Shortest paths in arbitrary plane domains

L. C. Hoehn, L. G. Oversteegen, and E. D. Tymchatyn

Abstract. Let Ω be a connected open set in the plane and $\gamma : [0,1] \to \overline{\Omega}$ a path such that $\gamma((0,1)) \subset \Omega$. We show that the path γ can be "pulled tight" to a unique shortest path which is homotopic to γ, via a homotopy h with endpoints fixed whose intermediate paths h_t, for $t \in [0,1)$, satisfy $h_t((0,1)) \subset \Omega$. We prove this result even in the case when there is no path of finite Euclidean length homotopic to γ under such a homotopy. For this purpose, we offer three other natural, equivalent notions of a "shortest" path. This work generalizes previous results for simply connected domains with simple closed curve boundaries.

1 Introduction

Bourgin and Renz [2] proved that given a simply connected plane domain (connected open set) Ω with simple closed curve boundary, and given any two points $p, q \in \overline{\Omega}$, there exists a unique *shortest* path in $\overline{\Omega}$ which is the uniform limit of paths which (except possibly for their endpoints p and q) are contained in Ω. If there is a rectifiable such curve (i.e. one with finite Euclidean length), then by shortest is meant the one with the smallest Euclidean length. If not, then by shortest is meant *locally shortest*, which means that every subpath not containing the endpoints p and q is of finite Euclidean length and is shortest among all such paths joining the endpoints of the subpath.

In Theorem 1, we extend the result of Bourgin and Renz to paths in multiply connected domains with arbitrary, in particular not necessarily locally connected, boundaries. Along the way, we characterize the concept of a shortest path in up to four different ways, outlined in the subsections below and stated in Theorem 2. In essence, one may imagine starting with a path representing a given homotopy class in the domain and pulling it tight by reeling it in at the endpoints. As we pull it tight, the path may come into contact with the boundary of the domain, but may not move across it. We remark that this is different from first taking the closure of the domain, and then considering shortest paths in that space. For example, if we consider a domain whose complement contains isolated points, such points vanish when one takes the closure of the domain, yet in the setting we consider here, such points are obstacles that the path will snag on as we pull it tight. For results about shortest paths in closed subsets of the plane, or in locally connected continua, see e.g. [3] and [8]. To state

our results precisely, we must fix some terminology and notation regarding paths and homotopies.

For notational convenience, we identify the plane $\mathbb{R}^2$ with the complex numbers $\mathbb{C}$. Fix a connected open set $\Omega \subset \mathbb{C}$ and points $p, q \in \overline{\Omega}$. We consider paths whose range is contained in Ω, except that the endpoints of the path may belong to $\partial\Omega$. For brevity, we use the abbreviation "e.p.e." to mean "except possibly at endpoints." Formally, a *path in Ω (e.p.e.) joining p and q* is a continuous function $\gamma : [0,1] \to \mathbb{C}$ such that $\gamma(0) = p$, $\gamma(1) = q$, and $\gamma(s) \in \Omega$ for all $s \in (0,1)$. When $p, q \in \partial\Omega$, the existence of such a path is equivalent to the statement that p and q are *accessible* from Ω.

We consider homotopies $h : [0,1] \times [0,1] \to \mathbb{C}$ such that for each $t \in [0,1)$, the path $h_t : [0,1] \to \mathbb{C}$ defined by $h_t(s) = h(s,t)$ is a path in Ω (e.p.e.) joining p and q. The path h_1 may meet $\partial\Omega$ in points other than $h_1(0)$ and $h_1(1)$. Given a path γ in Ω (e.p.e.) joining p and q, let $\overline{[\gamma]}$ denote the set of all paths homotopic to γ under such homotopies. Paths in $\overline{[\gamma]}$ may meet the boundary $\partial\Omega$ in more than just the endpoints, but for such a path, it must be possible to arrive at it via a homotopy of γ of the type described above. Let $[\gamma]$ denote the set of all paths in Ω (e.p.e.) which belong to $\overline{[\gamma]}$. Note that in spite of the notation, $\overline{[\gamma]}$ is not equal to the topological closure of $[\gamma]$ (e.g. in the function space).

