n = 1, 2, \dots$ with $x_{ij}^n \rightarrow \infty$ as $n \rightarrow \infty$. Without loss of generality, for any n we have $x_{ij}^n > 0$. By [Lemma 2.6](#), there exists a directed path from i to infinity consisting of arcs in $A(x^n)$. Since the graph does not have finite or infinite directed cycles, the flow on arc (i, j) in x^n must originate from a subset of supply nodes in the network, which we denote by $M_n \subset \mathcal{N}$. Note that $x_{ij}^n \leq \sum_{k \in M_n} b_k \leq \bar{b} \cdot |M_n|$. Since $x_{ij}^n \rightarrow \infty$, this implies $|M_n| \rightarrow \infty$ as $n \rightarrow \infty$. On the other hand, nodes of M_n are predecessors of i for any n , and by [Lemma 2.2](#) the cardinality of these sets must be bounded, a contradiction.

Hence, the flow on every arc (i, j) has a lower bound of 0 and an implied upper bound of $u_{ij} < \infty$. Without loss of generality, $u_{ij} \geq 1$. By Tychonoff's theorem, the "rectangle" $\prod_{(i,j) \in \mathcal{A}} [0, u_{ij}]$ is compact in the product topology.

It remains to show that the feasible region is a closed subset of this rectangle and thus compact. Each constraint in (1b) is of the form $\phi_i(x) = b_i$, where ϕ_i is a linear functional on $\mathbb{R}^{\mathcal{N}}$ with finite support. Hence, ϕ_i is continuous in the product topology (in fact, every linear topology on $\mathbb{R}^{\mathcal{N}}$) and so the set $\{x : \phi_i(x) = b_i\}$ is closed, being the pre-image of the closed set $\{b_i\}$ under a continuous function (see, e.g., [3, Theorem 2.27]). Hence, the feasible region of (1) is an intersection of closed sets and thus closed. \square

Remark 2.8. The upper bounds u_{ij} are implied by our assumptions and are not given explicitly as constraints in the formulation. Moreover, the implied bounds u_{ij} need not satisfy any uniformity properties across arcs. Thus, the instances of (P) we explore may fail the condition in Proposition 2.5 of [21].

Following [20], we define a new topology based on the u_{ij} 's defined in Lemma 2.7. Let $X_{ij} = [-u_{ij}, u_{ij}]$ for $(i, j) \in \mathcal{A}$ (observe that since $u_{ij} \geq 1$ we have $[-1, 1] \subseteq X_{ij}$), and let $X = \prod_{(i,j) \in \mathcal{A}} X_{ij}$. Note that X contains every feasible flow and the differences of any two feasible flows — a property we will need later (see the proof of Lemma 3.4). Adopting the relative topology σ inherited from the product topology on $\mathbb{R}^{\mathcal{N}}$, X is a compact topological space (via Tychonoff's Theorem, [3, Theorem 2.61]).

Theorem 2.9. The objective function $Z(x) = \sum_{(i,j) \in \mathcal{A}} c_{ij}x_{ij}$ is continuous in the relative topology σ over X .

Proof. This follows from Theorem 2.2 of [20] that depends on three assumptions (Assumptions A, B, and C) from that paper. Reordering nodes and arcs according to stages yields the "staircase" structure of [20]. Assumption A follows from the bounds u_{ij} in Lemma 2.7. Assumption B follows since the constraints in (P) are continuous in the product topology. Assumption C holds by linearity of the objective function $Z(x)$ and the staged structure of the costs, via a comment in [20] that immediately precedes the statement of proof of Theorem 2.2. \square

Corollary 2.10. Optimal cost Z^* is finite and achieved at an extremal flow.

Proof. Existence of an optimal extremal flow x^* is immediate from Bauer's minimum principle [3, Theorem 7.69]. Objective function $Z(x)$ is linear and continuous on the feasible region (via Theorem 2.9), and the feasible region is compact (via Lemma 2.7). By Lemma 2.5, $Z(x^*)$ is finite, and hence so is $Z^* = Z(x^*)$. \square

Corollary 2.10 shows that a simplex method that pivots between extremal flows has some hope of finding an optimal solution to (P). Moreover, it establishes a "finite target" Z^* for proving optimal value convergence. A similar result is established in [11] by relying on *explicit* bounds on all variables.

Remark 2.11. The construction of the topological space X and the relative topology σ are of critical importance here, due to the fact that the only continuous functions over the *whole* space $\mathbb{R}^{\mathcal{N}}$ with the product topology are those in the space of finite-support real sequences (for a discussion see [3, Chapter 16]). This is unacceptable for our applications, as limiting the discussion to such

functions would imply zero arc costs for all but finitely many arcs. However, use of a modified topology makes establishing duality via standard methods of [4] no longer possible. We are not working with the full topological vector space $\mathbb{R}^{\mathcal{N}}$ and its associated topological dual. This is the reason we must develop an alternative approach to duality in [Section 7](#) below.

Remark 2.12. Any result derived for a pure-supply problem can be applied to the case of pure demands, in which [Assumption 2\(i\)](#) is replaced with $b_i \leq 0$ for all $i \in \mathcal{N}$. Any instance of the latter is equivalent to a pure-supply problem with supplies $-b_i \geq 0$ for all i and with the direction of all arcs reversed. Our exposition sticks to pure supply problems acknowledging this correspondence.

3 Basic feasible flows

In this section we discuss the concepts of trees, basic feasible flows, and pivoting between basic feasible flows. At first glance, these concepts behave similarly to the finite-dimensional case. However, extensions to the infinite case require particular care, and frequently rely on our assumptions in [Section 2](#).

A *forest* is a subgraph of G that contains no finite or infinite cycles, and a *spanning tree* is a connected forest with arcs incident to all nodes in the graph (recall that we allow pairs of nodes to be connected at infinity). These definitions differ from the definitions in the infinite graph theory literature, where a forest only requires the absence of *finite* cycles (see, for instance, [8, Chapter 8]). We employ this stricter definition since it enables us to characterize extremal flows in networks. An important characterizing property of spanning trees is that, for each node $i \in \mathcal{N}$, there exists a unique path from i to infinity that uses only tree arcs. Indeed, if a node is not on a path to infinity, then it can only be connected to finitely-many nodes, violating the connectedness property of a tree, and multiple paths to infinity indicate a presence of a finite or infinite cycle.

Theorem 3.1 (cf. Theorem 3.13, [18]). Every forest in a locally-finite connected graph is contained in a spanning tree.

We call a vector x satisfying constraints (1b) a *basic flow* if the arcs $\{(i, j) \in \mathcal{A} : x_{ij} \neq 0\}$ form a forest in G . [Theorem 3.1](#) ensures that a spanning tree can be associated with every basic flow. A basic flow is a *basic feasible flow* if it also satisfies nonnegativity constraints (1c). In short, a basic feasible flow is a feasible flow whose active arcs form a forest.

We now justify the term *basic flow* in the infinite network setting. Given a basic flow x and a spanning tree T containing all arcs (i, j) with $x_{ij} \neq 0$, we call the arcs in T *basic arcs* (denoted $A(T)$), and the arcs not in T *nonbasic arcs* (denoted $\overline{A(T)}$). We also refer to the set $A(T)$ as a *basis*. The doubly-infinite node-arc incidence matrix M can be written as $M = (M_{A(T)}, M_{\overline{A(T)}})$, by possibly rearranging columns, where $M_{A(T)}$ and $M_{\overline{A(T)}}$ contain the columns of M associated with arcs in $A(T)$ and $\overline{A(T)}$, respectively. Analogous to the finite case, since T is a spanning tree, the submatrix $M_{A(T)}$ determines a bijective linear map, which can be shown by an argument similar to the proof of Theorem 7.3 in [5] and thus omitted here.

Basic feasible flows are tightly connected to extremal flows. In finite-dimensional network flow problems, basic feasible flows and extremal flows coincide. Unfortunately, as first pointed out by [18], this may not hold for infinite networks. However, this equivalence can be recovered due to [Assumption 2\(ii\)](#) and the following result.

Theorem 3.2 (cf. Theorems 3.13 and 3.14, [18]). If supplies b_i are integer for all $i \in \mathcal{N}$, a flow x is an extremal flow of (P) if and only if x is a basic feasible flow. Moreover, in every basic feasible flow x , x_{ij} is integer-valued for all $(i, j) \in \mathcal{A}$.

From Theorem 3.2, a feasible flow x is an extremal flow if and only if its active arcs form a forest. We call a basic feasible flow x *nondegenerate* if the forest $G(x) = (\mathcal{N}, A(x))$ is a spanning tree; otherwise we call x *degenerate*. As in the finite case, degeneracy can be problematic in the simplex method if there are multiple spanning trees containing $G(x)$.

The simplex method systematically *pivots* between spanning trees and their corresponding basic feasible flows. Let T be a spanning tree containing $G(x)$ for extremal flow x . Adding any nonbasic *entering arc* a_\uparrow to T creates a unique cycle C (in the undirected sense), which can be either finite or infinite. Removing any arc in C results in another spanning tree. We orient C so that a_\uparrow is a forward arc and let C^F and C^B denote the sets of forward and backward arcs in C under that orientation. This defines a flow vector h^C — called the *simple circulation* associated with C — where h_{ij}^C is 1 if $(i, j) \in C^F$, -1 if $(i, j) \in C^B$, and 0 otherwise.

