informs
http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 66, No. 4, July–August 2018, pp. 1058–1067
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Technical Note—The Competitive Facility Location Problem in a Duopoly: Advances Beyond Trees

Yonatan Gur,[a] **Daniela Saban,**[a] **Nicolas E. Stier-Moses**[b]

[a] Stanford University, Stanford, California, 94305; [b] Universidad Torcuato Di Tella, Buenos Aires (1428), Argentina

Contact: ygur@stanford.edu (YG); dsaban@stanford.edu (DS); nstier@utdt.edu, http://orcid.org/0000-0001-6380-1516 (NES-M)

Received: December 2, 2015
Revised: October 3, 2016; May 19, 2017
Accepted: August 3, 2017
Published Online in Articles in Advance: July 23, 2018

Subject Classifications: facilities/equipment planning: location: discrete; games/group decisions: noncooperative; networks/graphs: theory
Area of Review: Transportation

https://doi.org/10.1287/opre.2017.1694

Abstract. We consider a competitive facility location problem on a network where consumers located on vertices wish to connect to the nearest facility. Knowing this, each competitor locates a facility on a vertex, trying to maximize market share. We focus on the two-player case and study conditions that guarantee the existence of a pure-strategy Nash equilibrium for progressively more complicated classes of networks. For general graphs, we show that attention can be restricted to a subset of vertices referred to as the central block. By constructing trees of maximal bi-connected components, we obtain sufficient conditions for equilibrium existence. Moreover, when the central block is a vertex or a cycle (for example, in cactus graphs), this provides a complete and efficient characterization of equilibria. In that case, we show that both competitors locate their facilities in a solution to the 1-median problem, generalizing a well-known insight arising from Hotelling's model. We further show that an equilibrium must solve the 1-median problem in other classes of graphs, including grids, which essentially capture the topology of urban networks. In addition, when both players select a 1-median, the solution must be at equilibrium for strongly-chordal graphs, generalizing a previously known result for trees.

Funding: This work was partially supported by the Chazen Institute of International Business at Columbia Business School, by Conicet Argentina [Grant Resolución 4541/12] and by ANPCyT Argentina PICT-2012-1324.

Supplemental Material: The electronic companion is available at https://doi.org/10.1287/opre.2017.1694.

Keywords: competitive facility location • Hotelling competition • 1-median • Nash equilibrium • Voronoi game

1. Introduction

Facility location problems study how to best locate facilities, anticipating that consumers will be attracted by the facility that is most convenient to them. Such problems are related to several application domains that go beyond physical location of facilities; examples include voting (Taylor 1971, Hansen et al. 1990), rumor dissemination and seeding in social networks (Bharathi et al. 2007, Kostka et al. 2008), and product differentiation models (see overview in Tirole 1988).

While most previous research related to facility location problems considers a centralized optimization perspective, a significant fraction of it incorporates competition. In that case, players decide where to locate their facilities competing for market share. A common assumption has been to consider a continuous (usually linear) market in which consumers are distributed (in most cases, uniformly) over a low (mostly one) dimensional space (see, e.g., Eaton and Lipsey 1976, Salop 1979, de Palma et al. 1985, Gabszewicz and Thisse 1986, and the references therein). However, there have been relatively few attempts to study competitive location problems in more general topologies; we mention a few notable exceptions below.

Our results are motivated by the work of Hotelling (1929), who introduced an influential competitive facility location model in which each of two players selects a location in a linear segment, and consumers located uniformly in the segment select the closest location among the two. Assuming a common price, both players compete for the demand-maximizing location. The resulting equilibrium is that both competitors choose to locate their facilities in the center of the segment. This insight seems to extend to markets in the real world. For instance, it is common to observe in small towns or in neighborhoods of big cities that competing local businesses such as fast food chains, coffee shops, or banks are located one next to another. Gas stations tend to be located in opposite corners instead of being uniformly spaced in an area. A similar behavior is observed in product design selection, where many competing products in a market have similar design and comparable features. While other considerations besides spatial competition may contribute to such behavior, the purpose of this work is to offer a framework that can be used to formalize why competitors choose to cluster in close proximity to the "the center of the market" (to be formalized later), instead of being distributed throughout the network to serve the market more efficiently. We do this for a duopoly and progressively consider more complex underlying network topologies to show that these insights are quite robust.

We assume that the market is represented by a graph in which vertices represent both the possible locations of facilities and of demand. To the best of our knowledge, Wendell and McKelvey (1981) were the first to consider a graph version of the Hotelling model in which edges succinctly encode the possible consumer choices. Differently from us, they considered that facilities can be placed in the interior of the edges to keep players' choices continuous. (Slater 1975 and Hakimi 1983 had already considered a network context for a related model.) We study the pure-strategy Nash equilibria of this game and provide conditions that guarantee their existence for progressively more complicated network structures (the existence of a mixed-strategy equilibrium is guaranteed as the game at hand is zero sum with a discrete strategy space). We start with trees and cycles and continue to more general topologies, such as graphs whose *central block* is a vertex or cycle, as well as grids, median graphs, and strongly chordal graphs. Note that deciding if an instance of the game for a general number of players possesses a Nash equilibrium is $\mathcal{NP}$-complete, even for the case of vertices with equal demands (Dürr and Thang 2007).

