

Shortest Beer Path Queries in Outerplanar Graphs*

Joyce Bacic[†]

Saeed Mehrabi[‡]

Michiel Smid[†]

August 11, 2022

Abstract

A *beer graph* is an undirected graph G , in which each edge has a positive weight and some vertices have a beer store. A *beer path* between two vertices u and v in G is any path in G between u and v that visits at least one beer store.

We show that any outerplanar beer graph G with n vertices can be preprocessed in $O(n)$ time into a data structure of size $O(n)$, such that for any two query vertices u and v , (i) the weight of the shortest beer path between u and v can be reported in $O(\alpha(n))$ time (where $\alpha(n)$ is the inverse Ackermann function), and (ii) the shortest beer path between u and v can be reported in $O(L)$ time, where L is the number of vertices on this path. Note that the running time for (ii) does not depend on the number of vertices of G . Both results are optimal, even when G is a beer tree (i.e., a beer graph whose underlying graph is a tree).

1 Introduction

Imagine that you are going to visit a friend and, not wanting to show up empty handed, you decide to pick up some beer along the way. In this paper we determine the fastest way to go from your place to your friend's place while stopping at a beer store to buy some drinks.

A *beer graph* is a undirected graph $G = (V, E)$, in which each edge (u, v) has a positive weight $\omega(u, v)$ and some of the vertices are beer stores. For two vertices u and v of G , we define the *shortest beer path* from u to v to be the shortest (potentially non-simple) path that starts at u , ends at v , and visits at least one beer store. We denote this shortest path by $\text{SP}_B(u, v)$. The *beer distance* $\text{dist}_B(u, v)$ between u and v is the weight of the path $\text{SP}_B(u, v)$, i.e., the sum of the edge weights on $\text{SP}_B(u, v)$.

Observe that even though the shortest beer path from u to v may be a non-simple path, it is always composed of two simple paths: the shortest path from u to a beer store and

*A preliminary version was presented at the 32nd Annual International Symposium on Algorithms and Computation (ISAAC 2021). JB was supported by an NSERC Undergraduate Student Research Award. MS was supported by NSERC.

[†]School of Computer Science, Carleton University, Ottawa, Canada.

[‡]UMass Lowell, USA

the shortest path from this same beer store to v . Thus, when looking at the shortest beer path problem, we often need to consider the shortest path between vertices. We denote the shortest path in G from u to v by $\text{SP}(u, v)$ and we use $\text{dist}(u, v)$ to denote the weight of this path. We also say that $\text{dist}(u, v)$ is the *distance* between u and v in G .

To the best of our knowledge, the problem of computing shortest beer paths has not been considered before. Let s be a fixed source vertex of G . Recall that Dijkstra’s algorithm computes $\text{dist}(s, v)$ for all vertices v , by maintaining a “tentative distance” $\delta(v)$, which is the weight of the shortest path from s to v computed so far. If we also maintain a “tentative beer distance” $\delta_B(v)$ (which is the weight of the shortest beer path from s to v that has been found so far), then a modification of Dijkstra’s algorithm allows us to compute $\text{dist}_B(s, v)$ for all vertices v , in $O(|V| \log |V| + |E|)$ total time.

As far as we know, no non-trivial results are known for beer distance queries. In this case, we want to preprocess the beer graph G into a data structure, such that, for any two query vertices u and v , the shortest beer path $\text{SP}_B(u, v)$, or its weight $\text{dist}_B(u, v)$, can be reported.

Observe that, instead of beer stores, we could have used any type of facility, and consider shortest paths that visit at least one facility. As such, queries of the type considered in this paper are of relevance for practical applications. For example, when going home from work, we may have to pick up groceries in the supermarket, in which case we are interested in “supermarket paths”. In case we want to drop off a parcel at the post office, we are interested in “post office paths”. As a final example, if we have to charge our electric car, then we want to determine the “electric vehicle charging station distance”.

1.1 Our Results

We present data structures that can answer shortest beer path queries in outerplanar beer graphs. Recall that a graph G is *outerplanar*, if G can be embedded in the plane, such that all vertices are on the outer face, and no two edges cross.

Our first result is stated in terms of the inverse Ackermann function. We use the definition as given in [3]: Let $A_0(i) = i + 1$ and, for $\ell \geq 0$, $A_{\ell+1}(i) = A_\ell^{(i+1)}(i + 8)$, where $A_\ell^{(i+1)}$ is the function A_ℓ iterated $i + 1$ times. We define $\alpha(m, n)$ to be the smallest value of ℓ for which $A_\ell(\lfloor m/n \rfloor) > n$, and we define $\alpha(n) = \alpha(n, n)$.

For example, straightforward induction proofs show that $A_1(i) > 2i$, $A_2(i) > 2^i$, and

$$A_3(i) > 2^{2^{\cdot^{2^i}}},$$

where the number of 2’s is equal to i . Thus, if we let $\log^* n$ denote the number of times the function “log” must be applied, when starting with the value n , until the result is at most 1, then $\alpha(n \log^* n, n) \leq 3$. Let $\log^{**} n$ be the number of times the function “log^{*}” must be applied, again starting with n , until the result is at most 1. Since the function $\alpha(m, n)$ is non-decreasing in m , and since $\log^{**} n \leq \log^* n$, it also follows that $\alpha(n \log^{**} n, n) \leq 3$.

Theorem 1 *Let G be an outerplanar beer graph with n vertices. For any integer $m \geq n$, we can preprocess G in $O(m)$ time into a data structure of size $O(m)$, such that for any two query vertices u and v , both $\text{dist}(u, v)$ and $\text{dist}_B(u, v)$ can be computed in $O(\alpha(m, n))$ time.*

By taking $m = n$, both the preprocessing time and the space used are $O(n)$, and for any two query vertices u and v , both $\text{dist}(u, v)$ and $\text{dist}_B(u, v)$ can be computed in $O(\alpha(n))$ time. As another example, by taking $m = n \log^{**} n$, we obtain a data structure with space and preprocessing time $O(n \log^{**} n)$ that can answer both distance and beer distance queries in $O(1)$ time.

As we mentioned before, beer distance queries have not been considered for any class of graphs. In fact, the only result on (non-beer) distance queries in outerplanar graphs that we are aware of is by Djidjev *et al.* [5]. They show that an outerplanar graph with n vertices can be preprocessed in $O(n \log n)$ time into a data structure of size $O(n \log n)$, such that any distance query can be answered in $O(\log n)$ time. Our result in Theorem 1 significantly improves their result.

We also show that the result in Theorem 1 is optimal for beer distance queries, even if G is a *beer tree* (i.e., a beer graph whose underlying graph is a tree). We do not know if the query time is optimal for (non-beer) distance queries.

Our second result is on reporting the shortest beer path between two query vertices.

Theorem 2 *Let G be an outerplanar beer graph with n vertices. We can preprocess G in $O(n)$ time into a data structure of size $O(n)$, such that for any two vertices u and v , the shortest beer path from u to v can be reported in $O(L)$ time, where L is the number of vertices on this beer path.*

Observe that the query time in Theorem 2 does not depend on the number n of vertices of the graph. Again, we are not aware of any previous work on reporting shortest beer paths. Djidjev *et al.* [5] show that, after $O(n \log n)$ preprocessing and using $O(n \log n)$ space, the shortest (non-beer) path between two query vertices can be reported in $O(\log n + L)$ time, where L is the number of vertices on the path.

1.2 Preliminaries and Organization

Throughout this paper, we only consider outerplanar beer graphs G . The number of vertices of G is denoted by n . It is well known that G has at most $2n - 3$ edges. As in [5], we say that G satisfies the *generalized triangle inequality*, if for every edge (u, v) in G , $\text{dist}(u, v) = \omega(u, v)$, i.e., the shortest path between u and v is the edge (u, v) .

The outerplanar graph G is called *maximal*, if adding an edge between any two non-adjacent vertices of G results in a graph that is not outerplanar. In this case, the number of edges is equal to $2n - 3$. A maximal outerplanar graph G is 2-connected, each internal face of G is a triangle and the outer face of G forms a Hamiltonian cycle. In such a graph, edges on the outer face will be referred to as *external* edges, where all other edges will be referred to as *internal* edges.

The *weak dual* of a maximal outerplanar graph G is the graph $D(G)$ whose node set is the set of all internal faces of G , and in which (F, F') is an edge if and only if the faces F and F' share an edge in G ; see Figure 1. For simplicity, we will refer to $D(G)$ as the dual of G . Observe that $D(G)$ is a tree with $n - 2$ nodes, each of which has degree at most three.

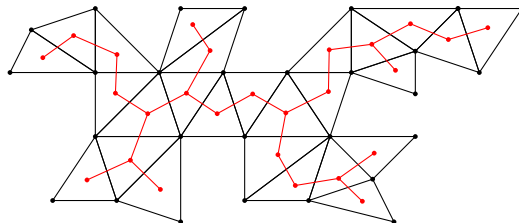


Figure 1: A maximal outerplanar graph shown in black. Its dual is shown in red.

If H is a subgraph of the beer graph G , and u and v are vertices of H , then $\text{dist}(u, v, H)$ and $\text{dist}_B(u, v, H)$ denote the distance and beer distance between u and v in H , respectively. The shortest beer path in H between u and v must be entirely within H . Observe that we use the shorthand $\text{dist}(u, v)$ for $\text{dist}(u, v, G)$, and $\text{dist}_B(u, v)$ for $\text{dist}_B(u, v, G)$.

It will not be surprising that the algorithms for computing shortest beer paths use the dual $D(G)$. Thus, our algorithms will need some basic data structures on trees. These data structures will be presented in Section 2.

In Section 3, we will prove Theorem 1 for maximal outerplanar beer graphs. We also prove that the result in Theorem 1 is optimal, even for beer trees. The proof of Theorem 2, again for maximal outerplanar beer graphs, will be presented in Section 4. Both Sections 3 and 4 will use the result in Lemma 4, whose detailed proof will be given in Section 5.

The extensions of Theorems 1 and 2 to arbitrary outerplanar beer graphs will be given in Section 6. Finally, Section 7 will present an $O(n)$ -time algorithm for computing the single-source shortest beer path tree for any given source vertex.

2 Query Problems on Trees

Our algorithms for computing beer shortest paths in an outerplanar graph G will use the dual of G , which is a tree. In order to obtain fast implementations of these algorithms, we need to be able to solve several query problems on this tree. In this section, we present all query problems that will be used in later sections.

Lemma 1 *Let T be a tree with n nodes that is rooted at an arbitrary node. We can preprocess T in $O(n)$ time, such that each of the following queries can be answered in $O(1)$ time:*

1. *Given a node u of T , return its level, denoted by $\text{level}(u)$, which is the number of edges on the path from u to the root.*
2. *Given two nodes u and v of T , report their lowest common ancestor, denoted by $\text{LCA}(u, v)$.*

3. Given two nodes u and v of T , decide whether or not u is in the subtree rooted at v .
4. Given two distinct nodes u and v of T , report the second node on the path from u to v .
5. Given three nodes u , v , and w , decide whether or not w is on the path between u and v .

Proof. The first claim follows from the fact that by performing an $O(n)$ -time pre-order traversal of T , we can compute $level(u)$ for each node u . A proof of the second claim can be found in Harel and Tarjan [6] and Bender and Farach-Colton [2]. The third claim follows from the fact that u is in the subtree rooted at v if and only if $LCA(u, v) = v$. A proof of the fourth claim can be found in Chazelle [4, Lemma 15]. The fifth claim follows from the following observations. Assume that u is in the subtree rooted at v . Then w is on the path between u and v if and only if $LCA(u, w) = w$ and w is in the subtree rooted at v . The case when v is in the subtree rooted at u is symmetric. Assume that $LCA(u, v) \notin \{u, v\}$. Then w is on the path between u and v if and only if w is on the path between u and $LCA(u, v)$ or w is on the path between v and $LCA(u, v)$. ■

2.1 Closest-Color Queries in Trees

Let T be a tree with n nodes and let \mathcal{C} be a set of colors. For each color c in \mathcal{C} , we are given a path P_c in T . Even though these paths may share nodes, each node of T belongs to at most a constant number of paths. This implies that the total size of all paths P_c is $O(n)$. We assume that each node u of T stores the set of all colors c such that u is on the path P_c .