A new flow \hat{x} is derived by sending flow θ around the simple circulation h^C ; that is, $\hat{x} = x + \theta h^C$, where $\theta = \inf_{(i,j) \in C^B} x_{ij}$ (and thus equal to ∞ if C^B is empty). Since every basic feasible flow is integer-valued (Theorem 3.2), the infimum defining θ is achieved (when $C^B \neq \emptyset$), and any arc that achieves it is a valid choice of *leaving arc* a_\downarrow from the basis. Every valid choice results in a new spanning tree \hat{T} by swapping out arc a_\downarrow for a_\uparrow . It is easy to see that \hat{x} is the basic feasible flow associated with \hat{T} . \hat{T} is said to be *adjacent* to T since they differ only by two arcs.

If $\theta > 0$, the new basic feasible flow \hat{x} is different from x (also called adjacent). If, however, $\theta = 0$, the pivot is *degenerate* and \hat{T} is an alternate spanning tree representation of the same extremal flow x . This implies that x itself is degenerate.

We now describe how the objective value changes with a pivot. For every nonbasic arc $(i, j) \notin A(T)$, let $C(i, j)$ denote the unique cycle formed when adding arc (i, j) to T . The *reduced cost* of arc $(i, j) \notin A(T)$ is

$$\bar{c}_{ij} = \sum_{(k,\ell) \in C(i,j)^F} c_{k\ell} - \sum_{(k,\ell) \in C(i,j)^B} c_{k\ell},$$

and $\bar{c}_{ij} = 0$ for $(i, j) \in A(T)$. Although traditionally not explicitly reflected in notation, the reduced cost of an arc is defined *with respect to* a specific spanning tree T . Since $c \in \ell_1(A)$, the reduced cost of every nonbasic arc is finite. After a pivot with a_\uparrow as the entering arc, the change in objective value is precisely $\theta \bar{c}_{a_\uparrow}$.

There is a simple optimality condition for infinite network flow problems involving the reduced costs of nonbasic variables (Theorem 3.6 below). Its proof requires the following lemmas. Recall that $C(i, j)$ denotes the unique cycle formed when adding (i, j) to spanning tree T .

Lemma 3.3. For any spanning T , arc $a \in A(T)$ is contained in finitely many of the cycles $C(i, j)$ for $(i, j) \notin A(T)$.

Proof. We call node i a *tree-predecessor* of arc a if a is contained in the unique tree path from i to infinity in T . Note that a belongs to $C(i, j)$ only if either i or j is a tree-predecessor of a . Moreover, since T contains no infinite cycles, a can have at most finitely-many tree-predecessors, each of which has finitely many adjacent arcs (due to the local finiteness of G). Taken together, this implies that a lies in at most finitely many cycles $C(i, j)$ for $(i, j) \notin A(T)$. \square

Lemma 3.4. Let y be a circulation in G , i.e., $y \in \mathbb{R}^A$ such that $My = 0$, that arises as the difference of two feasible flows. Let T be a spanning tree. Then y can be decomposed among simple circulations as follows:

$$y = \sum_{(i,j) \in \overline{A(T)}} y_{ij} h^{C(i,j)}. \quad (2)$$

Moreover, when y has finite cost,

$$Z(y) = \sum_{(i,j) \in \mathcal{A}} \bar{c}_{ij} y_{ij}. \quad (3)$$

Proof. Let $\tilde{y} = \sum_{(i,j) \in \overline{A(T)}} y_{ij} h^{C(i,j)}$. We claim this is a well-defined sum. For $\tilde{y}_{k\ell}$ where $(k, \ell) \notin A(T)$, only cycle $C_{k\ell}$ in the sum defining \tilde{y} contains arc (k, ℓ) and so $\tilde{y}_{k\ell}$ is clearly finite. For $\tilde{y}_{k\ell}$ where $(k, \ell) \in A(T)$, by Lemma 3.3, only finitely many of the cycles $C(i, j)$ contain arc (k, ℓ) and $\tilde{y}_{k\ell}$ also arises from a finite sum. This implies that $\tilde{y}_{k\ell}$ is well-defined for all $(k, \ell) \in \mathcal{A}$, and thus \tilde{y} is well-defined.

For any $(i, j), (k, \ell) \in \overline{A(T)}$, $h_{kl}^{C(i,j)} = 1$ if $(i, j) = (k, \ell)$ and 0 otherwise, since every basic simple circulation contains exactly one nonbasic arc. Moreover, for $(i, j) \in \overline{A(T)}$, $\tilde{y}_{ij} = \sum_{(k,\ell) \in \overline{A(T)}} y_{k\ell} h_{ij}^{C_{k\ell}} = y_{ij} h_{ij}^{C(i,j)} = y_{ij}$. Hence $\tilde{y}_{\overline{A(T)}} = y_{\overline{A(T)}}$. Combining this with the observation that $My = M\tilde{y} = 0$, we see that $M_{A(T)}\tilde{y}_{A(T)} = M_{A(T)}y_{A(T)}$. Since $M_{A(T)}$ is a bijection, $\tilde{y}_{A(T)} = y_{A(T)}$, and so $y = \tilde{y}$.

To derive (3), we appeal to the continuity of Z over the topological space X from Theorem 2.9. Since y is a difference of feasible flows, we have $y \in X$ since for any feasible flow x , $0 \leq x_{ij} \leq u_{ij}$ for all $(i, j) \in \mathcal{A}$. Thus $\tilde{y} = y \in X$. Moreover, $h^{C(i,j)} \in X$ since $[-1, 1]^A \subseteq X$, and hence $Z(\tilde{y}) = Z\left(\sum_{(i,j) \in \mathcal{A}} y_{ij} h^{C(i,j)}\right) = \sum_{(i,j) \in \mathcal{A}} y_{ij} Z(h^{C(i,j)})$, where the second equality holds by countable additivity, which is a consequence of continuity of Z in the topology over X (for a discussion of this property see [13], noting the fact from [3, Chapter 16] that continuity in the product topology over \mathbb{R}^N implies countable additivity and so specializes to the relative topology over X).

Finally,

$$Z(y) = Z(\tilde{y}) = \sum_{(i,j) \in \overline{A(T)}} y_{ij} Z(h^{C(i,j)}) = \sum_{(i,j) \in \overline{A(T)}} y_{ij} \bar{c}_{ij} = \sum_{(i,j) \in \mathcal{A}} \bar{c}_{ij} y_{ij},$$

where third equality follows from the definition of reduced costs and the fourth equality uses the fact that $\bar{c}_{ij} = 0$ for $(i, j) \in A(T)$. This completes the proof. \square

For the following two results, recall that any feasible flow has finite cost by Lemma 2.5, and optimal value Z^* is finite by Corollary 2.10.

Corollary 3.5. Let x be an extremal flow with an associated spanning tree T , and let f be an arbitrary feasible flow. Then $f - x = \sum_{(i,j) \in \overline{A(T)}} f_{ij} h^{C(i,j)}$.

Proof. $f - x$ is a circulation that satisfies assumptions of Lemma 3.4. Combining (2) with the observation that, for any $(i, j) \in \overline{A(T)}$, $x_{ij} = 0$ and so $f_{ij} - x_{ij} = f_{ij}$, gives the desired expression. \square

Theorem 3.6 (Optimality condition). Let x be an extremal flow and let T be a spanning tree containing $G(x)$ such that the reduced costs \bar{c}_{ij} are nonnegative for all nonbasic arcs (i, j) . Then $Z(x) = Z^*$ and x is an optimal solution to (P).

Proof. Let f be a feasible flow. **Corollary 3.5** implies $f - x = \sum_{(i,j) \in \overline{A(T)}} f_{ij} h^{C(i,j)}$. Applying (3) yields $Z(f - x) = \sum_{(i,j) \in \mathcal{A}} \bar{c}_{ij}(f_{ij} - x_{ij}) = \sum_{(i,j) \in \overline{A(T)}} \bar{c}_{ij} f_{ij} \geq 0$. That is, $Z(x) \leq Z(f)$ for every feasible f . Since Z^* is finite, $Z(x) = Z^*$ and x is an optimal solution to (P). \square

4 A simplex method

In this section we introduce a simplex method for pure supply problems. In general, each iteration of this algorithm can require an infinite amount of data and computation to execute. In this sense, it serves as an abstract template for practical, finitely implementable methods. Such a finite implementation is indeed possible for structured instances, as illustrated in Section 6 below and in the literature [10, 12, 21].

Degenerate basic feasible flows and degenerate pivots raise the possibility of *cycling*. We propose a simplex method that does not cycle on pure-supply networks and converges to optimality.

We begin by showing that a pure supply problem always has a spanning tree with the corresponding basic feasible flow where a simplex method can be initialized. Consider the following procedure for constructing an initial spanning tree:

Procedure 1 (Constructing an initial spanning tree). Given an instance of (P), (i) for every node i select a single outgoing arc a_i (such an arc is guaranteed to exist by **Proposition 2.1**), and (ii) construct subgraph S with arc set $\{a_i : i \in \mathcal{N}\}$.

A spanning tree S is an *in-tree rooted at infinity* if for each node $i \in \mathcal{N}$, the unique path from i to infinity in S contains only forward arcs.

Lemma 4.1. A subgraph S of a pure supply network is a spanning in-tree rooted at infinity if and only if it can be constructed by Procedure 1, i.e., S contains exactly one outgoing arc for each node $i \in \mathcal{N}$.