The main result of this paper is that if γ is a path in Ω (e.p.e.) joining p and q, then $\overline{[\gamma]}$ contains a unique *shortest* path. By "shortest," we mean any one of the equivalent notions introduced next and collected in Theorem 2 below.

First, if there is a path in $\overline{[\gamma]}$ of finite Euclidean length, then we may consider the path in $\overline{[\gamma]}$ of smallest Euclidean length.

Second, we adapt Bourgin and Renz's condition of locally shortest as follows:

Definition   Let γ be a path in Ω (e.p.e.) joining p and q, and let $\lambda \in \overline{[\gamma]}$.

- Given $0 < s_1 < s_2 < 1$, we define the class $\overline{[\lambda^\gamma_{[s_1,s_2]}]}$ as follows. Let $h : [0,1] \times [0,1] \to \overline{\Omega}$ be a homotopy such that $h_0 = \gamma$, $h_1 = \lambda$, and for each $t \in [0,1)$, h_t is a path in Ω (e.p.e.) joining p and q. We use the homotopy h to "pull off" the path $\lambda\restriction_{[s_1,s_2]}$ (e.p.e.) from $\partial\Omega$ as follows. Define the path $\lambda^\gamma_{[s_1,s_2]} : [s_1, s_2] \to \overline{\Omega}$ by

$$\lambda^\gamma_{[s_1,s_2]}(s) = h(s, 1 - (s - s_1)(s_2 - s)).$$

  See (Figure 1). Though this path is defined in terms of the homotopy h, the class $\overline{[\lambda^\gamma_{[s_1,s_2]}]}$ does not depend on the choice of h (see Corollary 12 below).
- A path $\lambda \in \overline{[\gamma]}$ is called *locally shortest* if for any $0 < s_1 < s_2 < 1$, the path $\lambda\restriction_{[s_1,s_2]}$ has finite Euclidean length, and this length is smallest among all paths in $\overline{[\lambda^\gamma_{[s_1,s_2]}]}$.

Two further notions of "shortest path" are given in the next subsections.

1.1 Efficient paths

As above, fix a connected open set $\Omega \subset \mathbb{C}$ and points $p, q \in \overline{\Omega}$.
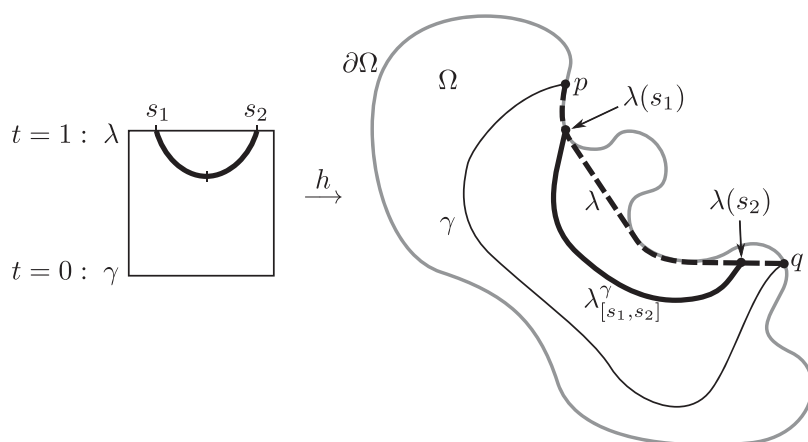
Figure 1: An illustration of the paths γ, λ, and $\lambda^{\gamma}_{[s_1,s_2]}$, for the definition of a locally shortest path.

Definition Let γ be a path in Ω (e.p.e.) joining p and q. A path $\lambda \in \overline{[\gamma]}$ is called *efficient* (in $\overline{[\gamma]}$) if the following property holds:

Given any $s_1, s_2 \in [0,1]$ with $s_1 < s_2$, let λ' be the path defined by $\lambda'(s) = \lambda(s)$ for $s \notin [s_1, s_2]$, and $\lambda'\!\restriction_{[s_1,s_2]}$ parameterizes the straight line segment $\overline{\lambda(s_1)\lambda(s_2)}$ or the constant path if $\lambda(s_1) = \lambda(s_2)$. If $\lambda' \in \overline{[\gamma]}$, then $\lambda = \lambda'$ (up to reparameterization).