In a *closest-color query*, we are given two nodes u and v of T , and a color c , such that u is on the path P_c . The answer to the query is the node on P_c that is closest to v . In other words, when following the path in T from u to v , the answer to the query is the last node on this path that is also on P_c . Refer to Figure 2 for an illustration.

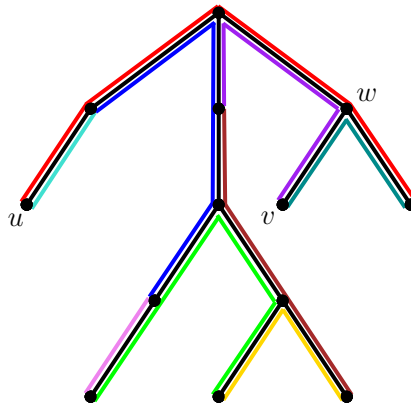


Figure 2: A tree T and a collection of colored paths. For a query with nodes u and v , and color “red”, the answer is the node w .

Lemma 2 *After an $O(n)$ -time preprocessing, we can answer any closest-color query in $O(1)$ time.*

Proof. We take an arbitrary node of T and make it the root. Then we preprocess T such that each of the queries in Lemma 1 can be answered in $O(1)$ time.

For each color c , let c^1 and c^2 be the end nodes of the path P_c , and let c^h be the highest node on P_c in the tree (i.e., the node on P_c that is closest to the root). With each node of P_c , we store pointers to c^1 , c^2 , and c^h .

Since each node of T is in a constant number of colored paths, we can compute the pointers for all the colored paths in $O(n)$ total time.

The query algorithm does the following. Let u and v be two nodes of T , and let c be a color such that u is on the c -colored path P_c .

If $u = v$ or v is also on P_c , then we return the node v . From now on, assume that $u \neq v$ and v is not on P_c . Below, we consider all possible cases, which are illustrated in Figure 3.

1. If $\text{LCA}(u, v) = v$, then u is in the subtree rooted at v . In this case, we return c^h , the highest c -colored node.
2. Assume that $\text{LCA}(u, v) = u$. Then v is in the subtree rooted at u .
 - (a) If $\text{level}(\text{LCA}(v, c^1)) \geq \text{level}(\text{LCA}(v, c^2))$, then we return $\text{LCA}(v, c^1)$.
 - (b) If $\text{level}(\text{LCA}(v, c^1)) < \text{level}(\text{LCA}(v, c^2))$, then we return $\text{LCA}(v, c^2)$.

We only justify the correctness for case (a), because the argument for case (b) is symmetric. Let u' be the child of u such that v is in the subtree rooted at u' .

First assume that the path P_c does not contain any node of the subtree rooted at u' . Then $\text{LCA}(v, c^1) = u$, which is the correct answer to the query.

Next assume that c^1 is in the subtree rooted at u' , but c^2 is not in this subtree. Then $\text{LCA}(v, c^1)$ is in this subtree and is the correct answer to the query.

If c^1 is not in the subtree rooted at u' , but c^2 is in this subtree, then $\text{LCA}(v, c^2)$ is also in this subtree, whereas $\text{LCA}(v, c^1)$ is on the path from u to the root of T . Thus, $\text{level}(\text{LCA}(v, c^1)) < \text{level}(\text{LCA}(v, c^2))$ and, therefore, this cannot happen.

We finally observe that c^1 and c^2 cannot both be in the subtree rooted at u' , because P_c is a path that contains u , c^1 , and c^2 .

3. Assume that $\text{LCA}(u, v) \neq u$ and $\text{LCA}(u, v) \neq v$. Then u and v are in different subtrees of $\text{LCA}(u, v)$.
 - (a) If $\text{level}(c^h) > \text{level}(\text{LCA}(u, v))$, then we return c^h .

For the correctness, let u' be the child of $\text{LCA}(u, v)$ whose subtree contains u . Then c^h is on the path in T from u to u' , and c^h is the correct answer to the query.

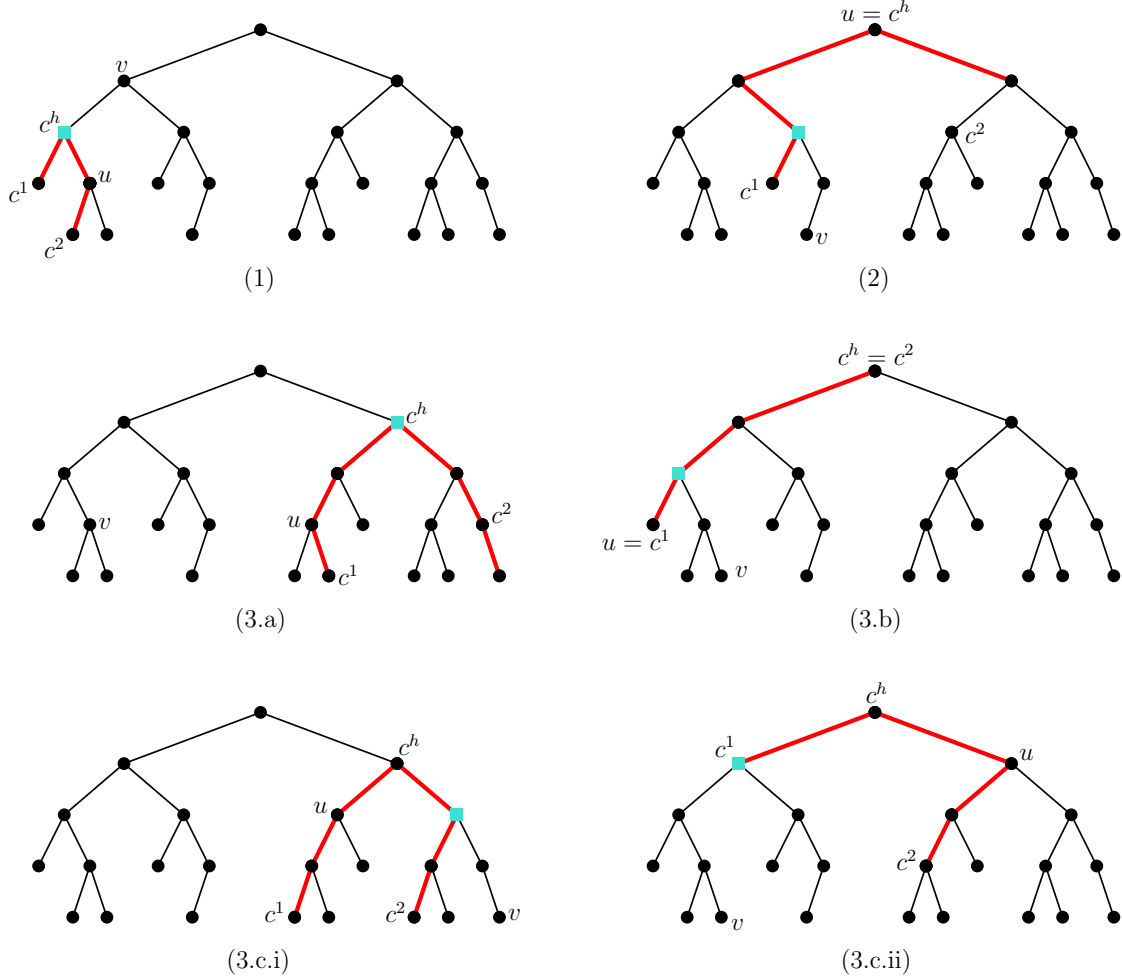


Figure 3: Illustrating all possible cases in the proof of Lemma 2. The path P_c is red and the blue square indicates the node that is returned by the closest-color query.

(b) If $\text{level}(c^h) < \text{level}(\text{LCA}(u, v))$, then we return $\text{LCA}(u, v)$.

For the correctness, let x be the child of c^h whose subtree contains u . Then $\text{LCA}(u, v)$ is also in this subtree and $\text{LCA}(u, v)$ is on P_c . Since $\text{LCA}(u, v)$ is the only node on the path from v to $\text{LCA}(u, v)$ that is on P_c , the node $\text{LCA}(u, v)$ is the correct answer to the query.

(c) Assume that $\text{level}(c^h) = \text{level}(\text{LCA}(u, v))$. Observe that exactly one end node of P_c is in the subtree rooted at u .

i. If c^1 is in the subtree rooted at u , then we return $\text{LCA}(v, c^2)$.

ii. If c^2 is in the subtree rooted at u , then we return $\text{LCA}(v, c^1)$.

For the correctness, assume that c^1 is in the subtree rooted at u . Then $\text{LCA}(u, v) = c^h$. If c^2 is on the path from v to $\text{LCA}(u, v)$, then $\text{LCA}(v, c^2) = c^2$, which is the correct answer to the query. Otherwise, $\text{LCA}(u, v)$ has at least three children,

and u , v , and c^2 are in different subtrees of $\text{LCA}(u, v)$. In this case, $\text{LCA}(v, c^2) = \text{LCA}(u, v)$, which is the correct answer to the query.

Using Lemma 1, each of these case takes $O(1)$ time. Therefore, the entire query algorithm takes $O(1)$ time. \blacksquare

2.2 Path-Sum Queries in Trees

Let (W, \oplus) be a semigroup. Thus, W is a set and $\oplus : W \times W \rightarrow W$ is an associative binary operator. We assume that for any two elements s and s' in W , the value of $s \oplus s'$ can be computed in $O(1)$ time.

Let T be a tree with n nodes in which each edge e stores a value $s(e)$, which is an element of W . For any two distinct nodes u and v in T , we define their *path-sum* $\text{PS}(u, v)$ as follows: Let e_1, e_2, \dots, e_k be the edges on the path in T between u and v . Then we define $\text{PS}(u, v) = \oplus_{i=1}^k s(e_i)$.

Chazelle [4] considers the problem of preprocessing the tree T , such that for any two distinct query nodes u and v , the value of $\text{PS}(u, v)$ can be reported. (See also Alon and Schieber [1], Thorup [9], and Chan *et al.* [3].) Chazelle's result is stated in terms of the inverse Ackermann function; see Section 1.1.

Lemma 3 *Let T be a tree with n nodes in which each edge stores an element of the semigroup (W, \oplus) . For any integer $m \geq n$, we can preprocess T in $O(m)$ time into a data structure of size $O(m)$, such that any path-sum query can be answered in $O(\alpha(m, n))$ time.*

Remark 1 Assume that (W, \oplus) is the semigroup, where W is the set of all real numbers and the operator \oplus takes the minimum of its arguments. In this case, we will refer to a query as a *path-minimum query*. For this semigroup, the result of Lemma 3 is optimal: Any data structure that can be constructed in $O(m)$ time has worst-case query time $\Omega(\alpha(m, n))$. To prove this, assume that we can answer any query in $o(\alpha(m, n))$ time. Then the on-line minimum spanning tree verification problem on a tree with n vertices and $m \geq n$ queries can be solved in $o(m \cdot \alpha(m, n))$ time, by performing a path-maximum query for the endpoints of each edge e and checking that the weight of e is larger than the path-maximum. This contradicts the lower bound for this problem proved by Pettie [8].