Proof. (if) Suppose S has exactly one outgoing arc for each node $i \in \mathcal{N}$. Clearly, every node i is in S . Furthermore, S has no finite or infinite cycles. Indeed, every node has only one outgoing arc in S and so any cycle in S would have to be directed. However, since G has no directed cycles, S cannot have any cycles. Finally, since S has no cycles and each node has one outgoing arc included in S , for any $i \in \mathcal{N}$, we can construct a directed path $P_{i\infty}^{\rightarrow}$ in S by starting in i and following the sequence of outgoing arcs selected at each subsequent nodes. Therefore, S is connected, and thus a spanning tree. Moreover, as observed above, all arcs on the unique path from any node to infinity are forward arcs. In other words, S is a spanning in-tree rooted at infinity.

(only if) Suppose S is a spanning in-tree rooted at infinity, but for some $i \in \mathcal{N}$, it includes arcs (i, j) and (i, k) , with $j \neq k$. By definition of a spanning in-tree rooted at infinity, S contains paths $P_{j\infty}^{\rightarrow}$ and $P_{k\infty}^{\rightarrow}$; combination of these paths with arcs (i, j) and (i, k) contains a cycle within S , resulting in a contradiction. \square

Procedure 2 (Constructing a basic flow from a tree). Given a spanning tree S , construct a flow x^S on S as follows: start with $x^S = 0$. For each $i \in \mathcal{N}$, identify the unique path $P_{i\infty}$ from i to infinity in S , with forward arcs $P_{i\infty}^F$ and backward arcs $P_{i\infty}^B$, and add flow of b_i to all arcs in $P_{i\infty}^F$ and remove flow of b_i from all arcs in $P_{i\infty}^B$.

Note that this procedure can be applied to any spanning tree (not necessarily a spanning in-tree rooted at infinity). It is easy to see that x^S always satisfies flow balance constraints (1b). Moreover, x^S has zero flow outside the arcs of the spanning tree S . Therefore, it is a basic feasible flow if and only if $x_a^S \geq 0$ for all arcs a in S .

Lemma 4.2. If S is a spanning in-tree rooted at infinity in a pure supply network, the flow x^S is a basic feasible flow.

Proof. If S is a spanning in-tree rooted at infinity, any path $P_{i\infty}$ in S contains only forward arcs. For any i , x_{a_i} can be calculated as the sum of supplies at node i and all nodes j that are tree-predecessors of i in S . For any $i \in \mathcal{N}$, x_{a_i} is finite (since G , and thus S , is locally finite) and $x_{a_i} \geq 0$ (since all supplies are nonnegative). Since $x_a^S \geq 0$ for all $a \in S$, x^S is a basic feasible flow. \square

Procedures 1 and 2 can be used to construct a basis that is a spanning in-tree rooted at infinity and the corresponding basic feasible solution, which can be used to initialize a simplex algorithm. Next, we discuss a particular pivoting rule we will use.

Procedure 3 (Pivot rule). Given a spanning in-tree rooted at infinity S , construct a new tree as follows: (i) Let $(i, j) \in \mathcal{A} \setminus A(S)$ be the entering arc. Add this arc to S to create an (undirected) cycle C in S . (ii) The leaving arc is $(i, j') \in A(S)$ with $j' \neq j$. Remove this arc from S to obtain a new tree, S' . There is only one such arc so this choice is unique.

Lemma 4.3. The output S' of Procedure 3 is a spanning in-tree rooted at infinity. Consequently, $x^{S'}$ is a basic feasible flow.

Proof. Since S is a spanning in-tree rooted at infinity, every node has exactly one outgoing arc included in S , by Lemma 4.1. Procedure 3 consists of replacing one outgoing arc at i with another, and therefore, the new tree S' has the same property. Thus, S' is also a spanning in-tree rooted at infinity (again, by Lemma 4.1), and $x^{S'}$ is a basic feasible flow by Lemma 4.2. \square

With Procedures 1, 2, and 3 in hand, we state our simplex method in Algorithm 1 (displayed below). Lemmas 4.1 and 4.3 imply that all of the spanning trees S^m produced in the algorithm (and hence also the trees T^n arising from line 10 of Algorithm 1) are spanning in-trees rooted at infinity and all x^m 's are basic feasible solutions by Lemma 4.2. All trees S^m are, in fact, strongly feasible trees, as originally defined in [6] with the root node being the “virtual” node at infinity. In fact, they have even more structure than strongly feasible trees since all paths to the root are directed in in-trees.

Our simplex algorithm shares the layer-based structure of the algorithm presented in [21], but our pivoting procedure is not explicitly considered in their development. Although they specify that their algorithm works if any anti-cycling method developed for finite problems is applied to select entering and exiting arcs inside the **while** loop, they only mention Bland’s rule specifically. Bland’s rule restricts the choice of both entering and leaving variables. In our algorithm we are free to choose the entering variable, e.g., we can choose a direction of steepest descent in the layer. Despite this flexibility, Theorem 4.4 below shows that the algorithm does not cycle. The similarity between our simplex method and that of [21], which applies to a complementary set of problems, demonstrates the broad power of simplex-like methods for countably-infinite networks.

Theorem 4.4. Algorithm 1 does not cycle, i.e., the spanning trees S^m in iterations $m = 0, 1, \dots$ (line 8) are all distinct.

Algorithm 1 The layered network simplex method for pure supply network flow problems

```

1: Input: A pure supply network  $N = (\mathcal{N}, \mathcal{A}, b, c)$ .
2: Initialization: Construct a spanning tree  $S$  and corresponding basic feasible flow  $x^S$  using
   Procedures 1 and 2.
3:  $S^0 \leftarrow S, x^0 \leftarrow x^S, n \leftarrow 0, m \leftarrow 1$ 
4: while there exists a nonbasic arc with negative reduced cost do
5:   if there exists a nonbasic arc  $a_\uparrow$  with negative reduced cost in  $A_n$  then
6:     Apply Procedure 3 to  $S$ , with  $a_\uparrow$  entering in step (i), producing  $S'$ 
7:     Use Procedure 2 to produce  $x^{S'}$ 
8:      $m \leftarrow m + 1, S \leftarrow S', S^m \leftarrow S, x^m \leftarrow x^S$ 
9:   else
10:     $n \leftarrow n + 1, T^n \leftarrow S$ 
return  $x^m$ 

```

Our approach to proving [Theorem 4.4](#) uses a correspondence between two related networks. Let $N = (\mathcal{N}, \mathcal{A}, b, c)$ denote the network of our original problem, and recall ([Proposition 2.1](#)) that, without loss of generality, each node has out-degree of at least 1. Let $N' = (\mathcal{N}, \mathcal{A}, b', c)$ denote the network with the same graph and arc costs, but with supply $b'_i = b_i$ if $b_i > 0$ and $b'_i = 1$ if $b_i = 0$. We say that the transshipment nodes of N have an *augmented supply* of 1 in N' . Observe that if N is a pure supply network, then so is N' .

Lemma 4.5. Let S be a spanning tree associated with a basic flow y^S in N' . Then (i) for every arc $(i, j) \in \mathcal{A}$, the reduced cost of arc (i, j) with respect to S in N' is equal to the reduced cost of arc (i, j) with respect to S in N and (ii) the basic flow x^S in N can be constructed from the basic flow y^S by removing the flow originating from the nodes with augmented supply. More precisely, for every transshipment node i in N there is a unique path to infinity $P_{i\infty}$ in S , and the flow x^S is equal to $y^S - \sum_{i \in \mathcal{N}: b_i = 0} \chi_{P_{i\infty}}$, where $\chi_{P_{i\infty}}$ is the characteristic vector of the path $P_{i\infty}$.

Proof. The proof of (a) is straightforward, since, given the tree of basic arcs, reduced costs depend only on the underlying graph and arc costs, which are identical across N and N' . To prove (b), note that another way to express Procedure 2 for a spanning tree S is $x^S = \sum_{i \in \mathcal{N}} b_i \chi_{P_{i\infty}}$, where $\chi_{P_{i\infty}}$ is the characteristic vector of the unique path $P_{i\infty}$ in S . Observe that

$$y^S = \sum_{i \in \mathcal{N}} b'_i \chi_{P_{i\infty}} = \sum_{i \in \mathcal{N}: b_i > 0} b_i \chi_{P_{i\infty}} + \sum_{i \in \mathcal{N}: b_i = 0} \chi_{P_{i\infty}} = x^S + \sum_{i \in \mathcal{N}: b_i = 0} \chi_{P_{i\infty}}.$$

It is then immediate that $y^S - \sum_{i \in \mathcal{N}: b_i = 0} \chi_{P_{i\infty}}$ yields x^S . □

Proof of Theorem 4.4. Suppose S^m , $m = 0, 1, \dots$ is the sequence of spanning trees visited by [Algorithm 1](#) applied to network N . We first argue that the same sequence of trees can be generated by applying [Algorithm 1](#) to network N' . Since S^0 is constructed using Procedure 1, it is a spanning in-tree rooted at infinity, and thus both x^0 and y^0 are basic feasible flows in their respective networks. Proceeding by induction, suppose S^m is the tree in iteration m of both algorithms. By [Lemma 4.5](#), arcs have the same reduced costs with respect to S^m in both networks, and thus entering arc a_\uparrow chosen by the algorithm applied to N can be chosen by the algorithm applied to N' . Since Procedure 3 prescribes a unique choice of leaving arc given a_\uparrow , the same tree will be generated by both algorithms after the pivot, concluding the inductive argument.