Below we provide a brief overview of several research efforts related to location theory, focusing on those most relevant to our proposed model. For additional references, we refer the reader to overviews by Mirchandani and Francis (1990), Drezner (1995b), and Eiselt and Marianov (2011). Recognizing that the existence of Nash equilibria in location games cannot be guaranteed in general, several research efforts focused on alternative solution concepts and multiple variants of game dynamics. A commonly employed equilibrium concept is to consider a model where agents locate facilities sequentially (Hay 1976, Prescott and Visscher 1977). This leads to Stackelberg games where players anticipate the moves of their opponents and play accordingly. Hakimi (1983, 1986) presented an algorithmic approach to compute the Stackelberg equilibrium. Other variants of the problem include using utility models (Leonardi and Tadei 1984), distributing demand with a logit model (Drezner et al. 1998), applying a gravity rule (Drezner 1995a), and defining a sphere of influence to capture that demand that is far from facilities may be lost (ReVelle 1986). In a Voronoi game, players place a facility in a domain and obtain the area corresponding to points in the domain that are closest to the facility (Ahn et al. 2004). Vetta (2002) considers ideas from auction design by allocating demand to the nearest facility and setting the price to the production cost of the second nearest facility that can serve the customer. He proves that the game always has a pure-strategy Nash equilibrium and, in the absence of fixed costs, the social welfare at equilibrium is never less than half of the maximum social welfare. Mirrokni and Vetta (2004) and Goemans et al. (2005) follow this up by studying the speed of convergence of best response dynamics and providing a more detailed analysis of the inefficiency induced by that game. Jain and Mahdian (2007) provide an overview of related research from the point of view of cooperative game theory, where the cost of opening facilities must be shared by participating agents.

Main Contributions. Our main contribution lies in putting forward connections between spatial competition and centralized optimization. Motivated by the result of Eiselt and Laporte (1991), our work establishes further links between Nash equilibria of the facility location problem in a duopoly and the 1-median problem for various classes of topologies. A 1-median is a vertex that minimizes the weighted distance to all other vertices. Since it is natural for players to choose a central location in the market, 1-medians would seem to be obvious candidates to locate facilities at equilibrium. Nevertheless, this intuition is not always right as there are instances that admit equilibria not located at 1-medians. This provides the motivation to understand the circumstances under which the players' incentive is to select solutions to the 1-median problem.

We provide necessary and sufficient conditions for equilibrium existence for duopolies on cycles with general positive weights and show that any vertex belonging to an equilibrium must be a 1-median. (This relates to Mavronicolas et al. 2008, who studied cycles with the restriction that all vertices have unit demand and characterized all possible equilibria for an arbitrary number of players.) To provide results for more general graphs, we rely on Eiselt and Laporte (1991) who showed that an equilibrium of a facility location game with two players on a tree with arbitrary weights is always guaranteed to exist and it consists of facilities located arbitrarily in the 1-medians of the tree. Using this, we reduce an arbitrary graph to a tree of maximal bi-connected components, also referred to as *blocks*, and show that the equilibrium must be located in the block that corresponds to the 1-median of that tree (referred to as the *central block*). This provides a general characterization of equilibria for graphs whose central block has a topology for which we can describe equilibria (for example, graphs with a central block that is a vertex or a cycle, a property we refer to as having a *simple central block*).

Our analysis allows one to verify existence of equilibria in linear time with respect to the number of vertices for progressively more complicated topologies (trees, cycles, sparse simple central block graphs), which is at least a cubic speedup compared to a naive approach of computing equilibria by exhaustive search. For general graphs, we reduce the decision problem to that in a bi-connected component that can be identified in linear time with respect to the number of edges. Although the worst-case complexity of the problem is the same as in the general graph, this allows one to focus attention to the "appropriate" part of the graph, which is typically smaller. In a nutshell, checking if an instance admits equilibria reduces to checking if its central block admits one itself.

To characterize the location of equilibria further, we show for arbitrary topologies that if a location is selected at equilibrium, it must be a *local optima* of the 1-median problem. For graphs whose local and global optima coincide, this implies that equilibria can only be located at a 1-median. Some families for which this happens include median, quasimedian, Helly graphs, and grids and lattices. The last ones capture the topology of many real urban networks, supporting the empirical observation of competitors locating close to each other in urban environments.

Finally, for strongly chordal graphs, which include trees, interval, and block graphs, we prove that both players selecting an arbitrary 1-median guarantees that the solution is an equilibrium. This result generalizes for the first time the same statement for trees (Eiselt and Laporte 1991).

Structure of the Paper. After introducing the facility location game in Section 2, Section 3 characterizes the equilibrium locations in progressively more complex topologies, and presents a complete equilibrium characterization on simple central block graphs. In Section 4 we explore the relations between equilibria and 1-medians, and we conclude with design recommendations and directions of future research in Section 5. The proofs and algorithms appear in the electronic companion, which also includes a study of equilibrium efficiency, model extensions and variations, and simulations based on real-world network topologies.

2. The Competitive Facility Location Game

We consider a finite undirected connected graph $G(V, E)$ whose vertices represent the locations of consumers and the potential locations of facilities. Each vertex $v \in V$ has an associated weight $w(v) > 0$ that measures the demand level of that location. (While having some vertices with $w(v) = 0$ does not change existence results, additional equilibria may exist; we discuss the implications of allowing vertices with zero weights in Section 5.) For any subset of vertices $S \subseteq V$, we define $W(S) := \sum_{v \in S} w(v)$, and set $W := W(V)$. Every edge $e \in E$ has a unit length. This assumption is made for the sake of simplicity, and we do not lose much generality since we can always subdivide edges (see Section 5). The game is played between two players $N := \{1, 2\}$. Each player selects a vertex at which to locate a facility. (Section EC.2.3 in the electronic companion considers a model with two players placing two facilities each.)

Given vertices $x, y \in V$, we denote the set of x-y paths by P_{xy}. We refer to the distance between x and y by $d_G(x, y) := \min\{|p|: p \in P_{xy}\}$, or simply by $d(x, y)$. We assume that demand is divisible. Then, given a solution $\bar{x} := (x_1, x_2) \in V^2$, each vertex v splits its demand $w(v)$ among facilities that are closest to it, which are defined by the set $F(v, \bar{x}) := \arg\min_{i \in N} d(v, x_i)$. Hence, player i will receive demand from vertices in $V_i(\bar{x}) := \{v \in V \mid d(x_i, v) \leqslant d(x_{3-i}, v)\}$, where x_{3-i} denotes the location chosen by the other player. The motivation of splitting demand between equidistant facilities arises from the assumption that consumers decide randomly between facilities that are equally attractive. One may consider alternative splitting proportions; Section EC.2.1 in the electronic companion considers heterogeneous players who get unequal demand fractions whenever demand is located equidistantly between them.