3 Beer Distance Queries in Maximal Outerplanar Graphs

Let G be a maximal outerplanar beer graph with n vertices that satisfies the generalized triangle inequality. We will show how to preprocess G , such that for any two vertices u and v , the weight, $\text{dist}_B(u, v)$, of a shortest beer path between u and v can be reported. Our approach will be to define a special semigroup (W, \oplus) , such that each element of W “contains” certain distances and beer distances. With each edge of the dual $D(G)$, we will store one element of the set W . As we will see later, a beer distance query can then be

reduced to a path-sum query in $D(G)$. Thus, by applying the results of Section 2.2, we will obtain a proof of Theorem 1.

We will need the first claim in the following lemma. The second claim will be used in Section 4.

Lemma 4 *Consider the beer graph G as above.*

1. *In $O(n)$ total time, we can compute $\text{dist}_B(u, u)$ for each vertex u of G , and $\text{dist}_B(u, v)$ for each edge (u, v) in G .*
2. *After an $O(n)$ -time preprocessing of G , we can report,*
 - (a) *for any query edge (u, v) of G , the shortest beer path between u and v in $O(L)$ time, where L is the number of vertices on this path,*
 - (b) *for any query vertex u of G , the shortest beer path from u to itself in $O(L)$ time, where L is the number of vertices on this path.*

Proof. We choose an arbitrary face R of G and make it the root of $D(G)$. Let (u, v) be any edge of G . This edge divides G into two outerplanar subgraphs, both of which contain (u, v) as an edge. Let G_{uv}^R be the subgraph that contains the face R , and let G_{uv}^{-R} denote the other subgraph. Note that if (u, v) is an external edge, then $G_{uv}^R = G$ and G_{uv}^{-R} consists of the single edge (u, v) . By the generalized triangle inequality, the shortest beer path between u and v is completely in G_{uv}^R or completely in G_{uv}^{-R} . The same is true for the shortest beer path from u to itself. Thus, for each edge (u, v) of G ,

$$\text{dist}_B(u, v) = \min(\text{dist}_B(u, v, G_{uv}^R), \text{dist}_B(u, v, G_{uv}^{-R})),$$

$$\text{dist}_B(u, u) = \min(\text{dist}_B(u, u, G_{uv}^R), \text{dist}_B(u, u, G_{uv}^{-R})).$$

By performing a post-order traversal of $D(G)$, we can compute $\text{dist}_B(u, v, G_{uv}^{-R})$ and $\text{dist}_B(u, u, G_{uv}^{-R})$ for all edges (u, v) , in $O(n)$ total time. After these values have been computed, we perform a pre-order traversal of $D(G)$ and obtain $\text{dist}_B(u, v, G_{uv}^R)$ and $\text{dist}_B(u, u, G_{uv}^R)$, again for all edges (u, v) , in $O(n)$ total time. The details will be given in Section 5. ■

In the rest of this section, we assume that all beer distances in the first claim of Lemma 4 have been computed.

For any two distinct internal faces F and F' of G , let $Q_{F, F'}$ be the union of the two sets

$$\{(u, v, \text{dist}(u, v), \text{D}) \mid u \text{ is a vertex of } F, v \text{ is a vertex of } F'\}$$

and

$$\{(u, v, \text{dist}_B(u, v), \text{BD}) \mid u \text{ is a vertex of } F, v \text{ is a vertex of } F'\},$$

where the “bits” D and BD indicate whether the tuple represents a distance or a beer distance. In words, $Q_{F, F'}$ is the set of all shortest path distances and all shortest beer distances between a vertex in F and a vertex in F' . Since each internal face has three vertices, the set $Q_{F, F'}$ has exactly 18 elements.

Observation 1 *Let u and v be vertices of G , and let F and F' be internal faces that contain u and v as vertices, respectively.*

1. *If $F = F'$, then we can determine both $\text{dist}(u, v)$ and $\text{dist}_B(u, v)$ in $O(1)$ time.*
2. *If $F \neq F'$ and we are given the set $Q_{F, F'}$, then we can determine both $\text{dist}(u, v)$ and $\text{dist}_B(u, v)$ in $O(1)$ time.*

Proof. First assume that $F = F'$. If $u = v$, then $\text{dist}(u, v) = 0$ and $\text{dist}_B(u, v)$ has been precomputed. If $u \neq v$, then (u, v) is an edge of G and, thus, $\text{dist}(u, v) = \omega(u, v)$ and $\text{dist}_B(u, v)$ has been precomputed.

Assume that $F \neq F'$. If we know the set $Q_{F, F'}$, then we can find $\text{dist}(u, v)$ and $\text{dist}_B(u, v)$ in $O(1)$ time, because these two distances are in $Q_{F, F'}$. ■

In the rest of this section, we will show that Lemma 3 can be used to compute the set $Q_{F, F'}$ for any two distinct internal faces F and F' .

Lemma 5 *For any edge (F, F') of $D(G)$, the set $Q_{F, F'}$ can be computed in $O(1)$ time.*

Proof. Let u be a vertex of F and let v be a vertex of F' . Consider the subgraph $G[F, F']$ of G that is induced by the four vertices of F and F' ; this subgraph has five edges. By the generalized triangle inequality, $\text{dist}(u, v) = \text{dist}(u, v, G[F, F'])$. Thus, $\text{dist}(u, v)$ can be computed in $O(1)$ time.

We now show how $\text{dist}_B(u, v)$ can be computed in $O(1)$ time. If $u = v$ or (u, v) is an edge of G , then $\text{dist}_B(u, v)$ has been precomputed. Assume that $u \neq v$ and (u, v) is not an edge of G . Let w and w' be the two vertices that are shared by F and F' . Since any path in G between u and v (and, in particular, any shortest beer path between u and v) contains at least one of w and w' , $\text{dist}_B(u, v)$ is the minimum of

1. $\text{dist}_B(u, w) + \omega(w, v)$,
2. $\omega(u, w) + \text{dist}_B(w, v)$,
3. $\text{dist}_B(u, w') + \omega(w', v)$,
4. $\omega(u, w') + \text{dist}_B(w', v)$.

Since (u, w) , (w, v) , (u, w') , and (w', v) are edges of G , all terms in these four sums have been precomputed. Therefore, $\text{dist}_B(u, v)$ can be computed in $O(1)$ time.

We have shown that each of the 18 elements of $Q_{F, F'}$ can be computed in $O(1)$ time. Therefore, this entire set can be computed in $O(1)$ time. ■

Lemma 6 *Let F , F' , and F'' be three pairwise distinct internal faces of G , such that F' is on the path in $D(G)$ between F and F'' . If we are given the sets $Q_{F, F'}$ and $Q_{F', F''}$, then the set $Q_{F, F''}$ can be computed in $O(1)$ time.*

Proof. Let u be a vertex of F and let v be a vertex of F'' . Since G is an outerplanar graph, any path in G between u and v must contain at least one vertex of F' . It follows that

$$\text{dist}(u, v) = \min\{\text{dist}(u, w) + \text{dist}(w, v) \mid w \text{ is a vertex of } F'\}.$$

Thus, since $(u, w, \text{dist}(u, w), D) \in Q_{F, F'}$ and $(w, v, \text{dist}(w, v), D) \in Q_{F', F''}$, the value of $\text{dist}(u, v)$ can be computed in $O(1)$ time.

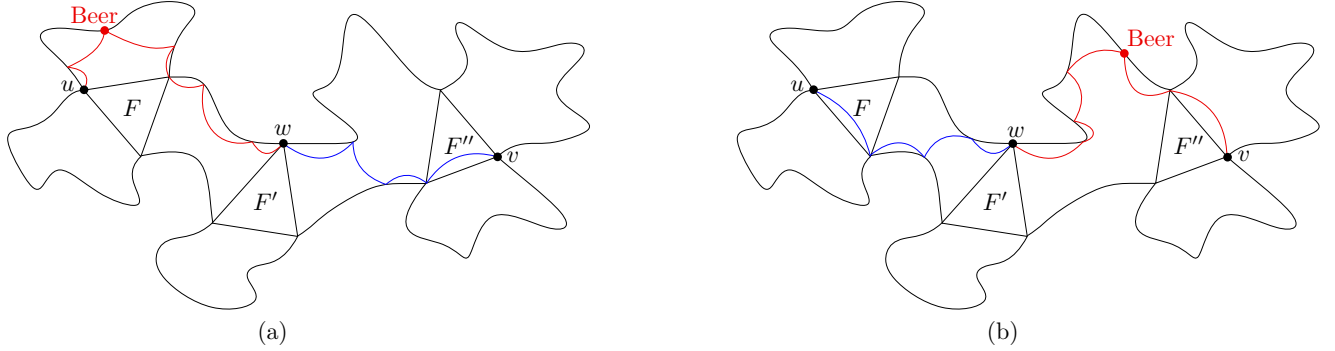


Figure 4: Any beer path from u to v contains at least one vertex of F' . In (a), we consider the shortest beer path from u to w , followed by the shortest path from w to v . In (b), we consider the shortest path from u to w , followed by the shortest beer path from w to v .

By a similar argument, $\text{dist}_B(u, v)$ is equal to (refer to Figure 4)

$$\min\{\min(\text{dist}_B(u, w) + \text{dist}(w, v), \text{dist}(u, w) + \text{dist}_B(w, v)) : w \text{ is a vertex of } F'\}.$$

All values $\text{dist}(u, w)$, $\text{dist}(w, v)$, $\text{dist}_B(u, w)$, and $\text{dist}_B(w, v)$ are encoded in the sets $Q_{F, F'}$ and $Q_{F', F''}$. Therefore, we can compute $\text{dist}_B(u, v)$ in $O(1)$ time.

Thus, since each of the 18 elements of $Q_{F, F''}$ can be computed in $O(1)$ time, the entire set can be computed in $O(1)$ time. \blacksquare

We define

$$W = \{Q_{F, F'} \mid F \text{ and } F' \text{ are distinct internal faces of } G\} \cup \{\perp\},$$

where \perp is a special symbol. We define the operator $\oplus : W \times W \rightarrow W$ in the following way.

1. If F and F' are distinct internal faces of G , then $Q_{F, F'} \oplus Q_{F, F'} = Q_{F, F'}$.
2. If F , F' , and F'' are pairwise distinct internal faces of G such that F' is on the path in $D(G)$ between F and F'' , then $Q_{F, F'} \oplus Q_{F', F''} = Q_{F, F''}$.
3. In all other cases, the operator \oplus returns \perp .

It is not difficult to verify that \oplus is associative, implying that (W, \oplus) is a semigroup. By Lemma 5, we can compute $Q_{F,F'}$ for all edges (F, F') of $D(G)$, in $O(n)$ total time.

Recall from Lemma 1 that, after an $O(n)$ -time preprocessing, we can decide in $O(1)$ time, for any three internal faces F , F' , and F'' of G , whether F' is on the path in $D(G)$ between F and F'' . Therefore, using Lemma 6, the operator \oplus takes $O(1)$ time to evaluate for any two elements of W .

Finally, let F and F' be two distinct internal faces of G , and let $F = F_0, F_1, F_2, \dots, F_k = F'$ be the path in $D(G)$ between F and F' . Then $Q_{F,F'} = \oplus_{i=0}^{k-1} Q_{F_i, F_{i+1}}$. Thus, if we store with each edge of the tree $D(G)$, the corresponding element of the semigroup, then computing $Q_{F,F'}$ becomes a path-sum query as in Section 2.2.

To summarize, all conditions to apply Lemma 3 are satisfied. As a result, we have proved Theorem 1 for maximal outerplanar graphs that satisfy the generalized triangle inequality.