Now, for any spanning in-tree S rooted at infinity, the corresponding basic feasible (in network N') flow y^S constructed via Procedure 2 has a flow of at least 1 on every tree arc. Therefore, all iterates y^{S^m} , $m = 0, 1, \dots$ of the simplex method applied to N' are nondegenerate, and every pivot of the algorithm applied to N' decreases the objective value by at least $|\bar{c}_{a^\dagger}| > 0$. Thus, all flows y^{S^m} are distinct, implying that all trees S^m , $m = 0, 1, \dots$ are distinct. \square

Theorem 4.6. The layered network simplex method either terminates with an optimal flow or generates an infinite sequence of adjacent extreme points x^0, x^1, x^2, \dots with nonincreasing objective values. If the algorithm does not terminate, then $n \rightarrow \infty$, and so entering variables from all layers in the graph are eventually considered.

Proof. If the algorithm escapes the **while** loop in lines 4–10, then the last basic feasible flow x^m has nonnegative reduced costs for all of its nonbasic variables. It follows by Theorem 3.6 that x^m is an optimal flow.

Now consider the case where the **while** loop is never escaped, i.e., the algorithm generates an infinite sequence of adjacent trees and corresponding extreme points x^0, x^1, x^2, \dots with nonincreasing objective values. We claim that **else** in line 9 of the algorithm is visited infinitely often, i.e., $n \rightarrow \infty$ and every layer of G is eventually reached. Suppose otherwise that line 9 is only visited a finite number of times, i.e., there is a point in the algorithm after which the **if** on line 5 of the algorithm is visited at every remaining iteration of the **while** loop (of which there are infinitely many), and n remains constant. Let N denote the final value of n . Note that all arcs entering the basis throughout the algorithm must be contained in A_N .

There are two types of pivots that can occur: *Type-1 pivots*, where the leaving arc is in A_N , and *Type-2 pivots*, where the leaving arc is not in A_N . In each Type-2 pivot, the number of basic arcs in A_N must increase by 1. Indeed, as mentioned above, only arcs in A_N enter the basis, so once an arc outside of A_N leaves, it is replaced by a new basic arc in A_N . Since A_N is finite, only a finite number of Type-2 pivots is possible.

Thus, after a finite number of iterations, only Type-1 pivots are possible, and an infinite number of them are performed. However, since both entering and leaving arcs must come from A_N , there are now only finitely many possible spanning trees that will be visited by pivoting (since all basic arcs outside of A_N are fixed during Type-1 pivots). Since there are infinitely-many consecutive Type-1 pivots and only finitely-many possible spanning trees to pivot to, eventually the algorithm must cycle. This contradiction to Theorem 4.4 completes the proof. \square

We next establish an important topological property of the trees generated by Algorithm 1 used in our convergence proof. We say a sequence of subgraphs S^k , converges in the *product discrete topology* to a subgraph S if for every arc $a \in \mathcal{A}$, there exists a sufficiently large K_a such that for $k \geq K_a$, $a \in S^k$ if and only if $a \in S$ (we refer to this behavior as “locking in” of the arcs). In other words, for any finite subset of \mathcal{A} , S^k ’s agree with S on this set of arcs for sufficiently large k .

Theorem 4.7. Let S^k denote any sequence of spanning in-trees rooted at infinity. There exists a convergent (in the product discrete topology) subsequence of these trees that converges to a spanning tree S . Moreover, S is a spanning in-tree rooted at infinity and x^S is a basic feasible solution.

Proof. The set of all subgraphs is compact in the product discrete topology by Tychonoff’s theorem, since node degrees are finite. Hence, S^k possesses a convergent subsequence with a subsequential limit S . For convenience, we refer to this subsequence again as S^k .

Consider an arbitrary node $i \in \mathcal{N}$. Due to the locking in behavior of the arcs, trees S^k agree with S on all arcs adjacent to i for sufficiently large k . Since each S^k is a spanning in-tree rooted at infinity, node i has exactly one outgoing arc in each S^k ; eventually a single outgoing arc is going to lock in, and be the only outgoing arc present at i in S . Thus by [Lemma 4.1](#), S is a spanning in-tree rooted at infinity and x^S is a basic feasible flow by [Lemma 4.2](#) \square

We are now ready to state and prove our main result.

Theorem 4.8. The iterates x^m generated by [Algorithm 1](#) converge in value to the optimal value of (P). That is, $Z(x^m) = \sum_{(i,j) \in A} c_{ij} x_{ij}^m \rightarrow Z^*$ as $m \rightarrow \infty$. Moreover, there exists a subsequence of simplex iterates that converges (in the relative topology σ on X described in [Section 2](#)) to an extremal optimal solution of (P).

Proof. If [Algorithm 1](#) terminates finitely, it clearly returns a tree and the corresponding optimal basic feasible flow; therefore, we consider the case where $m \rightarrow \infty$. Let T^m denote spanning trees generated in line 10 of the algorithm; they form an infinite sequence by [Theorem 4.6](#). Since all of these trees are spanning in-trees rooted at infinity, by [Theorem 4.7](#) there exists a subsequence T^{n_p} , $p = 1, 2, \dots$, that converges (in the product discrete topology) to a spanning tree T^* with associated basic feasible flow x^{T^*} . By conditions verified in [line 9](#), every tree T^{n_p} has the property that all arcs in layer A_{n_p} have nonnegative reduced costs with respect to T^{n_p} . The proof can be completed with the following three claims.

Claim 1. Arcs not in T^* have nonnegative reduced costs with respect to T^* .

We establish the claim by contradiction. Suppose there exists an (i, j) not in T^* with reduced cost $\bar{c}_{ij}^* < 0$, and let ℓ be such that A_ℓ is the smallest layer of arcs that contains (i, j) . To emphasize dependence of reduced costs on the specific basis, we will denote the reduced cost of (i, j) with respect to T^{n_p} by \bar{c}_{ij}^p .

Note that there exists an index P such that for $p \geq P$, $(i, j) \notin T^{n_p}$ (due to the locking in of the arcs in a convergent sequence of trees), and $\bar{c}_{ij}^p \geq 0$ (this happens as soon as $n_p \geq \ell$). Since the reduced cost of a non-basic arc can be calculated by considering the costs of all arcs in the cycle created by adding this arc to the tree, it will be convenient to denote the cycles generated by adding (i, j) to trees T^* and T^{n_p} by C_\star and C_p , respectively. We can take P to be large enough so that, in addition to the two properties above, for $p \geq P$, $A_\ell \cap C_\star = A_\ell \cap C_p$, i.e., cycles C_\star and C_p coincide within layer ℓ (this happens as soon as all arcs within layer ℓ lock in).

Let arc (i, j) determine the direction of cycles C_\star and C_p , and recall that C^F and C^B denote, respectively, the sets of forward and backward arcs of a cycle C . Then

$$\bar{c}_{ij}^* = \sum \{c_a : a \in C_\star^F\} - \sum \{c_a : a \in C_\star^B\} \quad (4)$$

$$= \sum \{c_a : a \in C_\star^F \cap A_\ell\} - \sum \{c_a : a \in C_\star^B \cap A_\ell\} \quad (5)$$

$$+ \sum \{c_a : a \in C_\star^F \setminus A_\ell\} - \sum \{c_a : C_\star^B \setminus A_\ell\} \quad (6)$$

and, for $p \geq P$,

$$\bar{c}_{ij}^p = \sum \{c_a : a \in C_p^F\} - \sum \{c_a : a \in C_p^B\} \quad (7)$$

$$= \sum \{c_a : a \in C_p^F \cap A_\ell\} - \sum \{c_a : a \in C_p^B \cap A_\ell\} \quad (8)$$

$$+ \sum \{c_a : a \in C_p^F \setminus A_\ell\} - \sum \{c_a : C_p^B \setminus A_\ell\}. \quad (9)$$

Note that both sums above are absolutely convergent by [Assumption 2\(iv\)](#) and, for $p \geq P$, expressions (5) and (8) coincide. Therefore, for $p \geq P$:

$$\begin{aligned} \bar{c}_{ij}^* - \bar{c}_{ij}^p &= \sum \{c_a : a \in C_\star^F \setminus A_\ell\} - \sum \{c_a : C_\star^B \setminus A_\ell\} \\ &\quad - \left(\sum \{c_a : a \in C_p^F \setminus A_\ell\} - \sum \{c_a : C_p^B \setminus A_\ell\} \right). \end{aligned}$$

If C_\star is a finite cycle, then the above becomes zero once all the arcs in C_\star lock in. Otherwise, since by [Assumption 2\(iv\)](#) $c \in \ell_1(\mathcal{A})$, each of the individual sums above goes to zero as ℓ goes to infinity, since it can be interpreted as a tail sum of an ℓ_1 sequence. This implies that $\bar{c}_{ij}^* = \lim_{p \rightarrow \infty} \bar{c}_{ij}^p \geq 0$, contradicting our assumption that $\bar{c}_{ij}^p < 0$ and proving [Claim 1](#).

Claim 2. $x^{T^{np}}$ converge to x^{T^*} as $p \rightarrow \infty$ in the relative product topology σ defined in [Section 2](#).

The proofs of this claim and [Claim 1](#) use similar logic, relying on the fact that trees T^* and T^p coincide on any finite set of arcs for sufficiently large p .

Relative topology σ on the topological space X described in [Section 2](#) is a topology of point-wise convergence, so we will argue that, for any arc $a \in \mathcal{A}$, $x_a^{T^{np}}$ and $x_a^{T^*}$ coincide for sufficiently large p (i.e., arc flows also “lock in”). Note that $x^{T^{np}}$ for any p , as well as x^{T^*} , belong to X , since they are feasible flows.

Any arc $a \notin T^*$ has $x_a^{T^*} = 0$; for sufficiently large p , this arc is also non-basic in T^{np} , and thus $x_a^{T^{np}} = 0$.