The *utility* of player i is given by the market share, defined as the total demand achieved at x_i:

$$u_i(\bar{x}) := \sum_{v \in V_i(\bar{x})} \frac{w(v)}{|F(v, \bar{x})|}.$$

A solution $\bar{x}$ is a *pure-strategy Nash equilibrium* of the facility location game if $u_i(x_i, x_{3-i}) \geqslant u_i(y, x_{3-i})$ for any $y \in V$ and for any $i \in N$.

Since both utilities must sum up to W, our model describes a constant-sum game that admits optimal strategies that are the solution to min-max problems (see, e.g., Osborne and Rubinstein 1994). When the game admits a pure-strategy Nash equilibrium, we refer to a vertex that represents an optimal strategy as a *winning strategy*. We make the following key observation.

Remark 1. In an equilibrium of a facility location game with two players, both of them experience the same utility. This holds because otherwise the player with the lower utility can select the location of the other player and consequently split the market equally. Hence, both utilities at equilibrium must equal $W/2$. This implies that the value of this constant-sum game is $W/2$, and that a vertex $x \in V$ is a winning strategy if $u_1(x, v) \geqslant W/2$ for all $v \in V$.

The utility obtained by a player when choosing a winning strategy is at least $W/2$ regardless of the selection of the other player. Remark 1 does not immediately generalize to more players. In fact, the characterization of equilibria with more players is much more involved and establishing conditions that guarantee existence is far from trivial. Moreover, deciding if an equilibrium exists is NP-hard for an arbitrary number of players (Dürr and Thang 2007). We indicate some directions of possible generalizations to more players in Section 5.

There is a direct connection between the location of winning strategies and that of equilibria. In fact, an equilibrium exists if and only if a winning strategy exists, and any equilibrium must consist of each player locating a facility in a winning strategy. This is formalized in the next result.

Proposition 1. *For arbitrary topologies, an equilibrium of a facility location game with two players exists if and only if there exists at least one winning strategy.*

Note that an equilibrium is unique if and only if there is a single winning strategy. For a vertex $x \in V$, the total *consumer cost* $C(x) := \sum_{z \in V} w(z)d(x,z)$ measures the weighted distance to x. We say that a vertex $v \in V$ is a 1-*median*, or simply a *median*, if it minimizes $C(x)$ over V.

Alternative Solution Concepts. Winning strategies are closely related to *dominant strategies*, which refer to actions that are optimal for an arbitrary action of the opponent. Although the facility location game does not necessarily admit equilibria in dominant strategies, a winning strategy guarantees that the player will not be worse off than her opponent *regardless what strategy the opponent selects*. Nevertheless, a winning strategy needs not be dominant since a better strategy may exist. To illustrate this, consider the path $(v_1, v_2, v_3, v_4, v_5)$ of five vertices with unit weights. The unique winning strategy is to choose v_3 (therefore, the only equilibrium is (v_3, v_3)); however, a best response to an opponent that chooses v_5 is to choose v_4, so v_3 is not a dominant strategy.

The *Stackelberg* version of the facility location game consists of one of the players (the *leader*) anticipating the actions of the second player (the *follower*) and placing her facility optimally. A profile $\bar{x} = (s, s')$ is a *Stackelberg solution* if s maximizes the leader's utility given that after she plays s, the follower chooses s' to maximize her utility. (When a topology does not admit a pure-strategy Nash equilibrium, a leader's location of a Stackelberg equilibrium is a natural candidate for a player to place a facility since it guarantees the maximum possible utility.) For an overview of Stackelberg solutions on discrete and continuous topologies, see Mirchandani and Francis (1990). When the game admits a winning strategy (and hence a Nash equilibrium), the Stackelberg leader also selects it. Indeed, the Stackelberg leader cannot get more than half of the market because the follower can select the same location. Hence, a vertex that is a winning strategy is an optimal decision for the leader. Since in every equilibrium the leader plays a winning strategy and both players obtain a utility of $W/2$, Nash equilibria are also Stackelberg solutions. The leader gets a utility of $W/2$ only if a winning strategy exists (and therefore, by Proposition 1, only if an equilibrium exists).

Figure 1. A Stackelberg Equilibrium That Is Not a Nash Equilibrium

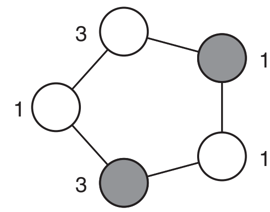

When a winning strategy (or, equivalently, an equilibrium) does not exist, the leader gets strictly less than $W/2$. Nevertheless, a Stackelberg solution need not be an equilibrium. For example, consider the instance in Figure 1. Both vertices with demand equal to 3 are winning strategies, hence in a Nash equilibrium each player must select one of them, possibly the same. Suppose that the Stackelberg leader selects the gray vertex with weight 3. While selecting the other grey vertex (with weight 1) is optimal for the follower, the resulting location profile is not a Nash equilibrium.

We finally note that following Proposition 1, it is immediate that if a winning strategy exists, the solution in which both players select that vertex is a symmetric equilibrium. Therefore, a symmetric equilibrium exists if and only if an equilibrium exists, and the complexity of computing each of these solution concepts is essentially the same.

3. Equilibrium Analysis

Equilibria in Trees and Cycles. Analogously to Hotelling's solution for two facilities on a line, a 1-median of a graph is a natural candidate to locate equilibria. Eiselt and Laporte (1991) showed that in a facility location game with two players on a tree, an equilibrium always exists. Moreover, a solution is an equilibrium if and only if both players select a median of the tree. In our context, that implies that medians of a tree and winning strategies coincide. This result extends the intuition arising from Hotelling's result by showing that the outcome of duopoly competition coincides with the single-facility centralized optimal solution, and is the starting point of the current study.