3.1 The Result in Theorem 1 is Optimal

In Section 2.2, see also Remark 1, we have seen path-minimum queries in a tree, in which each edge e stores a real number $s(e)$. In such a query, we are given two distinct nodes u and v , and have to return the smallest value $s(e)$ among all edges e on the path between u and v . Lemma 3 gives a trade-off between the preprocessing and query times when answering such queries.

Let D be an arbitrary data structure that answers beer distance queries in any beer tree. Let $P(n)$, $S(n)$, and $Q(n)$ denote the preprocessing time, space, and query time of D , respectively, when the beer tree has n nodes. We will show that D can be used to answer path-minimum queries. This reduction, combined with Remark 1, implies that the result of Theorem 1 is optimal (even for beer trees).

Consider an arbitrary tree T with n nodes, such that each edge e stores a real number $s(e)$. We may assume without loss of generality that $0 < s(e) < 1$ for each edge e of T .

By making an arbitrary node the root of T , the number of edges on the path in T between two nodes u and v is equal to

$$\text{level}(u) + \text{level}(v) - 2 \cdot \text{level}(\text{LCA}(u, v)).$$

Thus, by Lemma 1, after an $O(n)$ -time preprocessing, we can compute the number of edges on this path in $O(1)$ time.

We create a beer tree T' as follows. Initially, T' is a copy of T . For each edge $e = (u, v)$ of T' , we introduce a new node x_e and replace e by two edges (u, x_e) and (v, x_e) ; we assign a weight of 1 to each of these two edges. In the current tree T' , none of the nodes has a beer store. For every node x_e in T' , we introduce a new node x'_e , add the edge (x_e, x'_e) , assign a weight of $s(e)$ to this edge, and make x'_e a beer store. Finally, we construct the data structure D for the resulting beer tree T' . Since T' has $n + 2(n - 1) = 3n - 2$ nodes, it takes $P(3n - 2) + O(n)$ time to construct D from the input tree T . Moreover, the amount of space used is $S(3n - 2) + O(n)$.

Let u and v be two distinct nodes in the original tree T , let π be the path in T between u and v , and let ℓ be the number of edges on π . The corresponding path π' in T' between u and v has weight 2ℓ .

For any edge e of T , let π'_e be the beer path in T' that starts at u , goes to x_e , then goes to x'_e and back to x_e , and continues to v .

If e is an edge of π , then the weight of π'_e is equal to $2\ell + 2 \cdot s(e)$, which is less than $2\ell + 2$. On the other hand, if e is an edge of T that is not on π , then the weight of π'_e is at least $2\ell + 2 + 2 \cdot s(e)$, which is larger than $2\ell + 2$. It follows that the shortest beer path in T' between u and v visits the beer store x'_e , where e is the edge on π for which $s(e)$ is minimum.

Thus, by computing ℓ and querying D for the beer distance in T' between u and v , we obtain the smallest value $s(e)$ among all edges e on the path in T between u and v . The query time is $Q(3n - 2) + O(1)$.

By combining this reduction with Remark 1, it follows that the result of Theorem 1 is optimal.

4 Reporting Shortest Beer Paths in Maximal Outerplanar Graphs

Let G be a maximal outerplanar beer graph with n vertices that satisfies the generalized triangle inequality. In this section, we show that, after an $O(n)$ -time preprocessing, we can report, for any two query vertices s and t , the shortest beer path $\text{SP}_B(s, t)$ from s to t , in $O(L)$ time, where L is the number of vertices on this path. As before, $D(G)$ denotes the dual of G .

Observation 2 *Let v be a vertex of G . The faces of G containing v form a path of nodes in $D(G)$.*

Define P_v to be the path in $D(G)$ formed by the faces of G containing the vertex v . Let $G[P_v]$ be the subgraph of G induced by the faces of G containing v . Note that $G[P_v]$ has a fan shape. Let $\text{CW}(v)$ denote the clockwise neighbor of v in $G[P_v]$ and let $\text{CCW}(v)$ denote the counterclockwise neighbor of v in $G[P_v]$. We will refer to the clockwise path from $\text{CW}(v)$ to $\text{CCW}(v)$ in $G[P_v]$ as the v -chain and denote it by ρ_v . (Refer to Figure 5.)

Lemma 7 *After an $O(n)$ -time preprocessing, we can answer the following queries, for any three query vertices v , u , and w , such that both u and w are on the v -chain ρ_v :*

1. Report the weight $\text{dist}(u, w, \rho_v)$ of the path from u to w along ρ_v in $O(1)$ time.
2. Report the path $\text{SP}(u, w, \rho_v)$ from u to w along ρ_v in $O(L)$ time, where L is the number of vertices on this path.

Proof. For any vertex v and any vertex u on ρ_v , we store the weight of the path from u to $\text{CW}(v)$ along ρ_v . Observe that

$$\text{dist}(u, w, \rho_v) = |\text{dist}(u, \text{CW}(v), \rho_v) - \text{dist}(w, \text{CW}(v), \rho_v)|.$$

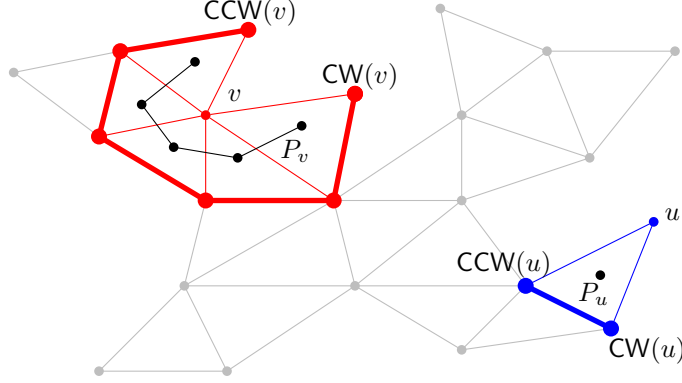


Figure 5: A maximal outerplanar graph G . The subgraphs $G[P_v]$ and $G[P_u]$ are shown in red and blue, respectively. Both the v -chain ρ_v and the u -chain ρ_u are shown in bold. Both paths P_v and P_u are shown in black. Observe that P_u is a single node.

Any exterior edge in G is in exactly one chain and any interior edge in G is in exactly two chains. Thus, the sum of the number of edges on each chain is proportional to the number of edges of G , which is $O(n)$. ■

Lemma 8 *After an $O(n)$ -time preprocessing, we can answer the following query in $O(1)$ time: Given three query vertices v , u , and w , such that both u and w are vertices of $G[P_v]$, report $\text{dist}(u, w)$, i.e., the distance between u and w in G .*

Proof. We get the following cases; the correctness follows from the generalized triangle inequality:

1. If $u = w$ then $\text{dist}(u, w) = 0$.
2. If $u = v$ then (u, w) is an edge and we return $\omega(u, w)$. Similarly if $w = v$, we return $\omega(u, w)$.
3. Otherwise u and w are both on ρ_v and we return $\min(\text{dist}(u, w, \rho_v), \omega(u, v) + \omega(v, w))$.

■

Lemma 9 *After an $O(n)$ -time preprocessing, we can report, for any three vertices v , u , and w , such that both u and w are vertices of $G[P_v]$, $\text{SP}(u, w)$ in $O(L)$ time, where L is the number of vertices on the path.*

Proof. Using Lemma 8, we can determine in $O(1)$ if the shortest path from u to w goes through v or follows the v -chain ρ_v . (Refer to Figure 6). If it goes through v , then $\text{SP}(u, w) = (u, v, w)$. Otherwise, $\text{SP}(u, w)$ takes the path along ρ_v and by Lemma 7, we can find this path in $O(L)$ time. ■

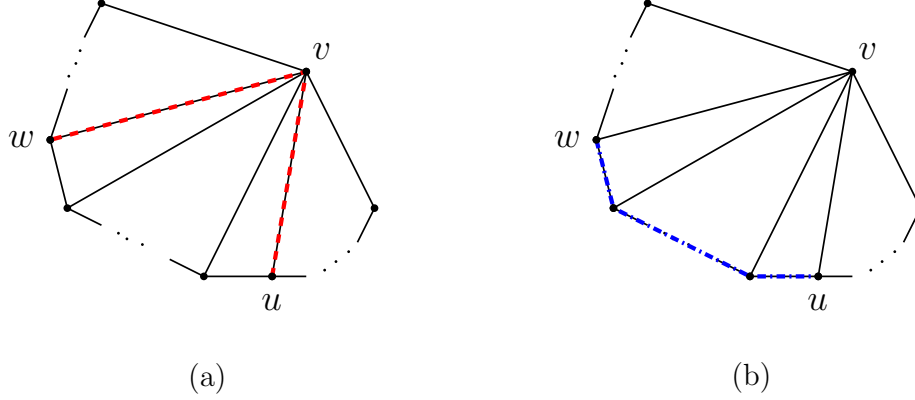


Figure 6: Two possible cases for the shortest path between u and w : (a) it goes through vertex v (shown in dashed red), or (b) it goes through the vertices of the v -chain between u and w (shown in dashed blue).

Lemma 10 *After an $O(n)$ -time preprocessing, we can report, for any three vertices v , u and w , such that both u and w are vertices of $G[P_v]$, the beer distance $\text{dist}_B(u, w)$ in $O(1)$ time. The corresponding shortest beer path $\text{SP}_B(u, w)$ can be reported in $O(L)$ time, where L is the number of vertices on the path.*

Proof. Recall from Lemma 4 that we can compute $\text{dist}_B(u, v)$ for every edge (u, v) in G , and $\text{dist}_B(v, v)$ for every vertex v in G , in $O(n)$ time.

Let $\rho_v = (\text{CW}(v) = u_1, u_2, \dots, u_N = \text{CCW}(v))$. Let $A_v[\]$ be an array of size $N - 1$. For $i = 1, \dots, N - 1$, we set $A_v[i] = \text{dist}_B(u_i, u_{i+1}) - \omega(u_i, u_{i+1})$. Recall that by the generalized triangle inequality, $\omega(u_i, u_{i+1}) = \text{dist}(u_i, u_{i+1})$. Therefore, $A[i]$ holds the difference between the weights of the shortest path from u_i to u_{i+1} and the shortest beer path from u_i to u_{i+1} . After preprocessing the array $A_v[\]$ in $O(N)$ time, we can conduct range minimum queries in $O(1)$ time. (Bender and Farach-Colton [2] show that these queries are equivalent to LCA-queries in the Cartesian tree of the array.) Thus, for each v -chain of N nodes, we spend $O(N)$ time processing the v -chain. Since every edge is in at most two chains, processing all v -chains takes $O(n)$ time and space.

Given two vertices u and w of $G[P_v]$, we determine the beer distance $\text{dist}_B(u, w)$ as follows:

1. If $u = w$ then $\text{dist}_B(u, w)$ has already been computed by Lemma 4.
2. If $u = v$ or $w = v$, then there is an edge from v to the other vertex. Thus, $\text{dist}_B(u, w)$ has already been computed by Lemma 4.
3. Otherwise, u , w and v are three distinct vertices. Assume without loss of generality that w is clockwise from u on the v -chain. We take the minimum of the following two cases:
 - (a) The shortest beer path from u to w that goes through v . Since a beer store must be visited before or after v , this beer path has a weight of $\min(\text{dist}_B(u, v) + \omega(v, w), \omega(u, v) + \text{dist}_B(v, w))$.