Consider now an arc $a \in T^*$, which belongs to all T^{np} for sufficiently large p . Consider the set of all arcs of G adjacent to T^* -predecessors of a (i.e., nodes i such that a is contained in the unique path from i to infinity in T^*). This is a finite set of arcs, which become locked in T^{np} for large enough p . Therefore, for sufficiently large p , a node i is a T^* -predecessor of a if and only if it is a T^{np} -predecessor of a . Since the flow on a in any T^{np} is equal to the sum of supplies at its tree-predecessors (cf. [Procedure 2](#)), $x_a^{T^{np}} = x_a^{T^*}$ for sufficiently large p , establishing [Claim 2](#).

Claim 3. The sequence of basic feasible solutions $x^{T^{np}}$ converges in value to $Z(x^{T^*})$.

As we have established in [Claim 2](#), the sequence $x^{T^{np}}$, $p = 1, 2, \dots$ converges to x^{T^*} in relative topology σ on X . Since $Z(\cdot)$ is continuous over X in this topology ([Theorem 2.9](#)), [Claim 3](#) follows immediately.

We can now combine the three claims above to complete the proof of the theorem. From [Theorem 3.6](#) and [Claim 1](#), $x^* = x^{T^*}$ is an optimal basic feasible solution; thus, by [Claim 2](#) a subsequence of iterates x^m of the algorithm converges to an optimal solution. Moreover, $Z(x^m)$ is non-decreasing in m (via [Theorem 4.6](#)), and [Claim 3](#) established convergence of values of a subsequence of these iterates. Hence, the entire sequence must converge in value to Z^* . \square

In the course of the above proof we established the following result, which we highlight as a corollary.

Corollary 4.9. Every convergent subsequence of bases (trees) T^m generated by [Algorithm 1](#) converges (in the product discrete topology) to an optimal spanning in-tree rooted at infinity.

We now show that the iterates of the simplex method become “arbitrarily close” to the set of the optimal solutions in the following sense. The set $\mathbb{R}^{\mathcal{A}}$ with the product topology is metrizable with a metric d [[3](#), [Theorem 16.2](#)]. This metric is inherited by X and its topology σ . Recall that

a sequence y^n converges to y in a metric space if $d(y^n, y) \rightarrow 0$ as $n \rightarrow \infty$. The distance from a point y to a set S is denoted $d(y, S) := \inf \{d(y, s) : s \in S\}$. We say a sequence y^n gets *arbitrarily close* to S if $d(y^n, S) \rightarrow 0$ as $n \rightarrow \infty$.

Theorem 4.10. The sequence of simplex iterates gets arbitrarily close to the set of optimal flows of (P).

Proof. Let F^* denote the set of optimal flows. Suppose there exists a subsequence x^{m_k} of simplex iterates and an $\epsilon > 0$ such that $d(x_k^{m_k}, F^*) > \epsilon$ for all k . By the proof of the previous theorem there exists a convergent subsequence that converges to an optimal feasible flow $x^* \in F^*$. However, this contradicts the supposition that $d(x_k^{m_k}, F^*) > \epsilon$ for all k . \square

5 All-to-infinity shortest path and dynamic programming

In this section we provide two examples of problems that fit our framework — the all-to-infinity shortest path problem, and deterministic nonstationary infinite horizon dynamic programming.

5.1 All-to-infinity shortest path problem

To motivate the all-to-infinity shortest path problem we first consider its *finite* counterpart: the all-to-one shortest path problem (see [5]). Given a *finite* directed graph with n nodes, where each arc (i, j) has length c_{ij} (which may be negative), the goal is to determine the shortest directed path to a designated node, say, n , from every node $i < n$, where the length of a directed path is the sum of the lengths of its arcs. This problem can be formulated as an uncapacitated network flow problem by assigning supply $b_i = 1$ to all nodes $i < n$, and $b_n = -(n - 1)$ to the destination node n . An optimal basis in this problem is a spanning in-tree rooted at n consisting of all-to-one shortest paths to node n .

For the infinite version, we are given a graph satisfying [Assumption 1](#) and [Assumption 3](#), with costs (arc lengths) that satisfy [Assumption 4](#). A natural counterpart of the above finite problem is the *all-to-infinity* shortest path problem, in which we seek to find a shortest directed path from each node to the virtual node at infinity.

Assign a supply of 1 to all nodes in the graph. The resulting network then satisfies [Assumption 2](#) and thus is a pure supply network, and our simplex method applies. According to [Corollary 4.9](#), the limit of every convergent subsequence of T^n is an optimal spanning in-tree rooted at infinity, and corresponds to an in-tree of all-to-infinity shortest paths.

[Assumption 2\(iv\)](#) is fairly natural here since we optimize over lengths of paths to infinity, and these lengths should be summable. The staging can correspond to some notion of “time” or “precedence,” depending on the application. As an illustration, we describe below a special case of the all-to-infinity shortest path problem where discounted costs are naturally interpreted in terms of discounting over time.

5.2 Non-stationary infinite horizon dynamic programming problem

Consider a deterministic non-stationary infinite horizon dynamic programming problems (DP) with finite state and action spaces. A system evolves over time periods $t = 0, 1, 2, \dots$ to be in one of finitely many states $s_t \in S_t$ each period. An action $a_t \in A_{s,t}$ is chosen from a finite set in each period and each state, and the system transitions to a new state according to the law

$s' = \tau_t(s_t, a_t) \in S_{t+1}$ yielding an immediate reward $r_t(s, a, s')$. We consider a non-stationary version of the problem, where state sets, available actions, immediate rewards, and transition laws all depend on t . The goal is to determine a *closed-loop* policy for maximizing total reward, i.e., a decision rule $d(\cdot, t)$, $t = 0, 1, \dots$, that prescribes the choice of action $d(s, t) \in A_{s,t}$ for every $s \in S_t$, to maximize the infinite sum of rewards starting in state s at time 0.

This problem can be formulated as a CINF in the following network. Consider a graph with nodes (s, t) where $t = 0, 1, 2, \dots$ and $s \in S_t$, each with supply of 1, and arcs $((s, t), (s', t+1))$ between pairs of nodes for which there exists $a \in A_{s,t}$ such that $\tau_t(s, a) = s'$ with costs $-r_t(s, a, s')$.

The resulting network naturally satisfies Assumptions 1 and 2(i–iii). Assumption 3 holds since states are only included in the set S_t for $t > 0$ if they are reachable from some state in S_0 by utilizing some sequence of actions. The leaves of the graph are precisely the nodes $(s, 0)$ where s is in finite set S_0 . Moreover, the t -th stage in this network consists precisely of nodes (s, t) , $s \in S_t$ — a re-interpretation of notation used in Section 2 — and the arcs connect nodes in t -th stage to nodes in $(t+1)$ -st stage. Cost structure of Assumption 4(i) is frequently assumed in DPs where maximum *total discounted reward* is sought, and Assumption 4(ii) is also commonly made (in fact, in many applications S_t is the same for every t). This results in a pure-supply network flow problem and our simplex method applies. An optimal basis is a spanning in-tree rooted at infinity that corresponds to a decision rule, with the outgoing tree arc for node (s, t) associated with the chosen action in the state s at time t .

Finally, observe that this class of countably-infinite network flow problems does not meet the sufficient criteria used in the analysis in [21]. For example, in their Proposition 2.5, they require the existence of not only explicit upper bounds on the flows, but also a uniform bound on the total flow between layers in the graph. In this formulation of DP, the flow between stages is increasing at least linearly in t , so there can be no uniform bound on the flow between layers, for any choice of root node.

6 Infinite-horizon dynamic lot sizing problem

In this section we consider another application that lends itself to analysis similar to that of Sections 2–4: the infinite-horizon version of the classic dynamic lot sizing problem in the special case of linear costs. Although the problem has many properties of the pure supply problem (more precisely, the pure demand problem of remark 2.12), it is not a special case of that class, and requires separate analysis.

6.1 Problem formulation

The finite-horizon version of the dynamic lot sizing dates back to [22]. There are N periods of demand d_t for $t = 1, \dots, N$ for a nonperishable product. In each period t , the decision-maker decides the amount of product to produce, x_t , and the amount of inventory, I_{t+1} , to send forward to period $t + 1$. We consider the linear cost version of the problem, where the cost of production in period t is $c_t x_t$, and the cost of holding inventory from period $t - 1$ to t is $h_t I_t$. We assume that the values of d_t , c_t , and h_t are all nonnegative and no backlogging is allowed, so demand in each period needs to be met with a combination of inventory available at the beginning of the period and production during the period. The problem can then be formulated as a finite network flow problem on the network in Figure 1 (adapted from [7]). Each node $t = 1, \dots, N$ in the horizontal




Figure 1: Illustration of the network flow representation of a finite-horizon dynamic lot sizing problem.

row represents a period with demand d_t , and node 0 represents an auxiliary “source” node with supply $S = \sum_{t=1}^N d_t$. The cost of arc $(0, t)$ (which has flow x_t) is c_t , and the cost of arc $(t - 1, t)$ (which has flow I_t) is h_t .