Moving on to cycle topologies, our game may or may not admit equilibria (Eaton and Lipsey 1976). As an example, a 6-cycle with weights that alternate between 1 and 100 does not admit winning strategies. However, replacing one weight of 100 by 200 makes that vertex be a winning strategy and hence an equilibrium exists. When the graph is a cycle, we observe that if two players locate their facilities on different vertices, each of them would obtain the demand corresponding to half of the vertices (up to shared vertices). Based on this observation, we use connected half-cycles to characterize the equilibria of the game. We define a *half-cycle* to be a set of adjacent vertices that span half of the cycle. More formally, suppose that the cycle has cardinality k. When k is even, a half-cycle may be either a set of $k/2$ adjacent vertices, or a set of $k/2+1$ ones where the two extreme vertices get only half of their original weight (we refer to them as *half-vertices*). When k is odd, each half-cycle contains $(k-1)/2$ adjacent vertices and one half-vertex.

The following result establishes a connection between winning strategies and 1-median solutions in the case of cycles, and it paves the road for the more general topologies studied below.

Proposition 2. *A vertex v is a winning strategy in a cycle if $W(S) \leqslant W/2$ for every half-cycle S that does not contain v. Furthermore, if a winning strategy v of a cycle exists, it must solve the* 1*-median problem.*

Based on this result, Section EC.3.1 in the electronic companion presents an algorithm to find all equilibria for cycles in linear time.

Bi-connected Components and Transformations Based on Them. To study equilibria in more general topologies, we transform general graphs into maximal bi-connected components trees, as defined by Harary (1969). This operation preserves the essence of the topology, insofar as pure-strategy equilibrium locations go. Using this transformation, we show that the location of equilibria can be narrowed down to a subset of vertices, which have a central location in the graph as we next describe.

Given a graph G, a *maximal bi-connected component* (also referred to as *block*) is a maximal subgraph in which each pair of vertices is connected by at least two vertex-disjoint paths; we denote by $\{G_i\}$ the set of blocks of G. A *cutoff vertex* is a vertex whose removal disconnects the graph; we denote by $\{c_j\}$ the set of cutoff vertices in G.

Using these definitions, any weighted graph G can be represented as a *maximal bi-connected components tree* $G^T(V^T, E^T)$ as follows. Each vertex in G^T represents either a cutoff vertex of G or a block of G; formally, we define $V^T := \{G_i\} \cup \{c_j\}$, where with some abuse of notation vertices are labeled by the cutoff vertex or block they represent. The weights of vertices in G^T are determined by the weights of vertices in the original graph G: (a) the weight of a vertex that represents a cutoff vertex $c_j \in G$ is the same as the original weight of that vertex; (b) the weight of a vertex that represents a block $G_i \subseteq G$ is equal to the sum of the weights of the vertices in G_i that *are not* cutoff vertices.

Finally, the set of edges E^T is formed by connecting each vertex $G_i \in G^T$ representing a block to each vertex $c_j \in G^T$ representing a cutoff vertex when $c_j \in G_i$, and by connecting all cutoff vertices in G^T that are adjacent in G but do not belong to the same block. By construction, the resulting graph G^T is a tree. Furthermore, the maximal bi-connected components tree of a given graph is unique because blocks are uniquely defined (Harary 1969).

Figure 2 illustrates the transformation of a unit-weight graph G (depicted on the left side of the figure) into its bi-connected components tree G^T (depicted in the center). Cutoff vertices in G are represented in G^T by vertices of the same weight, so all vertices in G^T that represent cutoff vertices have unit weight; these vertices are not labeled. Blocks in G are represented in G^T by the labeled vertices, where labels indicate the weights. Recall that the weight of a vertex that represents a block is equal to the sum of weights of vertices in that block, *excluding* cutoff vertices. (For instance, block A has six unit-weight vertices and four of these are cutoff vertices and therefore the weight of the vertex representing A in G^T is equal to 2.) Finally, edges in G^T are the subset of edges in G that connect two cutoff vertices.

We use this decomposition to establish a relationship between winning strategies and medians. In particular, we will see that (a) if there is a median of G^T that represents a block, any equilibria of the original graph G can only be located in vertices belonging to that block; and (b) all medians of G^T that represent a cutoff vertex define equilibria since winning strategies can only be located in cutoff vertices or blocks that are medians of G^T. We refer to a block that is represented by a median in G^T as a *central block*. Slater (1975) calls a vertex v a *security center* for a model with unit demands if it maximizes $\min_{w \in V \setminus \{v\}} f(v,w)$, where $f(v,w)$ is the number of vertices that are closer to v than to w minus those that are closer to w than to v. By generalizing

Figure 2. (*Left*) A Unit-Weight Graph G. (*Center*) G^T, the Bi-Connected Components Tree Representation of G. Vertices Without Labels Represent Cutoff Vertices and Have a Unit Demand. (*Right*) Projection of G Onto Blocks A and C

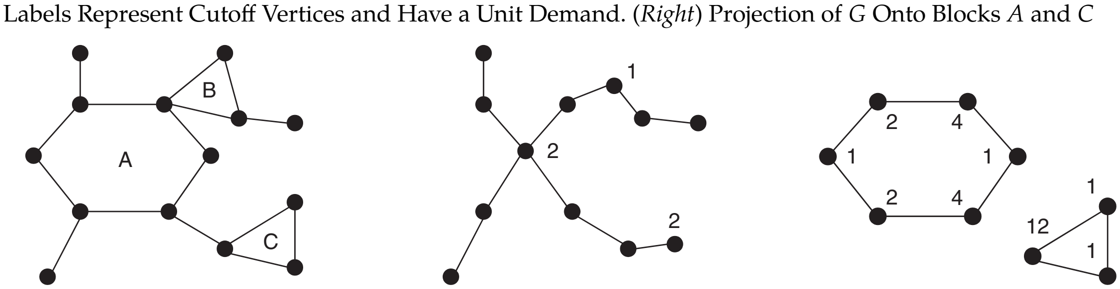

$f(v, w)$ to consider arbitrary demands, it can be seen that a winning strategy must be a security center. However, a security center v can only be a winning strategy if $\min_{w \in V \setminus \{v\}} f(v, w) \geqslant 0$.