The following is the first main result of this paper. The proof is in Section 3.

Theorem 1 *Let $\Omega \subset \mathbb{C}$ be a connected open set, let $p, q \in \overline{\Omega}$, and let γ be a path in Ω (except possibly at endpoints) joining p and q. Then $\overline{[\gamma]}$ contains a unique (up to parameterization) efficient path.*

1.2 Alternative notion of path length

Instead of the standard Euclidean path length, which can only distinguish between two paths if at least one of them is rectifiable, we can use an alternative notion of path length for which all paths have finite length, and which has other useful properties and many features in common with Euclidean length. A notion of length with such properties was first introduced in [11] and further developed in [15]. A similar notion is given in [3], and the authors of the present paper modified and extended that notion in [7].

Let len refer to either the path length function introduced in [7] or in [11]. The essential properties of len are:

(1) $\mathsf{len}(\gamma)$ is finite (in fact $\mathsf{len}(\gamma) < 1$) for any path γ;
(2) $\mathsf{len}(\gamma) = 0$ if and only if γ is constant;

(3) For any $p, q \in \mathbb{C}$, the straight line segment $\overline{pq}$ has smallest len length among all paths from p to q. Moreover, if γ is a path from p to q which deviates from the straight line segment $\overline{pq}$, or which is not monotone, then $\operatorname{len}(\gamma) > \operatorname{len}(\overline{pq})$;

(4) If $\Phi : \mathbb{C} \to \mathbb{C}$ is an isometry, then $\operatorname{len}(\Phi \circ \gamma) = \operatorname{len}(\gamma)$.

(5) Given $0 \le c_1 < c_2 \le 1$, $\operatorname{len}(\gamma\!\upharpoonright_{[c_1,c_2]}) \le \operatorname{len}(\gamma)$. This inequality is strict unless γ is constant outside of $[c_1, c_2]$;

(6) Given $c \in (0,1)$, $\operatorname{len}(\gamma) \le \operatorname{len}(\gamma\!\upharpoonright_{[0,c]}) + \operatorname{len}(\gamma\!\upharpoonright_{[c,1]})$. This inequality is strict unless one of the subpaths is constant; and

(7) len is a continuous function from the space of paths (with the uniform metric $d_{\sup}$) to $\mathbb{R}$.

Our second main result is the following. The proof is in Section 4.

Theorem 2 *Let $\Omega \subset \mathbb{C}$ be a connected open set, let $p, q \in \overline{\Omega}$, and let γ be a path in Ω (except possibly at endpoints) joining p and q. Then for a path $\lambda \in \overline{[\gamma]}$, the following are equivalent:*

(1) *λ is locally shortest;*

(2) *λ is efficient; and*

(3) *λ has smallest len length among all paths in $\overline{[\gamma]}$.*

Moreover, if $\overline{[\gamma]}$ contains a rectifiable path, then in addition to the above, we have

(4) *λ has smallest Euclidean length among all paths in $\overline{[\gamma]}$.*

2 Preliminaries on bounded analytic covering maps

Our arguments in Sections 3 and 4 make heavy use of the theory of complex analytic covering maps. In this section, we collect the results we will use later.

Denote $\mathbb{D} = \{z \in \mathbb{C} : |z| < 1\}$. It is a standard classical result (see e.g. [1]) that for any connected open set $\Omega \subset \mathbb{C}$ whose complement contains at least two points, and for any $z_0 \in \Omega$, there is a complex analytic covering map $\varphi : \mathbb{D} \to \Omega$ such that $\varphi(0) = z_0$. Moreover, this covering map φ is uniquely determined by the argument of $\varphi'(0)$.

Many of the results below hold only for analytic covering maps $\varphi : \mathbb{D} \to \Omega$ to bounded sets Ω. For the remainder of this subsection, let $\Omega \subset \mathbb{C}$ be a bounded connected open set, and let $\varphi : \mathbb{D} \to \Omega$ be an analytic covering map.

We first state three classic results about the boundary behavior of φ. All of the subsequent results in this section will be derived from these and from standard covering map theory.