In the infinite-horizon setting, the problem has the same structure but now with $t = 1, 2, \dots$. Total demand for the product is infinite, and there is no a priori bound on production in any one period, and so we cannot use an auxiliary source node (or nodes) with finite supply in this setting. However, the infinite network in [Figure 2a](#) can be used to model the infinite dynamic lot sizing problem as a countably-infinite network flow problem (see [18] for additional discussion). Here, nodes $t = 1, 2, \dots$ have demand d_t , and auxiliary nodes $s_{t,k}$ for $t = 1, 2, \dots$ and $k = 1, 2, \dots$ are transshipment nodes; this network structure corresponds to having a production node “at infinity” that can supply demand nodes without there being any explicit supply nodes in the graph. (We assume for simplicity that the initial inventory is 0.) The cost structure is as follows: arcs $(t - 1, t)$ have cost h_t , arcs $(s_{t,1}, t)$ have cost c_t , and the remaining arcs $(s_{t,k+1}, s_{t,k})$ have cost 0. For each t , the flow x_t on arc $(s_{t,1}, t)$ and arcs $(s_{t,k+1}, s_{t,k})$ for $k = 1, 2, \dots$ is the production in period t and the flow I_t on arc $(t - 1, t)$ is the inventory carried from period $t - 1$ to t (for $t \geq 2$). We denote a feasible solution to the dynamic lot sizing problem by (x, I) , and calculate its cost as $Z(x, I) := \sum_{t=1}^{\infty} c_t x_t + \sum_{t=2}^{\infty} h_t I_t$, noting that, without further assumptions, $Z(x, I)$ may not be well-defined or finite.

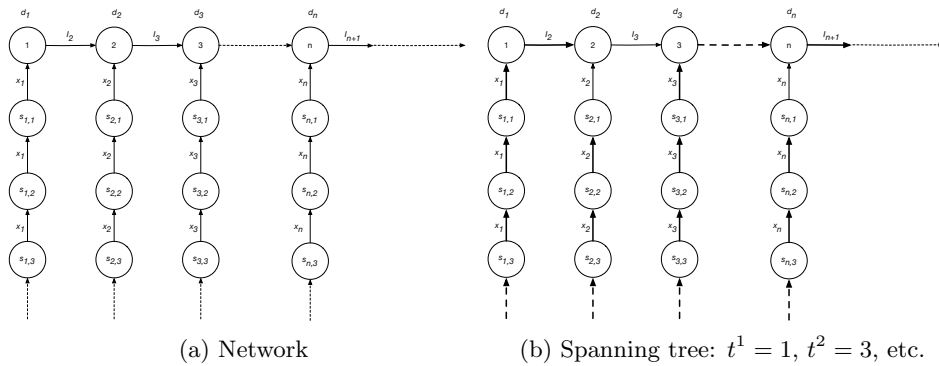


Figure 2: An infinite-horizon dynamic lot sizing problem: (a) network representation; (b) bold arcs form a spanning tree.

Assumption 5. The following hold:

- (i) Demands d_t are integer, and $d \in \ell_\infty(\mathbb{R}_+)$: $0 \leq d_t \leq D$ for all t for some $D > 0$;
- (ii) The remaining demand after time t (that is, $\sum_{s=t}^\infty d_s$) is infinite for all t ;
- (iii) Costs are “bounded with discounting”; i.e., there exist a $\beta \in (0, 1)$ and $\gamma, \delta > 0$, such that $0 \leq c_t \leq \gamma\beta^t$ and $0 \leq h_t \leq \delta\beta^t$ for all t ;
- (iv) For any t such that $c_t = 0$, $h_{t'} > 0$ for some $t' > t$.

Assumption 5(ii) is included so that the problem does not reduce to the finite case. **Assumption 5**(iv) is also quite natural: it precludes the pathological situation where, for some t , production is free, and there is no cost to holding inventory in any period afterwards, i.e., all demand in all future periods is met for free by infinitely-large production in period t .

The countably-infinite network flow formulation of the dynamic lot sizing problem is uncapacitated, and therefore does not satisfy the assumptions of [21]. We cannot bound feasible flows or their costs implicitly either: consider a sequence of feasible flows (x^n, I^n) where production in period 1 is used to meet demand up to period n (that is, $x_1^n = \sum_{t=1}^n d_t$) and $x_t^n = d_t$ for $t > n$. Clearly, as $n \rightarrow \infty$ the cost of feasible flow (x^n, I^n) tends to infinity. This simple example shows that assumption of [21, Proposition 2.5] is not satisfied either.

This network also violates assumptions of **Section 2**: note that the graph in **Figure 2a** contains directed cycles at infinity consisting of vertical arcs at any t together with horizontal arcs thereafter. Our analysis of the pure supply case used the absence of directed cycles at several key points: e.g., definition of stages (**Assumption 4**), finite objective value of feasible solutions (**Lemma 2.5**), compactness of the feasible region (**Lemma 2.7**), continuity of the objective function (**Theorem 2.9**), and convergence of spanning trees (**Theorem 4.7**). However, as we show in this section, analogous results can be established using different techniques.

6.2 Spanning trees and extreme point solutions

Under **Assumption 5**(i), **Theorem 3.2** implies that every extreme point of the feasible region of this network flow problem is a basic feasible solution, and vice versa. These basic feasible solutions are associated with spanning trees of the following structure:

- (i) there is an infinite increasing sequence of *production periods* t^j , $j = 1, 2, \dots$, with $t^1 = 1$,
- (ii) arcs $(s_{t,1}, t)$ (i.e., production arcs) are basic (i.e., in the tree) for $t = t^j$, $j = 1, \dots$,
- (iii) arcs $(t-1, t)$ (i.e., inventory arcs) are basic, *except* for $t = t^j$, $j = 1, 2, \dots$, and
- (iv) arcs $(s_{t,k+1}, s_{t,k})$, $k \geq 1$ are basic for all periods t (although they have flow 0 for non-production periods).

Setting $t^1 = 1$ ensures the tree is indeed spanning. Moreover, the sequence $\{t^j\}$ characterizing the spanning tree has to be infinite (otherwise, if N is the index of the largest production period, there is an infinite cycle consisting of the nodes $\dots, s_{N,3}, s_{N,2}, s_{N,1}, N, N+1, N+2, \dots$). **Figure 2b** illustrates a portion of a spanning tree where $t^1 = 1$ and $t^2 = 3$ (basic arcs are in bold). These observations establish the following:

Lemma 6.1. Every spanning tree is an out-tree rooted at infinity and associated with a basic feasible flow.

Another implication of this structure is as follows:

Lemma 6.2. If [Assumption 5](#) holds, and (x, I) is a basic feasible flow, then $x_t I_t = 0$ for all t . In other words, in every basic feasible flow, products are produced in period t only if there is zero incoming inventory into period t .

Proof. Suppose $x_t > 0$. If $I_t > 0$ then there must exist a previous period t' where $x_{t'} > 0$ and $I_{t'+1}, I_{t'+2}, \dots, I_{t-1} > 0$. However, this creates an infinite (undirected) cycle of active arcs, namely those connecting the nodes $\dots, s_{t',2}, s_{t',1}, t', t' + 1, \dots, t - 1, t, s_{t,1}, s_{t,2}, \dots$. Since active arcs must be basic, this violates the assumption that (x_t, I_t) is a basic feasible flow. \square

Remark 6.3. The property $x_t I_t = 0$ for all t is often called the *Wagner–Whitin property* (see [\[22\]](#) for its derivation for an optimal solution in the finite case). In the finite case the same is true when the objective function is, more generally, concave; for instance, in [\[22\]](#) production cost includes a fixed cost. In the finite case, it is well known that minimizing a concave function gives rise to optimal extreme point solutions which possess the Wagner–Whitin property. In the infinite case, the feasible region needs to be compact and objective function needs to be continuous in order to leverage the analogous result via Bauer’s minimum theorem (Theorem 7.69 in [\[3\]](#)). We treat these topics in the infinite setting later in [Section 6.4](#).

In light of the above tree structure, a basic feasible flow has $x_{t^j} = \sum_{i=t^j}^{t^{j+1}-1} d_i$ for all j , and $x_i = 0$ for $i = t^j + 1, \dots, t^{j+1} - 1$, and so we say that production period t^j *covers demand* in periods t^j to $t^{j+1} - 1$. The amount of inventory at the beginning of period t^j is $I_{t^j} = 0$, and the amount of inventory I_t for $t \neq t^j$ can be determined by reducing the amount of production in the most recent production period by the demands met at all the interceding periods.

6.3 A simplex method

We can apply the layered simplex method ([Algorithm 1](#)) adapted to work with out-, rather than in-trees, to the network flow formulation of the dynamic lot sizing problem. In fact, we show below that only a finite amount of computation is needed *per iteration* of the algorithm, in contrast to the general case.

We can initialize the algorithm with the flow (x^0, I^0) where $x_t^0 = d_t$ and $I_t^0 = 0$ for all t . Clearly, this flow is feasible. Moreover, it has finite cost under [Assumption 5](#):

$$Z^0 = Z(x^0, I^0) = \sum_{t=1}^{\infty} c_t d_t \leq \sum_{t=1}^{\infty} CD\beta^t = \frac{CD\beta}{1-\beta} < +\infty. \quad (10)$$

We now turn to the **while** loop in the algorithm (starting with [line 4](#)). Let the current spanning tree be specified by the infinite sequence t^j , $j = 1, 2, \dots$ with $t^1 = 1$. Note that there are two types of non-basic arcs:

Case 1: The first type of a non-basic arc is a production arc $(s_{t,1}, t)$ for $t^{j-1} < t < t^j$, where t^{j-1} is the production period that covers demand in period t (e.g., arc $(s_{2,1}, 2)$ in the graph in [Figure 2b](#)). Adding this arc to the tree and computing the cost of the resulting infinite cycle, we can compute the reduced cost of the arc as $c_t - c_{t^{j-1}} - \sum_{i=t^{j-1}+1}^t h_i$. If this arc is chosen to enter the basis, corresponding inventory arc $(t-1, t)$ will leave.