To characterize the winning strategies, we define the projection of a graph into a block as follows. Consider a weighted graph G and a block $G_j \subseteq G$. If a vertex v is connected to G_j through vertex $x \in G_j$, x is referred to as the projection of v onto G_j. Using that, we define the weighted graph G'_j as the *projection* of G onto G_j, where both G'_j and G_j have the same vertices and edges, and only differ in weights. The weight of each vertex $v \in G'_j$ is $w(v)$ plus the weights of all other vertices projected onto v. That is, each vertex in the projection G'_j gets its original weight in G plus the total weight of the rest of the graph that connects through it. The right-most panel of Figure 2 shows the projections of the graph G corresponding to blocks A and C.

The following result, which characterizes winning strategies in general graphs, implies that an equilibrium must be located in a central block or in cutoff vertices, and that a vertex is a winning strategy in a graph if and only if it is a winning strategy in the projection of this graph onto its central block.

Theorem 1. *Consider a facility location game with two players on a general graph G. If all medians of G^T are cutoff vertices, the winning strategies are exactly those medians. Otherwise, there is a median that represents a block. Referring to that block as G_j, the winning strategies of G coincide with the winning strategies of G'_j, the projection of G onto G_j*

The uniqueness of equilibria depends on the number of winning strategies in G. If there are multiple medians in G^T, the equilibrium may or may not be unique. Theorem 1 suggests a new approach to find the equilibria, which is summarized by the following conceptual algorithm. We highlight that the running time of the algorithm is polynomial under the unit-length edge assumption. However, if implemented naively, the consideration of general edge lengths involves a transformation of the graph that makes the suggested algorithm pseudo polynomial. A more careful implementation would take into account that there is a polynomial number of possible solutions and deviations that need to be evaluated. This implies that the complexity cannot exceed $O(|V|^3 SP(G))$, where $SP(G)$ is the complexity of computing a shortest path tree in G.

1. List the blocks of the graph (in $\mathcal{O}(|E|)$ time, see Hopcroft and Tarjan 1973, Aho et al. 1974)
2. Construct the tree G^T (in $\mathcal{O}(|E|)$ time, see Section EC.3.2 in the electronic companion)
3. Find the medians of G^T (in $\mathcal{O}(|V|)$ time, see Goldman 1971)
4. If a median represents a block, project G onto the central block (in $\mathcal{O}(|V|)$ time)
5. Find the equilibria in the central block

We define *simple central block* (SCB) graphs as those in which every central block is either a cycle or a vertex. This class of graphs includes topologies of interest such as cactus graphs (Geller and Manvel 1969). The combination of Theorem 1 with the results for cycles and trees provides a concrete characterization of equilibria for SCB graphs. In such a case, vertex v is a winning strategy if and only if v is a cutoff vertex that is a median of G^T or v is a winning strategy of the cycle formed when projecting G onto its unique central block. We note that a general graph can have more complicated central blocks, making it harder to find an equilibrium; yet, our results reduce the problem of finding an equilibrium in a graph to finding one in its central block.

4. Further Connections Between Winning Strategies and 1-Medians

As mentioned in the introduction, Eiselt and Laporte (1991) established a connection between 1-medians and equilibria in the facility location game in trees. We consider the following two natural questions: (i) Is it true that when an equilibrium exists every winning strategy must be a solution to the 1-median problem? (ii) Is it possible to generalize the aforementioned result obtained for trees? In other words, is there any other class of graphs for which an equilibrium always exists and for which selecting a median guarantees an equilibrium? Addressing the first question, we show that any vertex selected at equilibrium must be a *local* 1-*median*. Here, a vertex y is a local 1-median of G if $C(y) \leqslant C(v)$ for all $v \in N(y) := \{z \in V : yz \in E\}$. In words, the cost of y is not worse than that of its neighbors.

Theorem 2. *If y is a winning strategy of the facility location game in a general graph G, then y is a local* 1-*median of G.*

We note that Theorem 2 provides an alternative algorithm to identify equilibrium locations. It consists of finding local 1-medians and testing if they are equilibria by evaluating deviations. Even though it is natural to think that a more general property holds with (global) 1-medians, the example in Figure 3 shows this is not necessarily the case (Scoccola 2014). In that instance, there is a unique median and a unique winning strategy, and both are located in different vertices. This motivates characterizing classes of graphs for which winning strategies and 1-medians coincide. Interestingly, Bandelt and Chepoi (2002) proved that the following conditions are equivalent: (a) The median-set is connected for arbitrary weights w. (b) The set of local medians coincide with the median-set for arbitrary rational weights w. Based on this equivalence, we obtain the following corollary that implies that every winning strategy solves

Figure 3. A Winning Strategy May Not Be Located at the (Global) 1-Median

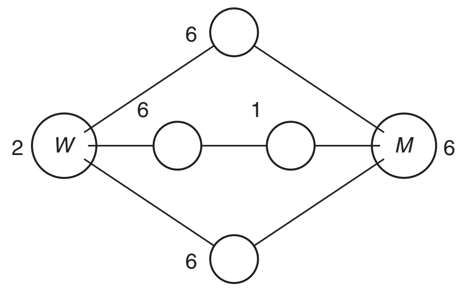

Notes. Numbers represent the weight of each vertex. The unique 1-median is labeled as M, and the unique winning strategy is labeled as W.

the 1-median problem. To describe the result compactly, let us refer to graphs for which equilibria must be located in a 1-median as *must play median* (MPM) graphs. Cycles and trees are examples of MPM graphs.