Theorem 3 (Fatou [5], see e.g. [4, p. 22]) *The radial limits $\lim_{r \to 1^-} \varphi(r\alpha)$ exist for all points $\alpha \in \partial\mathbb{D}$ except possibly for a set of linear measure zero in $\partial\mathbb{D}$.*

It can easily be seen that if $\alpha \in \partial\mathbb{D}$ is such that $\lim_{r \to 1^-} \varphi(r\alpha)$ exists, then the limit belongs to $\partial\Omega$.

Theorem 4 (Riesz [13, 14], see e.g. [4, p. 22]) *For each $z \in \partial\Omega$, the set of points $\alpha \in \partial\mathbb{D}$ for which $\lim_{r \to 1^-} \varphi(r\alpha) = z$ has linear measure zero in $\partial\mathbb{D}$.*

Theorem 5 (Lindelöf [10], see e.g. [4, p. 23]) *Let $\widehat{\gamma} : [0,1] \to \overline{\mathbb{D}}$ be a path such that $\widehat{\gamma}([0,1)) \subset \mathbb{D}$ and $\widehat{\gamma}(1) = \alpha \in \partial\mathbb{D}$. Suppose that $\lim_{t\to 1^-} \varphi \circ \widehat{\gamma}(t)$ exists. Then the radial limit $\lim_{r\to 1^-} \varphi(r\alpha)$ exists and is equal to $\lim_{t\to 1^-} \varphi \circ \widehat{\gamma}(t)$.*

We define the *extended covering map*

$$\overline{\varphi} : \mathbb{D} \cup \left\{ \alpha \in \partial\mathbb{D} : \lim_{r\to 1^-} \varphi(r\alpha) \text{ exists} \right\} \to \overline{\Omega}$$

by

$$\overline{\varphi}(\widehat{z}) = \begin{cases} \varphi(\widehat{z}) & \text{if } \widehat{z} \in \mathbb{D} \\ \lim_{r\to 1^-} \varphi(r\alpha) & \text{if } \widehat{z} = \alpha \in \partial\mathbb{D}. \end{cases}$$

By Theorem 3, this function is defined on $\mathbb{D}$ plus a full measure subset of $\partial\mathbb{D}$. Note that this function $\overline{\varphi}$ is not necessarily continuous at points where it is defined in $\partial\mathbb{D}$. It is, however, continuous by definition along each radial segment from the center 0 of $\mathbb{D}$ to any point $\alpha \in \partial\mathbb{D}$ where it is defined. In fact, more is true: if $\overline{\varphi}$ is defined at α, then its restriction to any *Stolz angle* at α is continuous (see e.g. [4, p. 23]); however, we will not need this concept in this paper.

As a general convention, we will put hats on symbols, as in $\widehat{z}$, $\widehat{A}$, or $\widehat{\gamma}$, to refer to points, subsets, or paths in $\overline{\mathbb{D}}$, and use symbols without hats to refer to points, subsets, or paths in $\overline{\Omega} \subset \mathbb{C}$, with the understanding that if both z and $\widehat{z}$ appear in an argument, they are related by $z = \overline{\varphi}(\widehat{z})$ (and likewise for sets and paths). If $z = \overline{\varphi}(\widehat{z})$ for points $z \in \overline{\Omega}$ and $\widehat{z} \in \overline{\mathbb{D}}$ (respectively, $A = \overline{\varphi}(\widehat{A})$ for sets $A \subset \overline{\Omega}$ and $\widehat{A} \subset \overline{\mathbb{D}}$), we say $\widehat{z}$ is a *lift* of z (respectively, $\widehat{A}$ is a *lift* of A). Similarly, if $\gamma = \overline{\varphi} \circ \widehat{\gamma}$ for paths γ in $\overline{\Omega}$ and $\widehat{\gamma}$ in $\overline{\mathbb{D}}$, we say $\widehat{\gamma}$ is a *lift* of γ.

The next result about lifts of paths is very similar to classical results for covering maps. Since our extended map $\overline{\varphi}$ is not necessarily continuous at points in $\partial\mathbb{D}$, it is not a simple consequence of basic covering map theory. It is proved in [6], and we also include a proof here to convey the flavor of working with the extended map $\overline{\varphi}$.