Case 2: The second type of a non-basic arc is an inventory arc $(t^j - 1, t^j)$ for some $j > 1$ (e.g., arc $(2, 3)$ in the graph in [Figure 2b](#)). Adding this arc to the tree and computing the cost of the

resulting infinite cycle, we can compute the reduced cost of the arc as $c_{t^j-1} - c_{t^j} + \sum_{t=t^j-1+1}^{t^j} h_t$. If this arc is chosen to enter the basis, corresponding production arc $(s_{t^j,1}, t^j)$ will leave.

In both cases, the reduced cost calculation can be performed in finite time using only requires information on inventory and production costs up to period t^j .

Lemma 6.4. Suppose (x, I) is a basic feasible solution with finite objective function value. After a pivot, the resulting basic feasible solution will also have finite objective function value.

Proof. The change in objective function value after a pivot can be computed as the product of the reduced cost of the entering arc (clearly finite according to the above discussion) and the flow on the leaving arc (also finite). \square

We conclude this subsection by collecting properties of the simplex algorithm that carry over to this setting from [Section 4](#). Indeed, a careful examination of the results in that section reveals that the assumptions violated in the dynamic lot sizing problem all stem from the presence of infinite directed cycles in the graph, and are not relevant for establishing the results captured here.

Theorem 6.5. Consider the layered simplex method for the infinite horizon dynamic lot sizing problem. Then (i) the algorithm does not cycle (cf. [Theorem 4.4](#)), (ii) each layer is eventually reached as the algorithm proceeds ($n \rightarrow \infty$) (cf. [Theorem 4.6](#)), (iii) costs of successive simplex iterates are nonincreasing (cf. [Theorem 4.6](#)), and (iv) all iterates of the algorithm are spanning out-trees rooted at infinity, and a subsequence of iterates converges to T^* , which is a spanning out-tree rooted at infinity with nonnegative reduced costs on all arcs (cf. [Claim 1](#)).

6.4 A compact representation

The remainder of the argument to establish optimal value convergence roughly follows the proof of [Theorem 4.8](#). We first need to establish compactness of the feasible region (cf. [Lemma 2.7](#)) and continuity of the objective (cf. [Theorem 2.9](#)). These results establish termination conditions for the simplex method: first, that an extreme point optimal solution exists, and second, that an optimal tree is characterized by nonnegative reduced costs ([Theorem 3.6](#)).

In [Section 2.3](#) we established existence of an optimal extreme point before studying the simplex algorithm. Here, we take the opposite approach, using properties of the simplex method to argue that an extremal optimal solution exists.

Recall that the simplex method was initialized with the basic feasible solution (x^0, I^0) with finite cost Z^0 (defined in [\(10\)](#)).

Lemma 6.6. The objective value of a feasible flow is either a nonnegative real number or $+\infty$. The optimal objective value Z^* is a nonnegative real number.

Proof. By [Assumption 5](#), the terms in the sum $Z(x, I) = \sum_{t=1}^{\infty} c_t x_t + \sum_{t=2}^{\infty} h_t I_t$ are nonnegative for any feasible flow. Thus, the sum is either a nonnegative real number or $+\infty$. Since $Z^* \leq Z(x^0, I^0) < +\infty$, the former holds. \square

We add the following additional constraint to our formulation:

$$\sum_{t=1}^{\infty} c_t x_t + \sum_{t=2}^{\infty} h_t I_t \leq Z^0 + 1, \quad (11)$$

and let F^0 denote the new feasible region. We make several observations about F^0 .

First, constraint (11) puts implicit bounds on all flows. Indeed, if $c_t > 0$ then from (11) we have $c_t x_t \leq Z^0 + 1$, and thus, $x_t \leq (Z^0 + 1)/c_t$. Similarly, if $h_t > 0$ then $I_t \leq (Z^0 + 1)/h_t$. If $c_t = 0$ for some t , there exists $t' > t$ with $h_{t'} > 0$ by [Assumption 5\(iv\)](#). Let $D = \sum_{\tau=t}^{t'-1} d_\tau$. If $x_t > D$, we have $I_{t'} \geq x_t - D > 0$ and therefore $x_t \leq D + (Z^0 + 1)/h_{t'}$. A similar bound on I_t can be derived when $h_t = 0$. This allows us to conclude that feasible region F^0 is compact in the product topology — a result analogous to [Lemma 2.7](#).

Second, all extreme points of the original problem with lower costs than (x^0, I^0) satisfy constraint (11) strictly. Therefore, such flows are extreme points of F^0 .

Third, the set F^0 may have extreme points that are not among the extremal flows of the original problem. However, (11) must be active at these extremal flows; that is, their cost must be greater than that of (x^0, I^0) .

Since the simplex method visits extreme points with nonincreasing costs ([Theorem 6.5\(iii\)](#)), this implies only extreme points of F^0 that were also extreme points of the original problem are visited. Navigating among improving extreme points of F^0 is equivalent to applying the simplex method to the original problem with initial basic feasible flow (x^0, I^0) . Thus, [Theorem 2.9](#) and [Corollary 2.10](#) can be adapted to the dynamic lot-sizing problem (by posing them over F^0 and using the relative topology of the product topology on that set) to conclude that objective function Z is continuous over F^0 and an extreme point optimal solution exists.

Since the formulation has a compact feasible region and a continuous objective function, we can adapt the development of [Section 3](#) to conclude that an optimal basis is characterized by nonnegative reduced costs ([Theorem 3.6](#)). Combining this with the conclusions of [Theorem 6.5\(iv\)](#) allows us to conclude that the limit tree T^* gives rise to an optimal feasible flow x^{T^*} . This establishes the following.

Theorem 6.7 (c.f. [Theorem 4.8](#)). Consider the infinite horizon dynamic lot sizing problem under [Assumption 5](#). The iterates (x^m, I^m) of [Algorithm 1](#) with initial basic feasible flow (x^0, I^0) converge in value to optimality. Moreover, there exists a subsequence of the simplex iterates that converges (in the topology of F^0) to an extremal optimal flow.

7 Duality

We return to the pure supply setting of [Sections 2 to 4](#) and provide a proof of strong duality using an argument based on the simplex method. Our argument is based on the strong characterization we have for an optimal tree and corresponding feasible flow, as summarized in [Theorem 3.6](#) and [Corollary 4.9](#).

We propose the following dual problem to associate with (P):

$$D^* = \sup_{\pi} D(\pi) := \sum_{i \in \mathcal{N}} b_i \pi_i \tag{12a}$$

$$(D) \quad \text{s.t. } \pi_i - \pi_j \leq c_{ij} \text{ for } (i, j) \in \mathcal{A} \tag{12b}$$

$$\pi \in c_0, \tag{12c}$$

where c_0 is the space of null sequences: $c_0 = \{\pi \in \mathbb{R}^\infty : \lim_{i \rightarrow \infty} \pi_i = 0\}$. The reason for this choice of dual space is that the transversality condition of [\[17\]](#) is then satisfied by construction. Note that c_0 (which contains π) and ℓ_∞ (which contains b) are not topological dual vector spaces (indeed, the

topological dual of c_0 is ℓ_1) and so weak duality is not a consequence of existing theory (e.g., [4]) and must be derived using specialized arguments.

In [21], the authors provide an example of a network such that the dual problem obtained by simple application of the formulation procedure familiar from the finite case (i.e., omitting space restrictions on the dual variables (12c)) did not produce a strong, or even a weak, dual problem. The uncapacitated version of their example is as follows: consider a graph with nodes $i = 1, 2, \dots$ and arcs $(i, i + 1)$, $i = 1, 2, \dots$. Consider a supply of 1 at node 1, and 0 elsewhere, and let the cost of arc $(i, i + 1)$ be $(1/2)^i$. The only feasible primal solution is $x_{i,i+1} = 1$ for all i , with the cost $Z^* = 1$. There is a variety of dual solutions satisfying (12a) and (12b). For example, we could take $\pi_i = \lambda$ for all i , for any number λ (this is the solution suggested in [21]); the objective value of this solution is λ , which can be made arbitrarily large or small.

What is at issue here is that, in the infinite case, total supply does not necessarily have to equal total demand, since the virtual node at infinity can serve as an infinite source or an infinite sink, provided the network has appropriate topology. To take advantage of this interpretation, define basic dual solutions as follows. Suppose T is a spanning tree, and x^T is the corresponding primal basic solution. For each node i , let $P_{i\infty}$ be the unique path from i to infinity in T , and let $P_{i\infty}^F$ and $P_{i\infty}^B$ be the set of forward and backward arcs in this path, respectively. We then define

$$\pi_i^T = Z(P_{i\infty}) := \sum_{(k,j) \in P_{i\infty}^F} c_{k,j} - \sum_{(k,j) \in P_{i\infty}^B} c_{k,j}. \quad (13)$$

π_i^T , which is well-defined since $c \in \ell_1$, can be interpreted as the total cost of the path from i to infinity in T , taking into account arc directions. (In the above example, this would correspond to taking $\pi_1 = 1$.) x^T and π^T are complementary, since $\pi_i^T - \pi_j^T = c_{ij}$ for $(i, j) \in A(T)$. Observe that $\pi^T \in c_0$ since $c \in \ell_1(\mathcal{A})$ and tail sums of a summable sequence converge to 0. It is trivial to see that reduced costs are equal to slacks in constraints (12b) and thus π^T is dual-feasible if and only if all reduced costs with respect to T are nonnegative.

The next lemma shows that x^T and π^T associated with a spanning in-tree T rooted at infinity have the same objective function values.