Corollary 1. *Let G be a graph such that the solutions to the 1-median problem for any rational weights $\{w(v)\}_{v\in V}$ induce a connected subgraph of G. Then, G is an MPM graph.*

Families of graphs satisfying this property include *median graphs*, *quasi-median graphs*, *pseudo-median graphs*, *Helly graphs*, and *strongly chordal graphs*; Bandelt and Chepoi (2002) provide a complete characterization of graphs with connected median sets. Among MPM graphs, we put special emphasis on median graphs because they are particularly relevant to us. A median graph is defined as a graph that satisfies that any three vertices a, b, and c have a unique median (which is a vertex that belongs to shortest paths between any two of a, b, and c). This class includes lattices, meshes, and grids, which encode the characteristics of networks found in the real world when dealing with location problems in urban areas. Next, we prove that arbitrary graphs whose central blocks satisfy the MPM property are MPM graphs.

Theorem 3. *Consider a facility location game with two players on a graph G whose central block is an MPM graph. Then, G is an MPM graph.*

Theorem 3 implies that simple central block graphs are MPM graphs, and that with much more generality than in the previous sections, one may search for equilibrium locations in an arbitrary MPM graph by checking if any of the 1-medians are a winning strategy.

We now address the second question raised at the beginning of this section, namely, whether there exists a class of graphs for which an equilibrium always exists and for which selecting a median is enough to guarantee that players will be at equilibrium. To that end, we focus on strongly chordal graphs, which generalize many well-known classes of graphs such as trees, block graphs, and interval graphs. A graph is *chordal* if every cycle with more than three vertices has a *chord* (an edge joining two nonconsecutive vertices of the cycle). A *p-sun* is a chordal graph with a vertex set $x_1, \dots, x_p, y_1, \dots, y_p$ such that $y_1, \dots, y_p$ is an independent set, $(x_1, \dots, x_p, x_1)$ is a cycle, and each vertex y_i has exactly two neighbors x_{i-1} and x_i, where $x_0 = x_p$. A graph is *strongly chordal* if it is chordal and contains no p-sun for $p \geqslant 3$.

Since the set of 1-medians of these graphs induce a connected component, we know from the previous section that they are MPM graphs, establishing a necessary condition for equilibria. Now we prove the converse to get a sufficient condition as well: a vertex that is a 1-median must be a winning strategy. This guarantees that strongly chordal graphs always admit equilibria and that the equilibrium locations and 1-medians coincide. This completely extends the results for trees of Eiselt and Laporte (1991) to this family. As opposed to the class of cacti, each edge may belong to many cycles. Because there are many chords, these graphs are densely populated with triangles.

Theorem 4. *Every connected strongly chordal graph has an equilibrium. Furthermore, there is a one-to-one correspondence between winning strategies and the solutions to the 1-median problem.*

The proof follows the main ideas of Theorem 1 in Lee and Chang (1994), where it is shown that the median-set of a connected strongly chordal graph is a clique. As a corollary, the set of winning strategies of a connected strongly chordal graph is not only connected (as previously discussed), but also a clique. Considering more general graphs, the ideas presented earlier imply that when the central block of an arbitrary graph is strongly chordal, this graph must have an equilibrium because a median must be a winning strategy.

5. Final Remarks

We have provided an exhaustive characterization of equilibria for duopolistic facility location games under different classes of topologies. Figure 4 provides an illustration of some of the classes of graphs that we have considered. For SCB graphs (including trees, cycles, and cacti) and for graphs whose 1-medians form a connected subgraph, we have established that equilibria must always coincide with 1-medians. For strongly chordal graphs, we have shown that every solution consisting of 1-medians must be an equilibrium. In addition, for SCB graphs, we have provided a procedure to decide if an equilibrium exists or not in $\mathcal{O}(|V|)$ time and characterized all their locations.

Network Design Recommendations. A system manager controlling a network may want to introduce changes to its topology, or even design it from scratch

Figure 4. Nesting of the Families of Graphs Considered in Our Results

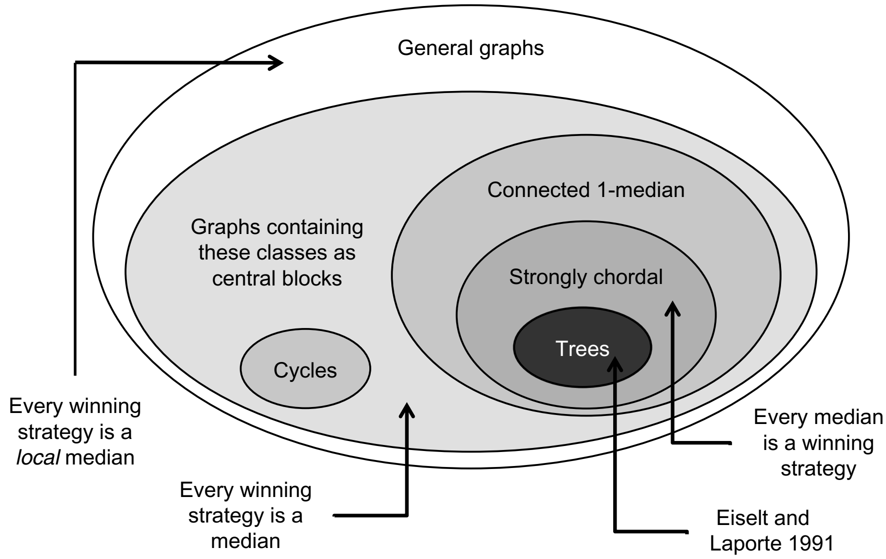

in a way that limits the negative effects of miscoordination and improve the efficiency of the facility location game at equilibrium. Network design challenges have been addressed from various centralized perspectives (see, e.g., Ravi and Sinha 2006 and references therein). In Section EC.1 of the electronic companion, we study the inefficiency of equilibria in facility location games, and show that in terms of consumer cost in equilibrium, star-shaped networks (in which all vertices are connected directly to one central vertices) are the most preferable among SCB networks. We establish this by showing that in SCB graphs the consumer cost at equilibrium cannot decrease when edges are removed. This monotonicity property allows one to measure and bound the inefficiency of equilibria for the class of monotone graphs, indicating that size, diameter, and demand variability are drivers of equilibria inefficiency. For a given size, one can keep the diameter small by having symmetry around the "center" of the instance, and the variability small by spreading the total demand W as uniformly as possible. Given the aforementioned monotonicity property, it suffices to minimize the consumer cost at equilibrium among all trees, and the solution of this problem is given by a star structure.