- (b) The shortest beer path through the vertices of the v -chain. Note that this beer path will visit each vertex on the v -chain between u and w , but may go off the v -chain to visit a beer store. On $\text{SP}_B(u, w)$, there is one pair of vertices, u_i and u_{i+1} , such that a beer path is taken between u_i and u_{i+1} , and u_i and u_{i+1} are adjacent on the v -chain; refer to Figure 7. The shortest path is taken between all other pairs of adjacent vertices on the v -chain. From Lemma 7, we can compute $\text{dist}(u, w, \rho_v)$ in $O(1)$ time. The shortest beer path through the vertices of the v -chain has a weight of $\text{dist}(u, w, \rho_v) + A_v[i]$, where $A_v[i]$ is the additional distance needed to visit a beer store between u_i and u_{i+1} . Let u be the j^{th} vertex on ρ_v and let w be the k^{th} vertex in ρ_v . Then $A_v[i]$ is the minimum value in $A_v[j, \dots, k-1]$. We can determine $A_v[i]$ in constant time using a range minimum query.

Note that in case 1 and case 2, $\text{SP}_B(u, w)$ can be constructed in $O(L)$ time by Lemma 4. For case 3 (a) let $p = (u, v, w)$ and for case 3 (b) let $p = \text{SP}(u, w, \rho_v)$. Let u_i, u_{i+1} be the pair of adjacent vertices on p between which a beer path was taken. Using Lemma 4 we can find $\text{SP}_B(u_i, u_{i+1})$ in $O(L)$ time. We obtain $\text{SP}_B(u, w)$ by replacing the edge (u_i, u_{i+1}) in p with $\text{SP}_B(u_i, u_{i+1})$. ■

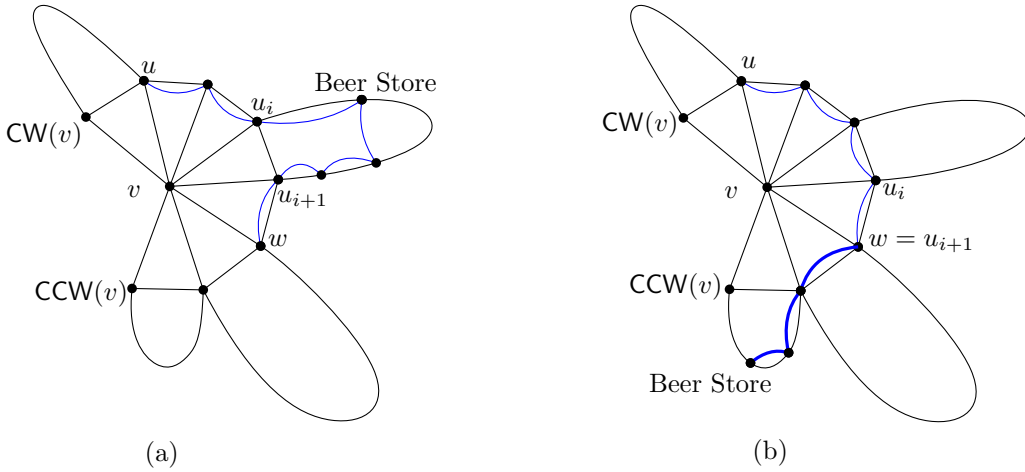


Figure 7: Both figures show a shortest beer path from u to w through the vertices on the v -chain. Thicker edges on the blue beer path are edges that are traversed twice; once in each direction.

4.1 Answering Shortest Beer Path Queries

Recall that, for any vertex v of G , P_v denotes the path in $D(G)$ formed by the faces of G containing v . Moreover, $G[P_v]$ denotes the subgraph of G induced by these faces.

Consider two query vertices s and t of G . Our goal is to compute the shortest beer path $\text{SP}_B(s, t)$.

Let F_s and F_t be arbitrary faces containing s and t , respectively. If t is in $G[P_s]$ then, by Lemma 10, we can construct $\text{SP}_B(s, t)$ in $O(L)$ time. For the remainder of this section, we assume that t is not in $G[P_s]$. To find $\text{SP}_B(s, t)$, we start by constructing a directed acyclic graph (DAG), H . In this DAG, vertices will be arranged in columns of constant size, and all edges go from left to right between vertices in adjacent columns. In H , each column will contain one vertex that is on $\text{SP}_B(s, t)$. First we will construct H and then we will show how we can use H to construct $\text{SP}_B(s, t)$. The entire construction is illustrated in Figure 8.

Observation 3 *Any interior edge (a, b) of G splits G into two subgraphs such that if s is in one subgraph and t is in the other, then any path in G from s to t must visit at least one of a and b .*

Let P be the unique path between F_s and F_t in $D(G)$. Consider moving along P from F_s to F_t . Let F_1 be the node on P_s that is closest to F_t , and let F'_1 be the successor of F_1 on P . Note that, by Lemmas 1 and 2, we can find F_1 and F'_1 in $O(1)$ time.¹ Let $e_1 = (a_1, b_1)$ be the edge in G shared by the faces F_1 and F'_1 . Since $\text{SP}_B(s, t)$ must visit both of these faces, by Observation 3, at least one of a_1 or b_1 is on the shortest beer path.

We place s in the first column of H and a_1 and b_1 in the second column of H . We then add two directed edges from s to a_1 , one with weight $\text{dist}(s, a_1)$ and the other with weight $\text{dist}_B(s, a_1)$. Similarly, we add two directed edges from s to b_1 with weights $\text{dist}(s, b_1)$ and $\text{dist}_B(s, b_1)$.

When $i \geq 2$ we construct the $(i + 1)^{\text{th}}$ column of H in the following way. Let $e_{i-1} = (a_{i-1}, b_{i-1})$ be the edge shared by the faces F_{i-1} and F'_{i-1} . The i^{th} column of H contains the vertices a_{i-1} and b_{i-1} . Note that F'_{i-1} is in both $P_{b_{i-1}}$ and $P_{a_{i-1}}$. Using Lemma 2, we find the node F_i^b on $P_{b_{i-1}}$ that is closest to F_t . If the vertex a_{i-1} is not in F_i^b , then we let $F_i = F_i^b$. Otherwise, we let F_i be the node on $P_{a_{i-1}}$ that is closest to F_t .

If t is not a vertex of F_i , then let F'_i be the node that follows F_i on P ; we find F'_i using Lemma 1. Let $e_i = (a_i, b_i)$ be the edge of G shared by the faces F_i and F'_i . In the $(i + 1)^{\text{th}}$ column, we place a_i and b_i . For each $u \in \{a_{i-1}, b_{i-1}\}$ and each $v \in \{a_i, b_i\}$ we add two directed edges (u, v) to the DAG, one with weight $\text{dist}(u, v)$ and the other with weight $\text{dist}_B(u, v)$. If F_i is in $P_{a_{i-1}}$, all these vertices are in $G[P_{a_{i-1}}]$; otherwise, F_i is in $P_{b_{i-1}}$, and all these vertices are in $G[P_{b_{i-1}}]$. Thus, by Lemmas 8 and 10, we can find the distances and beer distances to assign to these edges in constant time.

If t is in F_i , then in the $(i + 1)^{\text{th}}$ column we only place the vertex t . In this case, for each $u \in \{a_{i-1}, b_{i-1}\}$, we add two directed edges (u, t) to the DAG with weights $\text{dist}(u, t)$ and $\text{dist}_B(u, t)$. At this point we are done constructing H .

We define a *beer edge* to be an edge of H that was assigned a weight of a beer path during the construction of H . We find the beer distance from s to t in G using the following dynamic programming approach in H .

¹To apply Lemma 2, we consider each vertex of G to be a color. For each vertex v of G , the v -colored path in the tree $D(G)$ is the path P_v . The face F_1 is the answer to the closest-color query with nodes F_s and F_t and color s .

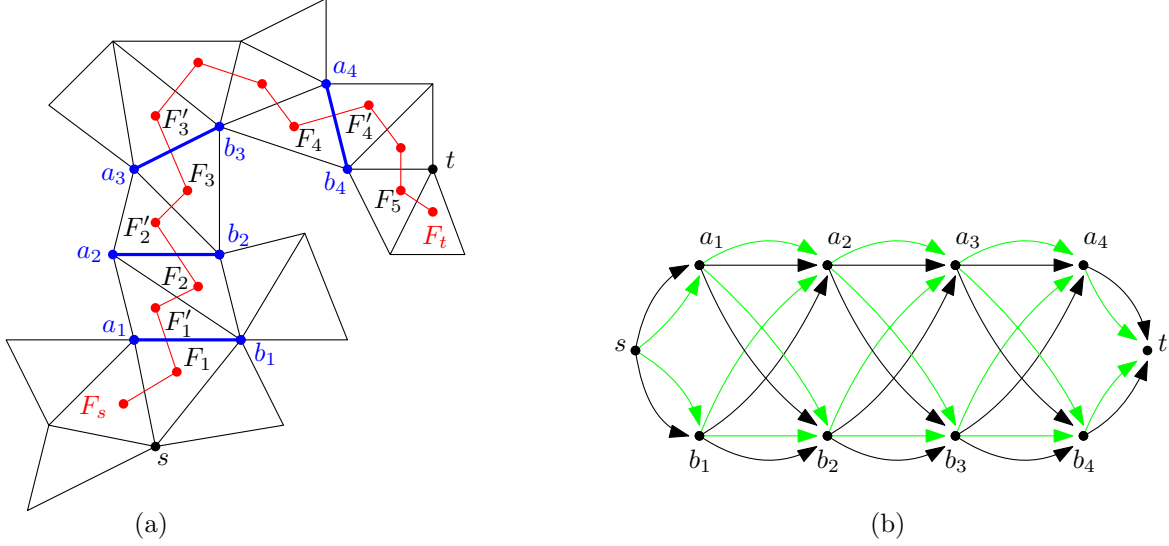


Figure 8: An outerplanar graph G (a) and the DAG H constructed for the shortest beer path query from s to t (b). The path P from F_s to F_t is shown in red. Each edge $e_i = (a_i, b_i)$ such that e_i is shared by F_i and F'_i is shown in blue. The green edges of H represent the beer edges.

Let M denote the number of columns in H . For $i = 3, \dots, M$ and for all u in the i^{th} column of H , compute

$$\text{dist}_B(s, u) = \min \begin{cases} \text{dist}_B(s, a_{i-2}) + \text{dist}(a_{i-2}, u), \\ \text{dist}(s, a_{i-2}) + \text{dist}_B(a_{i-2}, u), \\ \text{dist}_B(s, b_{i-2}) + \text{dist}(b_{i-2}, u), \\ \text{dist}(s, b_{i-2}) + \text{dist}_B(b_{i-2}, u) \end{cases}$$

and

$$\text{dist}(s, u) = \min \begin{cases} \text{dist}(s, a_{i-2}) + \text{dist}(a_{i-2}, u), \\ \text{dist}(s, b_{i-2}) + \text{dist}(b_{i-2}, u). \end{cases}$$

The vertices a_{i-2} and b_{i-2} occur in the $(i-1)^{\text{th}}$ column. Thus, $\text{dist}_B(s, a_{i-2})$, $\text{dist}_B(s, b_{i-2})$, $\text{dist}(s, a_{i-2})$, and $\text{dist}(s, b_{i-2})$ will be computed before computing the values for the i^{th} column. We get $\text{dist}(a_{i-2}, u)$, $\text{dist}_B(a_{i-2}, u)$, $\text{dist}(b_{i-2}, u)$ and $\text{dist}_B(b_{i-2}, u)$ from the weights of the DAG-edges between the $(i-1)^{\text{th}}$ and i^{th} columns of H .

By keeping track of which expression produced $\text{dist}_B(s, u)$ and $\text{dist}(s, u)$, we can backwards reconstruct the shortest beer path in the DAG. Knowing the shortest beer path in the DAG enables us to construct the corresponding beer path in G as follows.

1. Define P_{st} to be an empty path.
2. For each edge (w, v) of the shortest beer path in the DAG.

- (a) If (w, v) was a beer edge, let $P_{wv} = \text{SP}_B(w, v)$, which can be constructed in time proportional to its number of vertices via Lemma 10.
- (b) Otherwise, let $P_{wv} = \text{SP}(w, v)$ which can be constructed in time proportional to its number of vertices as seen in Lemma 9.