Lemma 7.1. If T is a spanning in-tree rooted at infinity then $Z(x^T) = D(\pi^T)$.

Proof. Construct x^T via Procedure 2 and π^T via (13). Then, by the continuity of Z (Theorem 2.9), $Z(x^T) = \sum_{i \in \mathcal{N}} b_i \cdot Z(P_{i\infty}) = \sum_{i \in \mathcal{N}} b_i \pi_i^T = D(\pi^T)$. \square

To prove weak duality we establish the following two lemmas.

Lemma 7.2. Let π be a dual-feasible solution. Then for any i , $\pi_i \leq \sum_{(k,\ell) \in P_{i\infty}^{\rightarrow}} c_{k\ell}$ for any directed path $P_{i\infty}^{\rightarrow}$ from i to infinity.

Proof. Let $j_1 = i, j_2, j_3, \dots$ be the sequence of nodes forming $P_{i\infty}^{\rightarrow}$, and let us denote $c_{j_t, j_{t+1}}$ by c_t to simplify notation. Dual feasibility of π^T implies that $\pi_{j_t} \leq c_t + \pi_{j_{t+1}}$ for all $t \geq 1$. Invoking this inequality k times reveals

$$\pi_i \leq \sum_{k=1}^K c_k + \pi_{j_{K+1}} \quad (14)$$

for all $K \geq 1$. Taking the limit as $K \rightarrow \infty$ on both sides of (14) yields $\pi_i \leq \sum_{k=1}^{\infty} c_k + \lim_{K \rightarrow \infty} \pi_{j_{K+1}}$. The second limit on the right-hand side is 0 since $\pi \in c_0$. Hence, $\pi_i \leq \sum_{k=1}^{\infty} c_k = \sum_{(k,\ell) \in P_{i\infty}^{\rightarrow}} c_{k,\ell}$, as required. \square

Theorem 7.3 (Weak duality). In a pure-supply network, every primal feasible x and dual feasible π satisfy $Z(x) \geq D(\pi)$.

Proof. By [Corollary 4.9](#), there exists an optimal spanning in-tree rooted at infinity, T^* , and the corresponding optimal basic feasible flow x^* . Let $P_{i\infty}^*$ denote the directed path to infinity from $i \in \mathcal{N}$ in T^* . Recall that we can write $Z(x^*) = \sum_{i \in \mathcal{N}} b_i Z(P_{i\infty}^*)$. By [Lemma 7.2](#), we know that $\pi_i \leq Z(P_{i\infty}^*)$ for every dual feasible π , which, combined with expression above, yields $Z(x^*) = \sum_{i \in \mathcal{N}} b_i Z(P_{i\infty}^*) \geq \sum_{i \in \mathcal{N}} b_i \pi_i = D(\pi)$. Since $Z(x) \geq Z(x^*)$ for all feasible flows x this immediately implies $Z(x) \geq D(\pi)$ for all feasible x and π . \square

Theorem 7.4 (Strong duality). There are optimal solutions x^* to (P) and π^* to (D) such that $Z(x^*) = D(\pi^*)$.

Proof. Let T^* be an optimal spanning in-tree rooted at infinity, $x^* = x^{T^*}$ and $\pi^* = \pi^{T^*}$. We know that x^* is a primal-feasible solution, $Z(x^*) = D(\pi^*)$ (via [Lemma 7.1](#)), and π^* satisfies [\(12b\)](#).

To conclude that π^* is a feasible dual solution, it remains to show that $\pi^* \in c_0$. Label the nodes and the arcs in accordance with the stages in the graph, as discussed in [Section 2](#) (since reordering of the elements of a sequence does not have an effect on its limit). Recall that $c \in \ell_1$, i.e., $\sum_{a=1}^{\infty} |c_a| < \infty$. This implies that tail sums $\sum_{a=k}^{\infty} |c_a|$ converge to 0 as $k \rightarrow \infty$.²

For any $i \in \mathcal{N}$, $|\pi_i^*| = |Z(P_{i\infty}^*)|$ by [\(13\)](#). As i increases, the cost of the corresponding path is equal to the sum of costs of arcs connecting higher-numbered stages. In particular, since the number of nodes in each stage is finite, as $i \rightarrow \infty$, the stage labels, and hence the labels of the arcs included in the sum, also tend to infinity. In other words, $|\pi_i^*|$ can be bounded above by the tail sums $\sum_{a=k}^{\infty} |c_a|$ with $k \rightarrow \infty$ as $i \rightarrow \infty$, establishing the desired result. \square

Remark 7.5. Observe that the key to [Theorem 7.3](#) is that spanning tree T^* is an in-tree rooted at infinity. This same result (for out-trees) holds for the problem studied in [Section 6](#). Hence, the theory of this section can be adjusted in a straightforward way to provide weak and strong duality results for the infinite-horizon dynamic lot sizing problem as well.

8 Conclusion

In this paper we devised a simplex method for infinite network flow problems that addresses potential cycling in the degenerate case and has convergence guarantees, without relying on uniform capacity bounds. Our algorithm produces a sequence of monotone improving adjacent extreme points that converges in value to the optimum, and converges to an extreme point optimal solution on a subsequence. A variety of applied problems are amenable to this method. Our approach is “primal,” in contrast to all known previous results on countably-infinite linear programs (CILPs), that argue through analysis of a dual. Consequently, our approach provides new tools and insights to study CILPs.

There is scope for extending our results. Our simplex method is not necessarily *finitely implementable*. Calculating reduced costs in [line 5](#) may involve infinite cycles with infinite data. However, as demonstrated in our analysis of the dynamic lot sizing problem, it is possible that looking at only a finite amount of data could suffice for sufficiently structured problems. Uncovering this structure is an area for future investigation.

²Labeling arcs by the natural numbers respecting stages, as discussed in [Section 2.2](#) (preceding the statement of [Assumption 4](#)).

Extensions to more general network flow problems are also a topic of future work. For instance, including a mix of demand and supply nodes in a structured setting could allow modeling of infinite versions of matching and transportation problems. There are new challenges associated with this generalization, but we believe the core methodology developed in this paper can serve as a blueprint.

Acknowledgements

Christopher Thomas Ryan would like to thank the University of Chicago Booth School of Business for its generous research support. Robert Smith and Marina Epelman are supported in part by National Science Foundation grant CMMI-1333260.

References

- [1] R. Aharoni, E. Berger, A. Georgakopoulos, A. Perlstein, and P. Sprüßel. The max-flow min-cut theorem for countable networks. *J. Combin. Theory Ser. B*, 101(1):1–17, 2011.
- [2] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- [3] C.D. Aliprantis and K.C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer-Verlag, Berlin Heidelberg, 3rd edition, 2006.
- [4] E.J. Anderson and P. Nash. *Linear Programming in Infinite-Dimensional Spaces: Theory and Applications*. Wiley, New York, NY, 1987.
- [5] D. Bertsimas and J.N. Tsitsiklis. *Introduction to Linear Optimization*. Athena, Belmont, MA, 1997.
- [6] W.H. Cunningham. A network simplex method. *Math. Program.*, 11(1):105–116, 1976.
- [7] E.V. Denardo. *Dynamic Programming: Models and Applications*. Dover, Mineola, NY, 2003.
- [8] R. Diestel. *Graph Theory*. Springer-Verlag, New York, NY, 4th edition, 2010.
- [9] A. Ghate and R.L. Smith. Characterizing extreme points as basic feasible solutions in infinite linear programs. *Oper. Res. Lett.*, 37(1):7–10, 2009.
- [10] A. Ghate and R.L. Smith. A linear programming approach to non stationary infinite-horizon Markov decision processes. *Oper. Res.*, 61:413–425, 2013.
- [11] A. Ghate, D. Sharma, and R.L. Smith. A shadow simplex method for infinite linear programs. *Oper. Res.*, 58(4):865–877, 2010.
- [12] I. Lee, M.A. Epelman, H.E. Romeijn, and R.L. Smith. Simplex algorithm for countable-state discounted Markov decision processes. *Oper. Res.*, to appear, 2017.
- [13] J.P. Ponstein. On the use of purely finitely additive multipliers in mathematical programming. *J. Optim. Theory Appl.*, 33(1):37–55, 1981.

- [14] M.C. Pullan. A duality theory for separated continuous linear programs. *SIAM J. Control Optim.*, 34:931, 1996.
- [15] L. Rademacher, A. Toriello, and J.P. Vielma. On packing and covering polyhedra in infinite dimensions. *Oper. Res. Lett.*, 44(2):225–230, 2016.
- [16] H.E. Romeijn and R.L. Smith. Shadow prices in infinite-dimensional linear programming. *Math. Oper. Res.*, 23(1):239–256, 1998.
- [17] H.E. Romeijn, R.L. Smith, and J.C. Bean. Duality in infinite dimensional linear programming. *Math. Program.*, 53(1-3):79–97, 1992.
- [18] H.E. Romeijn, D. Sharma, and R.L. Smith. Extreme point characterizations for infinite network flow problems. *Networks*, 48(4):209–22, 2006.
- [19] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 3rd edition, 1976.
- [20] I.E. Schochetman and R.L. Smith. Finite dimensional approximation in infinite dimensional mathematical programming. *Math. Program.*, 54(1-3):307–333, 1992.
- [21] T.C. Sharkey and H.E. Romeijn. A simplex algorithm for minimum-cost network-flow problems in infinite networks. *Networks*, 52(1):14–31, 2008.
- [22] H.M. Wagner and T.M. Whitin. Dynamic version of the economic lot size model. *Management Sci.*, 5(1):89–96, 1958.