Theorem 6 *Let $p, q \in \overline{\Omega}$, let $z \in \Omega$, and let $\widehat{z} \in \mathbb{D}$ such that $\varphi(\widehat{z}) = z$.*

(1) *Let γ be a path in Ω (e.p.e.) joining p and q, and suppose $\gamma(s_0) = z$ for some $s_0 \in [0,1]$. Then there exists a unique path $\widehat{\gamma}$ in $\mathbb{D}$ (e.p.e.) such that $\overline{\varphi} \circ \widehat{\gamma} = \gamma$ and $\widehat{\gamma}(s_0) = \widehat{z}$.*

(2) *Let $h : [0,1] \times [0,1] \to \mathbb{C}$ be a homotopy such that for each $t \in [0,1]$, the path $h_t(s) = h(s,t)$ is a path in Ω (e.p.e.) joining p and q, and suppose $h(s_0, 0) = z$ for some $s_0 \in [0,1]$. Then there exists a unique homotopy $\widehat{h} : [0,1] \times [0,1] \to \overline{\mathbb{D}}$ such that $\widehat{h}(s,t) \in \mathbb{D}$ and $\varphi \circ \widehat{h}(s,t) = h(s,t)$ for each $(s,t) \in (0,1) \times [0,1]$, $\widehat{h}(\{0\} \times [0,1])$ is a single point $\widehat{p}$ in $\overline{\mathbb{D}}$ with $\overline{\varphi}(\widehat{p}) = p$, $\widehat{h}(\{1\} \times [0,1])$ is a single point $\widehat{q}$ in $\overline{\mathbb{D}}$ with $\overline{\varphi}(\widehat{q}) = q$, and $\widehat{h}(s_0, 0) = \widehat{z}$.*

Proof Observe that (1) follows from (2), by using the constant homotopy $h(s,t) = \gamma(s)$ for all $t \in [0,1]$. Thus, it suffices to prove (2).

Since φ is a covering map and $h((0,1) \times [0,1]) \subset \Omega$, it follows from standard covering space theory that there is a unique homotopy $\widehat{h} : (0,1) \times [0,1] \to \mathbb{D}$ such that $\varphi \circ \widehat{h}(s,t) = h(s,t)$ for each $(s,t) \in (0,1) \times [0,1]$, and $\widehat{h}(s_0, 0) = \widehat{z}$. It remains to prove that there are points $\widehat{p}, \widehat{q} \in \mathbb{D}$ such that defining $\widehat{h}(\{0\} \times [0,1]) = \{\widehat{p}\}$ and $\widehat{h}(\{1\} \times [0,1]) = \{\widehat{q}\}$ makes $\widehat{h}$ into a continuous function from $[0,1] \times [0,1]$ to $\overline{\mathbb{D}}$. This is immediate if p and q belong to Ω, so we assume $p, q \in \partial\Omega$.

Observe that the set $\widehat{h}((0, \frac{1}{2}] \times [0,1])$ compactifies on a continuum $K \subset \partial\mathbb{D}$. We need to prove that K is a single point $\{\widehat{p}\}$. Suppose for a contradication that K contains more than one point. Then there exists by Theorem 3 a set E of positive measure in the interior of K so that for each $\alpha \in E$, the radial limit $\lim_{r \to 1^-} \varphi(r\alpha)$ exists. Since the set $\widehat{h}((0, \frac{1}{2}] \times [0,1])$ compactifies on K, we can choose, for each $\alpha \in E$, a sequence (s_n, t_n) in $(0, \frac{1}{2}] \times [0,1]$ such that $s_n \to 0$ and $\widehat{h}(s_n, t_n) = r_n \alpha$, with $r_n \to 1$. It follows that the radial limit $\lim_{r \to 1^-} \varphi(r\alpha) = p$ for each $\alpha \in E$, a contradiction with Theorem 4. Thus K is a single point $\{\widehat{p}\}$, and so we can continuously extend $\widehat{h}$ to $\{0\} \times [0,1]$ by defining $\widehat{h}(\{0\} \times [0,1]) = \{\widehat{p}\}$. By Theorem 5 it follows that $\overline{\varphi}(\widehat{p}) = p$.