Let $P_{st} = P_{st} \cup P_{wv}$.

3. Return P_{st} , which is equal to $\text{SP}_B(w, v)$.

Let L denote the number of vertices on $\text{SP}_B(s, t)$. In order for the above query algorithm to take $O(L)$ time, the size of the DAG must be $O(L)$. The following three lemmas will show this to be true.

Lemma 11 *For $2 \leq i < M - 1$, F_i contains either a_{i-1} or b_{i-1} , but not both.*

Proof. Recall that we defined F_i^b to be the last node on P that is also on $P_{b_{i-1}}$. We similarly define F_i^a to be the last node on P that is also on $P_{a_{i-1}}$. From the way we choose F_i , F_i is either F_i^b or F_i^a . We only choose $F_i = F_i^b$ after having checked that a_{i-1} is not in F_i^b ; thus in this case we can be sure that F_i only contains b_{i-1} .

Assume for the purpose of contradiction that we choose $F_i = F_i^a$ and b_{i-1} is also in F_i . Let the third vertex of F_i be c . Let the face on P immediately following F_i be F'_i . The edge shared by F_i and F'_i is either (b_{i-1}, c) or (a_{i-1}, c) . If (b_{i-1}, c) is the shared edge, then F'_i is a face closer to F_t that contains b_{i-1} and not a_{i-1} , so we would have chosen $F_i = F_i^b$, which is a contradiction. Otherwise, (a_{i-1}, c) is the edge shared by F_i and F'_i , which implies that there is a face containing a_{i-1} closer to F_t in P than F_i^a , which contradicts the definition of F_i^a . ■

Lemma 12 *Every vertex of G appears in at most one column of H .*

Proof. Since (a_1, b_1) is an edge shared by both the last face of P containing s and the first face of P that does not contain s it is not possible for either of these vertices to be the vertex s . Thus, s will only be represented by the vertex in the first column of H . By stopping the construction of H as soon as we add a vertex representing t , we ensure that H only contains one vertex corresponding to the vertex t in G .

For $2 \leq i \leq M - 2$, consider the vertex a_{i-1} in G represented by a vertex in the i^{th} column of H . If $F_i = F_i^a$ then by definition of F_i^a , F'_i does not contain a_{i-1} . Since (a_i, b_i) is an edge of F'_i , $a_i \neq a_{i-1}$ and $b_i \neq a_{i-1}$. Because the face F'_i is closer to F_t than F_i^a , a_{i-1} is not a vertex on any of the faces on the path from F'_i to F_t . Thus, subsequent columns of H will not contain vertices representing the vertex a_{i-1} in G .

If $F_i = F_i^b$ then by Lemma 11, a_{i-1} is not in F_i and since (a_i, b_i) is an edge of F_i , $a_i \neq a_{i-1}$ and $b_i \neq a_{i-1}$. Because F_i is a face on P closer to F_t than F_{i-1} (a face that contains a_{i-1}) it follows from Observation 2 that none of the faces on P from F_{i-1} to F_t will have the vertex a_{i-1} on their face and, thus, a_{i-1} will not be represented by vertices in subsequent columns of H .

By switching the roles of a_{i-1} with b_{i-1} in the above reasoning we can see that this also holds for b_{i-1} . \blacksquare

Lemma 13 *The number of vertices and edges of H is $O(L)$.*

Proof. Recall that M denotes the number of columns of the DAG H , whereas L denotes the number of vertices on $\text{SP}_B(s, t)$. The vertices s and t (which are in the first and last columns of H , respectively) are obviously on $\text{SP}_B(s, t)$. By Observation 3, this beer path must contain at least one vertex from each of the columns $2, 3, \dots, M-1$ of H . By Lemma 12, the vertices in all columns of H are pairwise distinct. Therefore, $L \geq M$. Since each column has at most two vertices, and each of these vertices has at most four outgoing edges, the total number of vertices and edges of H is $O(M)$, which is $O(L)$. \blacksquare

Observe that the total preprocessing time is $O(n)$. For two query vertices s and t , the DAG, H , can be constructed in $O(L)$ time. Finally, the dynamic programming algorithm on H takes $O(L)$ time. Thus, we have proved Theorem 2 for maximal outerplanar graphs that satisfy the generalized triangle inequality.

5 Proof of Lemma 4

Let G be a maximal outerplanar beer graph with n vertices that satisfies the generalized triangle inequality. We will first show how to compute $\text{dist}_B(u, u)$ for each vertex u of G , and $\text{dist}_B(u, v)$ for each edge (u, v) of G . Consider again the dual $D(G)$ of G . We choose an arbitrary face of G and make it the root of $D(G)$.

Let (u, v) be any edge of G . This edge divides G into two outerplanar subgraphs, both of which contain (u, v) as an edge. Let G_{uv}^R be the subgraph that contains the face represented by the root of $D(G)$, and let G_{uv}^{-R} denote the other subgraph. Note that if (u, v) is an external edge, then $G_{uv}^R = G$ and G_{uv}^{-R} consists of the single edge (u, v) .

By the generalized triangle inequality, the shortest beer path between u and v is completely in G_{uv}^R or completely in G_{uv}^{-R} . The same is true for the shortest beer path from u to itself. This implies:

Observation 4 *For each edge (u, v) of G ,*

1. $\text{dist}_B(u, v) = \min(\text{dist}_B(u, v, G_{uv}^R), \text{dist}_B(u, v, G_{uv}^{-R}))$,
2. $\text{dist}_B(u, u) = \min(\text{dist}_B(u, u, G_{uv}^R), \text{dist}_B(u, u, G_{uv}^{-R}))$,
3. $\text{dist}_B(v, v) = \min(\text{dist}_B(v, v, G_{uv}^R), \text{dist}_B(v, v, G_{uv}^{-R}))$.

Thus, it suffices to first compute $\text{dist}_B(u, v, G_{uv}^{-R})$, $\text{dist}_B(u, u, G_{uv}^{-R})$, and $\text{dist}_B(v, v, G_{uv}^{-R})$ for all edges (u, v) , and then compute $\text{dist}_B(u, v, G_{uv}^R)$, $\text{dist}_B(u, u, G_{uv}^R)$, and $\text{dist}_B(v, v, G_{uv}^R)$, again for all edges (u, v) .

5.1 Recurrences for $\text{dist}_B(u, v, G_{uv}^{-R})$, $\text{dist}_B(u, u, G_{uv}^{-R})$, and $\text{dist}_B(v, v, G_{uv}^{-R})$

Let (u, v) be an edge of G . Item 1. below presents the base cases, whereas item 2. gives the recurrences.

1. Assume that (u, v) is an external edge of G .
 - (a) If both u and v are beer stores, then $\text{dist}_B(u, v, G_{uv}^{-R}) = \omega(u, v)$, $\text{dist}_B(u, u, G_{uv}^{-R}) = 0$, and $\text{dist}_B(v, v, G_{uv}^{-R}) = 0$.
 - (b) If exactly one of u and v , say u , is a beer store, then $\text{dist}_B(u, v, G_{uv}^{-R}) = \omega(u, v)$, $\text{dist}_B(u, u, G_{uv}^{-R}) = 0$, and $\text{dist}_B(v, v, G_{uv}^{-R}) = 2 \cdot \omega(u, v)$.
 - (c) If neither u nor v is a beer store, then $\text{dist}_B(u, v, G_{uv}^{-R}) = \infty$, $\text{dist}_B(u, u, G_{uv}^{-R}) = \infty$, and $\text{dist}_B(v, v, G_{uv}^{-R}) = \infty$.
2. Assume that (u, v) is an internal edge of G . Let w be the third vertex of the face of G_{uv}^{-R} that contains (u, v) as an edge. All possible cases are illustrated in Figure 9.
 - (a) The value of $\text{dist}_B(u, v, G_{uv}^{-R})$ is the minimum of
 - i. $\text{dist}_B(u, w, G_{uw}^{-R}) + \omega(w, v)$,
 - ii. $\omega(u, w) + \text{dist}_B(w, v, G_{vw}^{-R})$,
 - iii. $\text{dist}_B(u, u, G_{uw}^{-R}) + \omega(u, v)$,
 - iv. $\omega(u, v) + \text{dist}_B(v, v, G_{vw}^{-R})$.
 - (b) The value of $\text{dist}_B(u, u, G_{uv}^{-R})$ is the minimum of
 - i. $\text{dist}_B(u, u, G_{uw}^{-R})$,
 - ii. $2 \cdot \omega(u, w) + \text{dist}_B(w, w, G_{vw}^{-R})$,
 - iii. $2 \cdot \omega(u, v) + \text{dist}_B(v, v, G_{vw}^{-R})$.

The value of $\text{dist}_B(v, v, G_{uv}^{-R})$ is obtained by swapping u and v in i., ii., and iii.

These recurrences express $\text{dist}_B(u, v, G_{uv}^{-R})$, $\text{dist}_B(u, u, G_{uv}^{-R})$, and $\text{dist}_B(v, v, G_{uv}^{-R})$ in terms of values that are “further down” in the tree $D(G)$. Therefore, by performing a post-order traversal of $D(G)$, we obtain all these values, for all edges (u, v) of G , in $O(n)$ total time.

5.2 Recurrences for $\text{dist}_B(u, v, G_{uv}^R)$, $\text{dist}_B(u, u, G_{uv}^R)$, and $\text{dist}_B(v, v, G_{uv}^R)$

Let (u, v) be an edge of G . Item 1. below presents the base cases, whereas item 2. gives the recurrences.

1. Assume that (u, v) is an edge of the face representing the root of $D(G)$. Let w be the third vertex of this face.
 - (a) The value of $\text{dist}_B(u, v, G_{uv}^R)$ is the minimum of
 - i. $\text{dist}_B(u, u, G_{uw}^{-R}) + \omega(u, v)$,

- ii. $\text{dist}_B(u, w, G_{uw}^{-R}) + \omega(w, v)$,
- iii. $\omega(u, v) + \text{dist}_B(v, v, G_{vw}^{-R})$,
- iv. $\omega(u, w) + \text{dist}_B(w, v, G_{vw}^{-R})$.

(b) The value of $\text{dist}_B(u, u, G_{uw}^R)$ is the minimum of²

- i. $\text{dist}_B(u, u, G_{uw}^{-R})$,
- ii. $2 \cdot \omega(u, w) + \text{dist}_B(w, w, G_{vw}^{-R})$,
- iii. $2 \cdot \omega(u, v) + \text{dist}_B(v, v, G_{vw}^{-R})$.

The value of $\text{dist}_B(v, v, G_{uv}^R)$ is obtained by swapping u and v in i., ii., and iii.

2. Assume that (u, v) is not an edge of the face represented by the root of $D(G)$. Let w be the third vertex of the face of G_{uw}^R that contains (u, v) as an edge. We may assume without loss of generality that (v, w) is an edge of the face represented by the parent of the face representing (u, v, w) . All possible cases are illustrated in Figure 10.

(a) The value of $\text{dist}_B(u, v, G_{uv}^R)$ is the minimum of

- i. $\text{dist}_B(u, u, G_{uw}^{-R}) + \omega(u, v)$,
- ii. $\text{dist}_B(u, w, G_{uw}^{-R}) + \omega(w, v)$,
- iii. $\omega(u, v) + \text{dist}_B(v, v, G_{vw}^R)$,
- iv. $\omega(u, w) + \text{dist}_B(w, v, G_{vw}^R)$.