In some cases, although managers may not be able to control the topology of the network, they can get value from being able to anticipate the players' locations in advance. When a pure strategy Nash equilibrium exists, the locations associated to it provide a sensible forecast. Otherwise, a natural forecast is given by the leader's location at a Stackelberg solution since that location maximizes the utility of the player. Under the assumptions we have made, a Stackelberg solution is guaranteed to exist, and Hakimi (1983, 1986) presented an algorithmic approach to compute it. We have elaborated about this natural interpretation of Stackelberg solutions when discussing equilibrium existence in Section 2.

Robustness Analysis and Practical Impact. We have tested the robustness of our results when demand is not split equally between facilities in case of a tie. Although in theory even a minimal change in the splitting rule may modify the existence of equilibria, in Section EC.2.2 of the electronic companion, we provide evidence that such cases are rare in real-world urban networks. We run numerous experiments using random portions of the city of Buenos Aires, Argentina, using street corners as vertices (with uniform or normal weight distribution) and streets as edges, inducing a total of approximately 17,000 vertices and 30,000 edges. Although real-world urban networks are not necessarily grid graphs, in more than 99.5% of the tested instances that admit an equilibrium, winning strategies and 1-medians coincided, even under the asymmetric splitting rule.

Model Primitives. Allowing zero weights on vertices does not affect the existence and characterization results in this work. For trees, it is possible that there are more than two medians (which may happen only when there is a path of zero-weight vertices between two positive-weight medians), leading to additional equilibria. Considering cycles, the definition of a winning strategy does not assume strictly positive weights, and therefore results hold without changes. Hence, we can extend our results to SCB graphs with zero weights. Furthermore, in a cycle with arbitrary edge lengths, one may subdivide edges with zero-weight vertices at unit distances from each other. These artificial vertices do not constitute feasible locations for facilities, and should not be taken into account when determining if a vertex is a winning strategy. The computational effort in looking for winning strategy vertices does not change, and Proposition 2 holds as well. As results for trees hold for general edge lengths, our results extend to SCB graphs.

Another direction of further research is to characterize equilibria for an arbitrary number of players in trees and in cycles with arbitrary weights. One difficulty that arises when considering more than two players is that an equilibrium may not split the market in equal parts, as Remark 1 indicates it happens for duopolies. In the cases of trees, an equilibrium may not exist, and if it does it may not be located on a median. For the case of three players, one can show that there is an equilibrium only if any median of that tree has a degree strictly larger than two, which generalizes a result of Lerner and Singer (1937) whereby an equilibrium does not exist for three players on a unit-weight path graph. More generally, Eiselt and Laporte (1993) provided sufficient conditions and a characterization of equilibria on a tree with three players, under various assumptions. While a generalization to an arbitrary number of players on a tree is difficult, such generalization may be successful for simpler structures. In the case of cycles, our analysis is based on how players split the vertices and the total weight of the cycle; this property does not hold for more than two players. Nevertheless, the intuition gained in the treatment for two players may help when considering necessary and sufficient conditions for equilibria for more than two players.

Finally, one may consider an extension of our model where each player can control multiple facilities. We note that our results cannot be trivially extended in that direction. Even when considering tree (or even path) networks with players controlling two facilities each, an equilibrium in pure strategies might not exist, and even when it does, the solution might not be one where both players select a 2-median of the tree. Through a computational study, we analyzed various tree instances with two players placing two facilities each. We found that, when demands are randomly selected, an equilibrium is likely to exist, especially in small trees. Furthermore, in about 85%–90% of the cases in which the equilibrium did exist, players selected the 2-median. Further details can be found in Section EC.2.3 of the electronic companion. Overall, this simulation provides some partial evidence that even though the theoretical results might not generalize exactly, our high-level insights apply beyond the assumptions of the current paper.

Acknowledgments

The authors would like to thank Ehud Lehrer, Miquel Oliu Barton, Iván Rapaport, Luis Scoccola, and the editorial team who reviewed this paper for insightful discussions and enriching comments.

References

Ahn H-K, Cheng S-W, Cheong O, Golin M, van Oostrum R (2004) Competitive facility location: The Voronoi game. *Theoret. Comput. Sci.* 310(1–3):457–467.

Aho AV, Hopcroft JE, Ullman JD (1974) *The Design and Analysis of Computer Algorithms* (Addison-Wesley, Reading, MA).

Bandelt H, Chepoi V (2002) Graphs with connected medians. *SIAM J. Discrete Math.* 15(2):268–282.

Bharathi S, Kempe D, Salek M (2007) Competitive influence maximization in social networks. Deng X, Graham FC, eds. *Proc. Third Internat. Workshop Internet and Network Econom. WINE '07* (Springer, Berlin), 306–311.

de Palma A, Ginsberg V, Papageorgiou YY, Thisse JF (1985) The principle of minimum differential holds under sufficient heterogeneity. *Econometrica* 53(4):767–781.

Drezner T (1995a) Competitive facility location in the plane. Drezner Z, ed. *Facility Location: A Survey of Applications and Methods* (Springer, New York), 285–300

Drezner Z (1995b) *Facility Location: A Survey of Applications and Methods* (Springer, New York).

Drezner Z, Wesolowsky GO, Drezner T (1998) On the logit approach to competitive facility location. *J. Regional Sci.* 38(2):313–327.

Dürr C, Thang NK (2007) Nash equilibria in Voronoi games on graphs. Arge L, Hoffmann M, Welzl E, eds. *Proc. 15th Eur. Sympos. Algorithms ESA '07* (Springer, Berlin, 17–28.