Likewise, by considering $\widehat{h}([\frac{1}{2}, 1) \times [0,1])$, we obtain by the same argument a point $\widehat{q} \in \overline{\mathbb{D}}$ such that $\widehat{h}$ extends continuously to $\{1\} \times [0,1]$ by defining $\widehat{h}(\{1\} \times [0,1]) = \{\widehat{q}\}$, and $\overline{\varphi}(\widehat{q}) = q$. ∎

Theorem 6(2) implies that if $p, q \in \overline{\Omega}$, γ is a path in Ω (e.p.e.) joining p and q, $\widehat{\gamma}$ is a lift of γ with endpoints $\widehat{p}, \widehat{q} \in \overline{\mathbb{D}}$, and $\lambda \in [\gamma]$, then there exists a lift $\widehat{\lambda}$ of λ with the same endpoints $\widehat{p}, \widehat{q}$ as $\widehat{\gamma}$. We will prove a stronger statement in Theorem 11 below.

The next result follows immediately from Theorem 6(1).

Corollary 7   *Let L be an open arc in Ω whose closure is an arc with distinct endpoints in $\partial\Omega$. Let $\widehat{L}$ be a component of $\varphi^{-1}(L)$. Then $\widehat{L}$ is an open arc in $\mathbb{D}$ whose closure in $\overline{\mathbb{D}}$ is an arc with distinct endpoints in $\partial\mathbb{D}$.*

We next prove a result about the existence of small "crosscuts" in $\mathbb{D}$ straddling any point $\alpha \in \partial\mathbb{D}$ for which $\overline{\varphi}(\alpha) \in \partial\Omega$ is not isolated in $\partial\Omega$. For one-to-one analytic maps $\varphi : \mathbb{D} \to \mathbb{C}$, this result is standard; see e.g. [12]. We were unable to find a reference for the case of a bounded analytic covering map φ, so we include a proof for completeness.

Theorem 8   *Let $\alpha \in \partial\mathbb{D}$ be such that the radial limit $\lim_{r \to 1^-} \varphi(r\alpha)$ exists, and let $p \in \partial\Omega$ be the limit.*

(1) *If p is not isolated in $\partial\Omega$, then for any sufficiently small simple closed curve S in $\mathbb{C}$ containing p in its interior, there is a component $\widehat{S}$ of $\varphi^{-1}(S)$ whose closure is an arc separating α from the center 0 of $\mathbb{D}$ in $\overline{\mathbb{D}}$.*

(2) *If p is isolated in $\partial\Omega$, then for any sufficiently small simple closed curve S in Ω containing p in its interior, there is a component $\widehat{S}$ of $\varphi^{-1}(S)$ whose closure is a circle whose intersection with $\partial\mathbb{D}$ is $\{\alpha\}$.*

Moreover, in both cases, the diameter of $\widehat{S}$ can be made arbitrarily small by choosing S sufficiently small.

Proof   Let S be small enough so that $\varphi(0)$ is not in the closed topological disk bounded by S. Let $r_0 < 1$ be close enough to 1 so that $\varphi(r\alpha)$ is in the interior of S for all $r \in [r_0, 1)$. The closed (in $\mathbb{D}$) set $\varphi^{-1}(S)$ must separate $r_0\alpha$ from 0 in $\mathbb{D}$, since otherwise there would be a path from 0 to $r_0\alpha$ which would project to a path in Ω from $\varphi(0)$ to $\varphi(r_0\alpha)$ without intersecting S, a contradiction. Therefore, there must be a component $\widehat{S}$ of $\varphi^{-1}(S)$ which separates $r_0\alpha$ from 0 in $\mathbb{D}$ (see e.g. [9, p. 438, Section 57 III, Theorem 1]).

For (1), suppose that p is not isolated in $\partial\Omega$. Assume first that $S \cap \partial\Omega \neq \varnothing$. Let C be the component of $S \cap \Omega$ containing $\varphi(\widehat{S})$. The closure of this C is a path in Ω (e.p.e.) joining two points (not necessarily distinct) $a, b \in \partial\Omega$ with $a \neq p \neq b$. By Theorem 6 and the fact that C is an open arc in Ω, we have that the closure of each component of $\varphi^{-1}(C)$ is an arc in $\mathbb{D}$ joining points $\widehat{a}, \widehat{b} \in \partial\mathbb{D}$ at which the radial limits exist and are equal to a and b, respectively; in particular, we have $\widehat{a} \neq \alpha \neq \widehat{b}$. It follows that the closure of $\widehat{S}$ is an arc with endpoints distinct from α, which separates α from 0 in $\overline{\mathbb{D}}$, as desired.