(b) The value of $\text{dist}_B(u, u, G_{uv}^R)$ is the minimum of

- i. $\text{dist}_B(u, u, G_{uw}^{-R})$,
- ii. $2 \cdot \omega(u, v) + \text{dist}_B(v, v, G_{vw}^R)$,
- iii. $2 \cdot \omega(u, w) + \text{dist}_B(w, w, G_{vw}^R)$.

(c) The value of $\text{dist}_B(v, v, G_{uv}^R)$ is the minimum of

- i. $\text{dist}_B(v, v, G_{vw}^R)$,
- ii. $2 \cdot \omega(v, u) + \text{dist}_B(u, u, G_{uw}^R)$,
- iii. $2 \cdot \omega(v, w) + \text{dist}_B(w, w, G_{uw}^R)$.

These recurrences express $\text{dist}_B(u, v, G_{uv}^R)$, $\text{dist}_B(u, u, G_{uv}^R)$, and $\text{dist}_B(v, v, G_{uv}^R)$ in terms of values that are “higher up” in the tree $D(G)$ and values that involve graphs with a superscript “ $-R$ ”. These latter values have been computed already. Thus, by performing a pre-order traversal of $D(G)$, we obtain all values $\text{dist}_B(u, v, G_{uv}^R)$, $\text{dist}_B(u, u, G_{uv}^R)$, and $\text{dist}_B(v, v, G_{uv}^R)$, for all edges (u, v) of G , in $O(n)$ total time. This completes the proof of the first claim in Lemma 4.

To prove the second claim in Lemma 4, consider $\text{dist}_B(u, v)$ where $u = v$ or (u, v) is an edge of G . If u or v is a beer store, then store *nil* with $\text{dist}_B(u, v)$.

²We do not have to consider $W := \omega(u, v) + \text{dist}_B(v, w, G_{vw}^{-R}) + \omega(w, u)$, because the sum of the values in ii. and iii. is at most $2W$. Therefore, the smaller of the values in ii. and iii. is at most W .

The values $\text{dist}_B(u, v, G_{uv}^R)$ and $\text{dist}_B(u, v, G_{uv}^{-R})$ are computed as the minimum of a set of path weights from u to v through a vertex x such that x is adjacent to both u and v or x is equal to one of these vertices and adjacent to the other. Either the subpath from u to x is a beer path or the subpath from x to v is a beer path. Whenever we take the minimum of a set of path weights in the above computation, we store with that distance the vertex x and a bit to indicate which subpath is the beer path. When $u = v$ we can arbitrarily choose which subpath is the beer path. After taking the minimum of $\text{dist}_B(u, v, G_{uv}^R)$ and $\text{dist}_B(u, v, G_{uv}^{-R})$, we are left with a vertex, x , on the shortest beer path from u to v and the bit indicating which subpath is a beer path.

We recursively compute $\text{SP}_B(u, v)$ where either (u, v) is an edge of G or $u = v$ as follows.

1. If nil is stored with $\text{dist}_B(u, v)$ and $u = v$, then $\text{SP}_B(u, v) = (u)$.
2. If nil is stored with $\text{dist}_B(u, v)$ and (u, v) is an edge, then $\text{SP}_B(u, v) = (u, v)$.
3. If a vertex x is stored with (u, v) and the subpath from u to x is a beer path, then recursively compute $\text{dist}_B(u, x)$, and set $\text{SP}_B(u, v) = (\text{SP}_B(u, x), v)$.
4. Otherwise x is stored with (u, v) and the subpath from x to v is a beer path. Recursively compute $\text{dist}_B(x, v)$, and set $\text{SP}_B(u, v) = (u, \text{SP}_B(x, v))$.

Note that a constant amount of work is done at each level of the recurrence excluding the time spent in recursive calls. In each recursive call, except potentially the last call, we get one new vertex on the shortest beer path. Thus, constructing the whole path requires a total of $O(L)$ time.

6 Extension to Arbitrary Outerplanar Graphs

6.1 Conversion to a Maximal Outerplanar Graph

Let G be an outerplanar beer graph with n vertices. Assume that the outer face of G is not a Hamiltonian cycle. We traverse G along the outer face in a clockwise manner, and mark each vertex when we encounter it for the first time. Each time we visit a marked vertex v , we take note of v 's current counterclockwise neighbor, $\text{CCW}(v)$. Then we continue from v to the next clockwise vertex on the outer face and add an edge from this vertex to $\text{CCW}(v)$. We continue this process until we have returned to the vertex we started from and all vertices have been marked.

At this moment, the outer face is a Hamiltonian cycle. For every interior face that is not a triangle, we pick a vertex u on that face and add edges connecting u with all vertices of the face that are not already adjacent to u .

The resulting graph is a maximal outerplanar graph. Each edge that has been added is given a weight of infinity. Observe that each shortest (beer) path in the resulting graph corresponds to a shortest (beer) path in the original graph, and vice versa.

6.2 Guaranteeing the Generalized Triangle Inequality

Let G be a maximal outerplanar graph with an edge weight function ω . In order to convert G to a graph that satisfies the generalized triangle inequality, we need to compute $\text{dist}(u, v)$ for every edge (u, v) in G , and we need to be able to construct $\text{SP}(u, v)$ for each edge (u, v) .

Let $D(G)$ be the dual of G rooted at an arbitrary interior face of G . For each edge (u, v) , we initialize $\delta(u, v) = \omega(u, v)$. (At the end, $\delta(u, v)$ will be equal to $\text{dist}(u, v)$.) For each edge (u, v) , we also maintain a parent vertex, $p(u, v)$, initialized to nil .

We first conduct a post-order traversal of $D(G)$, processing each associated face in G . Then we conduct a pre-order traversal of $D(G)$, again processing each associated face in G . Let F be a face of G . We process F as follows. Let (u, v) be the edge of F that is shared with the predecessor face F' in the traversal, and let w be the third vertex of F' . If $\delta(u, v) > \delta(u, w) + \delta(w, v)$, we set $\delta(u, v) = \delta(u, w) + \delta(w, v)$ and $p(u, v) = w$.

After these traversals, $\delta(u, v) = \text{dist}(u, v)$ for every edge (u, v) in G . If $p(u, v) = \text{nil}$, then $\text{SP}(u, v) = (u, v)$; otherwise, $\text{SP}(u, v)$ is the concatenation of $\text{SP}(u, p(u, v))$ and $\text{SP}(p(u, v), v)$, both of which can be computed recursively.

7 Single Source Shortest Beer Path

In this section we will describe how to compute the single source shortest beer path from a source vertex s on a maximal outerplanar graph G that satisfies the generalized triangle inequality. In order to do this we first precompute (i) $\text{dist}_B(u, v)$ for every edge (u, v) and $\text{dist}_B(u, u)$ for every vertex u and (ii) $\text{dist}(s, v)$ for every vertex v . By Lemma 4, we can compute (i) in $O(n)$ time and in [7], Maheshwari and Zeh present a single source shortest path algorithm for undirected outerplanar graphs which gives us (ii) in $O(n)$ time.

Let $D(G)$ be the dual of G and let F_s be an arbitrary interior face of G containing s . Root $D(G)$ at the node F_s and then conduct a pre-order traversal of $D(G)$. Let F be the current node of $D(G)$ being processed during this traversal.

1. If $F = F_s$, let u and v be the vertices of F_s that are not s . Since (s, u) and (s, v) are both edges of G , $\text{dist}_B(s, s)$, $\text{dist}_B(s, u)$ and $\text{dist}_B(s, v)$ were precomputed in (i).
2. If $F \neq F_s$, let a and b be the vertices of F shared with the face F' , where F' is the parent of F in $D(G)$. This implies that by this step we have already computed $\text{dist}_B(s, a)$ and $\text{dist}_B(s, b)$. Let c be the third vertex of F . The value of $\text{dist}_B(s, c)$ is the minimum of:
 - (a) $\text{dist}(s, a) + \text{dist}_B(a, c)$,
 - (b) $\text{dist}_B(s, a) + \omega(a, c)$,
 - (c) $\text{dist}(s, b) + \text{dist}_B(b, c)$,
 - (d) $\text{dist}_B(s, b) + \omega(b, c)$.

Since (a, c) and (b, c) are edges of G , we precomputed $\text{dist}_B(a, c)$ and $\text{dist}_B(b, c)$ in (i). Lastly, we computed $\text{dist}(s, a)$ and $\text{dist}(s, b)$ in (ii), so each of the values listed above can be computed in constant time.

The correctness of this algorithm follows from Observation 3 and the generalized triangle inequality. Since we do a constant amount of work at each face in the traversal of $D(G)$ and the number of interior faces of a maximal outerplanar graph is $n - 2$, this algorithm takes $O(n)$ time.

Let L be the number of vertices on the shortest beer path from s to v . If $\text{dist}_B(s, v)$ was found in step 1, then by Lemma 4, $\text{SP}_B(s, v)$ can be constructed in $O(L)$ time. If this is not the case, then $v = c$ in some iteration of step 2. At this step we store a vertex $p(v)$ such that $p(v) = a$ if (a) or (b) was the minimum of step 2 and $p(v) = b$ otherwise. We also store a bit to indicate if the subpath from s to $p(v)$ is the shortest path (as in cases (a) and (c)) or the shortest beer path (as in cases (b) and (d)). If the subpath from s to $p(v)$ is the shortest path, we use the method described in [7] to find $\text{SP}(s, p(v))$ and use Lemma 4 to find $\text{SP}_B(p(v), v)$ and then concatenate $\text{SP}(s, p(v))$ and $\text{SP}_B(p(v), v)$ to get $\text{SP}_B(s, v)$. Both $\text{SP}(s, p(v))$ and $\text{SP}_B(p(v), v)$ are found in time proportional to the number of vertices on their paths, so this takes $O(L)$ time. If the subpath from s to $p(v)$ is the shortest beer path, then we recursively find $\text{SP}_B(s, p(v))$ and concatenate it with the edge $(p(v), v)$ (which is $\text{SP}(p(v), v)$ by the generalized triangle inequality). Each iteration of the recursive step takes time proportional to the number of new vertices of the path found in that step. Thus, we find $\text{SP}_B(s, v)$ in a total of $O(L)$ time.

8 Concluding Remarks

We have introduced a variant of the shortest path problem in weighted graphs, in which a subset of the vertices store a facility, say a beer store, and we want to compute the shortest path between two vertices, or the length of this path, that visits at least one facility. For the class of outerplanar beer graphs, we have presented data structures that are optimal in terms of preprocessing time, space usage, and query times.

As a byproduct, we have also obtained improved results for standard shortest path queries in outerplanar graphs. In particular, after an $O(m)$ -time preprocessing of an n -vertex beer outerplanar graph, we can determine the length of a shortest path between any two query vertices in $O(\alpha(m, n))$ time. We have shown that this is optimal for beer paths, even if the graph is a tree. Unfortunately, our reduction does not work for standard shortest path queries. Observe that, in any tree with n vertices, after $O(n)$ preprocessing, we can report the length $\text{dist}(u, v)$ of the path between any two query vertices u and v in $O(1)$ time, because

$$\text{dist}(u, v) = \text{dist}(u, r) + \text{dist}(v, r) - 2 \cdot \text{dist}(\text{LCA}(u, v), r),$$

where r is an arbitrarily chosen root of the tree.

Open Problem 1 *Is there a linear-size data structure that can answer standard distance queries in outerplanar graphs in $O(1)$ time?*

Since shortest beer path queries are of relevance for practical applications, such queries should be studied in larger families of graphs. Obvious examples are general planar graphs and graphs of bounded tree width.