Eaton B, Lipsey R (1976) The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *Rev. Econom. Stud.* 42(1):7–49.

Eiselt HA, Laporte G (1991) Locational equilibrium of two players on a tree. *Recherche Opérationnelle* 25(1):5–18.

Eiselt HA, Laporte G (1993) The existence of equilibria in the 3-facility Hotelling model in a tree. *Transportation Sci.* 27(1): 39–43.

Eiselt HA, Marianov V (2011) *Foundations of Location Analysis*. International Series in Operations Research and Management Science (Springer, New York).

Gabszewicz JJ, Thisse JF (1986) Spatial competition and location of firms. Lesourne J, Sonnenschein H, eds. *Location Theory*. Fundamentals of Pure Appl. Econom., 5 (Harwood Academic Publishers, London), 1–71.

Geller D, Manvel B (1969) Reconstruction of cacti. *Canadian J. Math.* 21(6):1354–1360.

Goemans MX, Mirrokni VS, Vetta A (2005) Sink equilibria and convergence. *Proc. 46th Annual IEEE Sympos. Foundations of Comput. Sci., FOCS '05* (IEEE Computer Society, Washington, DC), 142–154.

Goldman AJ (1971) Optimal center locations in simple networks. *Transportation Sci.* 5(2):212–221.

Hakimi SL (1983) On locating new facilities in a competitive environment. *Eur. J. Oper. Res.* 12(1):29–35.

Hakimi SL (1986) P-Median theorems for competitive locations. *Ann. Oper. Res.* 6(4):79–98.

Hansen P, Thisse J-F, Wendell R (1990) Equilibrium analysis for voting and competitive location problems. Mirchandani PB, Francis RL, eds. *Discrete Location Theory* (John Wiley & Sons, New York), 479–501.

Harary F (1969) *Graph Theory* (Addison-Wesley, Reading, MA).

Hay DA (1976) Sequential entry and entry-deterring strategies in spatial competition. *Oxford Econom. Papers* 28(2):240–257.

Hopcroft J, Tarjan R (1973) Efficient algorithms for graph manipulation. *Commun. ACM* 16(6):372–378.

Hotelling H (1929) Stability in competition. *Econom. J.* 39(153):41–57.

Jain K, Mahdian M (2007) Cost sharing. Nisan N, Roughgarden T, Tardos É, Vazirani VV, eds. *Algorithmic Game Theory* (Cambridge University Press, New York), 385–410.

Kostka J, Oswald YA, Wattenhofer R (2008) Word of mouth: Rumor dissemination in social networks. Shvartsman AA, Felber P, eds. *Proc. 15th Internat. Colloquium on Structural Inform. Commun. Complexity, SIROCCO '08*. Lecture Notes in Computer Science, Vol. 5058 (Springer, Berlin), 185–196.

Lee H, Chang GJ (1994) The w-median of a connected strongly chordal graph. *J. Graph Theory* 18(7):673–680.

Leonardi G, Tadei R (1984) Random utility demand models and service location. *Regional Sci. Urban Econom.* 14(3):399–431.

Lerner AP, Singer HW (1937) Some notes on duopoly and spatial competition. *J. Political Econom.* 45(2):145–186.

Mavronicolas M, Monien B, Papadopoulou VG, Schoppmann F (2008) Voronoi games on cycle graphs. Ochmanski E, Tyszkiewicz J, eds. *Proc. 33rd Internat. Sympos. Math. Foundations Comput. Sci., MFCS '08*. Lecture Notes in Computer Science, Vol. 5162 (Springer, Berlin), 503–514.

Mirchandani PB, Francis RL (1990) *Discrete Location Theory* (John Wiley & Sons, New York).

Mirrokni VS, Vetta A (2004) Convergence issues in competitive games. *7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems APPROX*. Lecture Notes in Computer Science, Vol. 3122 (Springer, Berlin), 183–194.

Osborne MJ, Rubinstein A (1994) *A Course in Game Theory* (MIT Press, Cambridge, MA).

Prescott EC, Visscher M (1977) Sequential location among firms with foresight. *Bell J. Econom.* 8(2):378–393.

Ravi R, Sinha A (2006) Approximation algorithms for problems combining facility location and network design. *Oper. Res.* 54(1): 73–81.

ReVelle CS (1986) The maximum capture or sphere of influence location problem: Hotelling revisited on a network. *J. Regional Sci.* 26(2):343–357.

Salop SC (1979) Monopolistic competition with outside goods. *Bell J. Econom.* 10(1):141–156.

Scoccola L (2014) A counterexample for the competitive facility location problem provided via personal communication with the authors, March 12.

Slater PJ (1975) Maximin facility location. *J. Res. National Bureau of Standards B* 79(3–4):107–115.

Taylor M (1971) Review article: Mathematical political theory. *British J. Political Sci.* 1(3):339–381.

Tirole J (1988) *The Theory of Industrial Organization* (MIT Press, Cambridge, MA).

Vetta A (2002) Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. *Proc. 43rd Annual IEEE Sympos. Foundations Comput. Sci., FOCS '02* (IEEE Computer Society Press, Washington, DC), 416–428.

Wendell RE, McKelvey RD (1981) New prospectives in competition location theory. *Eur. J. Oper. Res.* 6(2):174–182.

Yonatan Gur is an assistant professor of operations, information, and technology at the Graduate School of Business, Stanford University. His research interests include operations management, revenue management, and service systems, with a particular emphasis on data-driven, dynamic, online environments.

Daniela Saban is an assistant professor of operations, information, and technology at the Graduate School of Business, Stanford University. Her research interests include operations and revenue management, with a particular emphasis on the design of dynamic, online markets.

Nicolas E. Stier-Moses is an associate professor with tenure at Di Tella Business School and an independent researcher at CONICET Argentina. His research interests include combinatorial and network optimization, game theory, and operations management with applications in transportation, telecommunication, and high-tech industries.