References

- [1] N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. Technical Report 71/87, Tel-Aviv University, 1987.
- [2] M. A. Bender and M. Farach-Colton. The LCA problem revisited. In *Proceedings of the 4th Latin American Symposium on Theoretical Informatics*, volume 1776 of *Lecture Notes in Computer Science*, pages 88–94, Berlin, 2000. Springer-Verlag.
- [3] T. M. Chan, M. He, J. I. Munro, and G. Zhou. Succinct indices for path minimum, with applications. *Algorithmica*, 78(2):453–491, 2017.
- [4] B. Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987.
- [5] H. Djidjev, G. E. Pantziou, and C. D. Zaroliagis. Computing shortest paths and distances in planar graphs. In *Automata, Languages and Programming, 18th International Colloquium, ICALP91*, volume 510 of *Lecture Notes in Computer Science*, pages 327–338. Springer, 1991.
- [6] D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.
- [7] A. Maheshwari and N. Zeh. I/O-optimal algorithms for outerplanar graphs. *Journal of Graph Algorithms and Applications*, 8(1):47–87, 2004.
- [8] S. Pettie. An inverse-Ackermann type lower bound for online minimum spanning tree verification. *Combinatorica*, 26(2):207–230, 2006.
- [9] M. Thorup. Parallel shortcutting of rooted trees. *Journal of Algorithms*, 32:139–159, 1997.

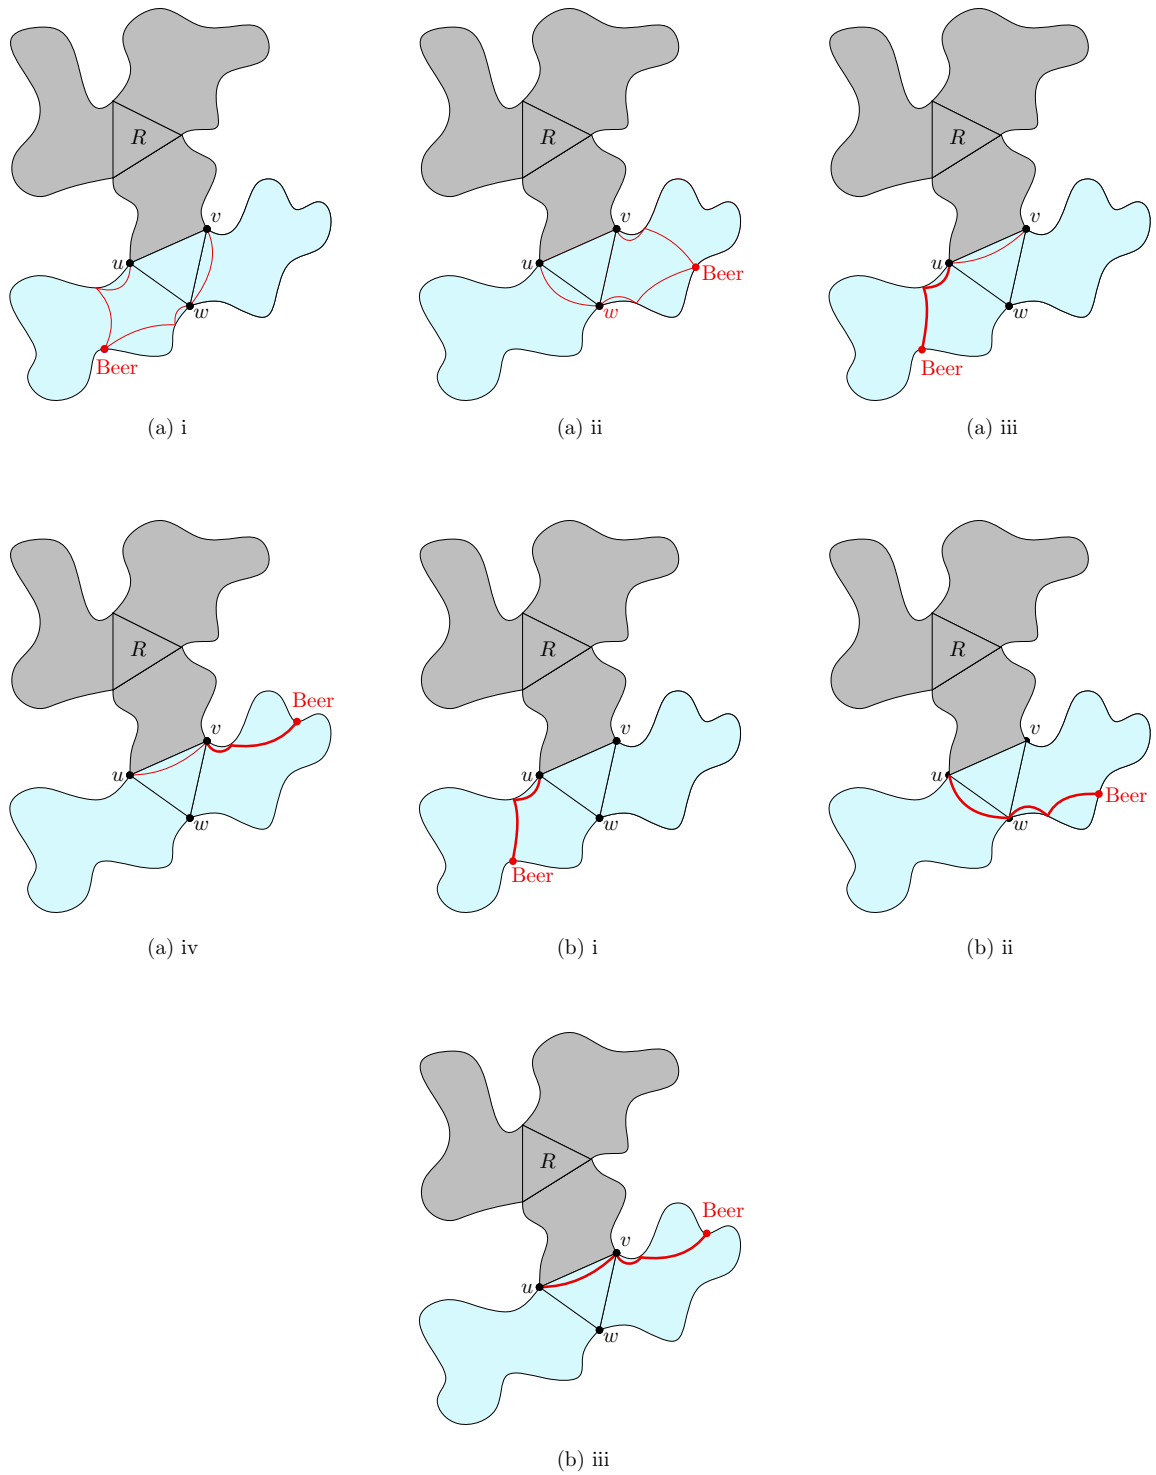


Figure 9: Illustrating the post-order traversal for all cases in item 2.

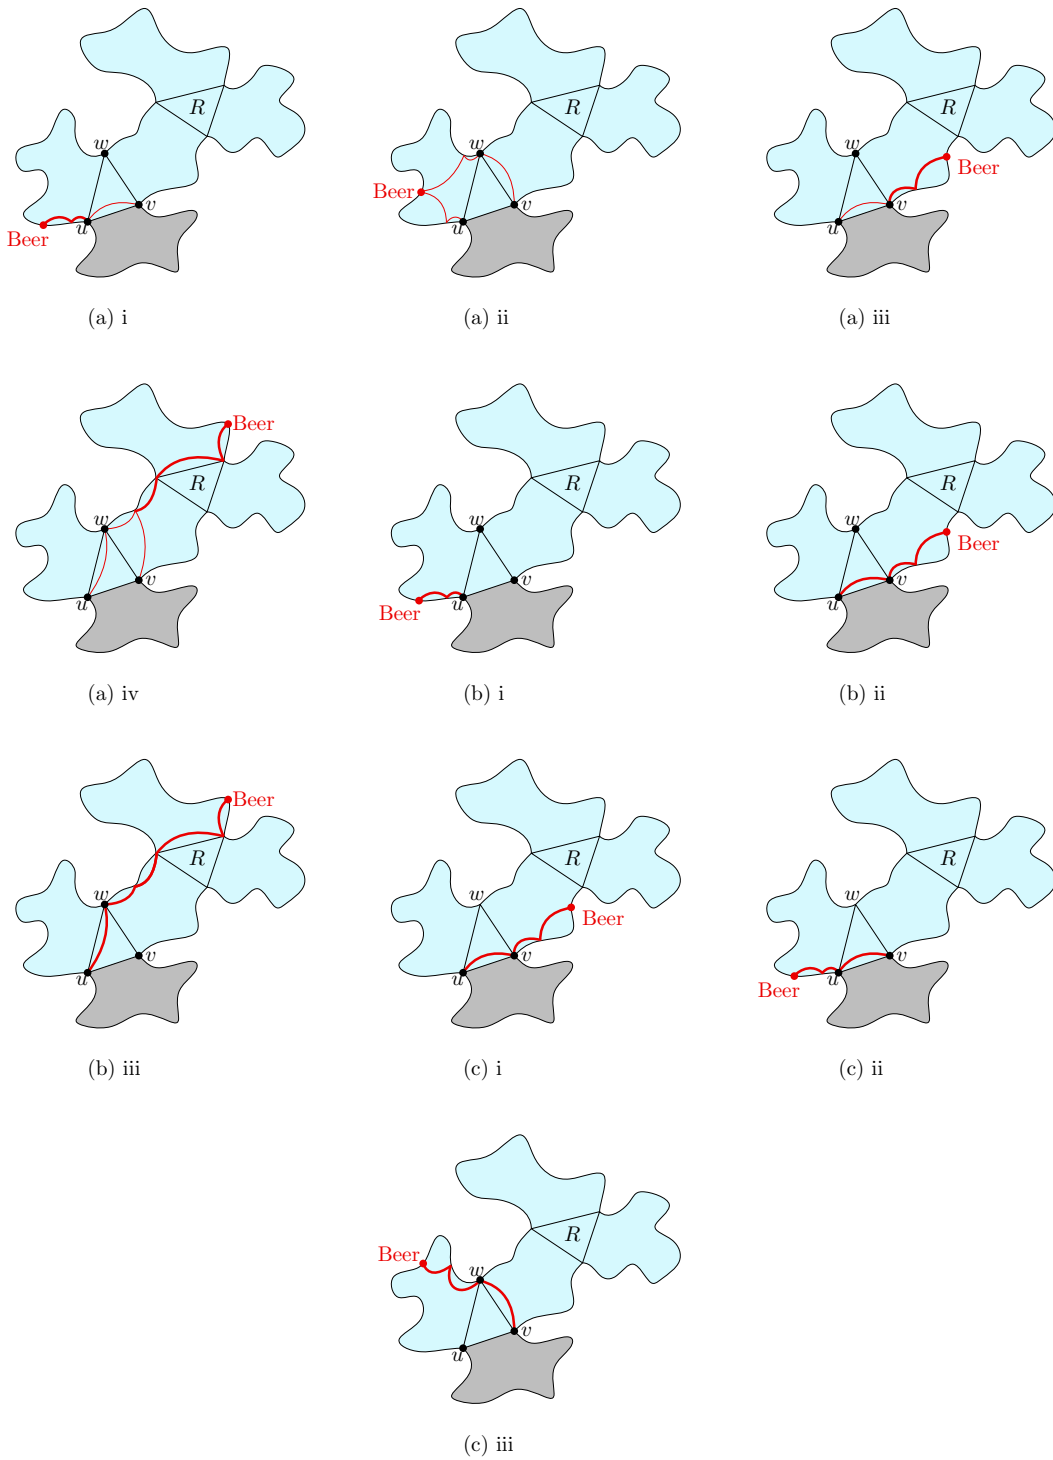


Figure 10: Illustrating the pre-order traversal for all cases in item 2.