Now choose $\varepsilon_1 > 0$ small enough so that $\varphi(0)$ is not in the closed disk $\overline{B}(p, \varepsilon_1)$ and such that $\partial B(p, \varepsilon_1) \cap \partial\Omega \neq \varnothing$. Assume that $S \subset B(p, \varepsilon_1)$. If $S \cap \partial\Omega \neq \varnothing$, then we are done by the previous paragraph; hence, suppose that $S \cap \partial\Omega = \varnothing$. Choose $\varepsilon_0 > 0$ small enough so that $B(p, \varepsilon_0)$ is contained in the interior of S, and such that $\partial B(p, \varepsilon_0) \cap \partial\Omega \neq \varnothing$. By the previous paragraph, there are components $\widehat{S}_0$ and $\widehat{S}_1$ of $\varphi^{-1}(\partial B(p, \varepsilon_0))$ and $\varphi^{-1}(\partial B(p, \varepsilon_1))$, respectively, which are arcs separating α from 0. As above, there is a component $\widehat{S}$ of $\varphi^{-1}(S)$ which separates a tail end of the radial segment at α from 0 in $\mathbb{D}$, and which separates $\widehat{S}_0$ from $\widehat{S}_1$ in $\mathbb{D}$ (i.e. lies between $\widehat{S}_0$ and $\widehat{S}_1$). This implies the endpoints of the closure of $\widehat{S}$ are distinct, hence the closure of $\widehat{S}$ is an arc, as desired.

For (2), let S be small enough so that p is the only point of $\partial\Omega$ in the closed topological disk bounded by S. As above, we obtain a component $\widehat{S}$ of $\varphi^{-1}(S)$ which separates a tail end of the radial segment at α from 0 in $\mathbb{D}$. Since φ is a covering map, this $\widehat{S}$ must be an open arc in $\mathbb{D}$ whose two ends both accumulate on continua K_1 and K_2 in $\partial\mathbb{D}$.

We first argue that K_1 and K_2 are single points. Suppose $\beta_1, \beta_2 \in K_1$ with $\beta_1 \neq \beta_2$. Then by Theorem 3, there exists $\beta \in K_1$ between β_1 and β_2 where the radial limit $\lim_{r\to 1^-} \varphi(r\beta)$ exists. However, this radial segment meets $\widehat{S}$ arbitrarily close to β, hence $\lim_{r\to 1^-} \varphi(r\beta) \notin \partial\Omega$, a contradiction. Thus, K_1 is a single point $\{\widehat{a}\}$. Similarly, K_2 is a single point $\{\widehat{b}\}$.

If $\widehat{a} \neq \alpha$, by Theorem 3 and Theorem 4, we can find $\beta \in \partial\mathbb{D}$ between $\widehat{a}$ and α such that the radial limit $\lim_{r\to 1^-} \varphi(r\beta)$ exists and is different from p. The path $\overline{\varphi}(r\beta)$, $0 \leq r \leq 1$, is homotopic to one which does not enter the interior of S. By Theorem 6, we can lift this homotopy, to obtain a path from 0 to β in $\mathbb{D}$ which does not meet $\widehat{S}$. But this is a contradiction since $\widehat{S}$ separates β from 0. Therefore, $\widehat{a} = \alpha$, and likewise $\widehat{b} = \alpha$. Thus the closure of $\widehat{S}$ is a circle meeting $\partial\mathbb{D}$ at α only, as desired.

For the moreover part, we argue as in the proof of Theorem 6 that if these components $\widehat{S}$ did not converge to 0 in diameter as the diameter of S is shrunk toward 0, then they would accumulate on a nondegenerate continuum $K \subset \partial\mathbb{D}$, and we would obtain a contradiction by Theorem 3 and Theorem 4. The details are left to the reader. ∎

Lemma 9 *Let S be a straight line or round circle in $\mathbb{C}$, and let $\varepsilon > 0$. Then there are only finitely many lifts of components of $S \cap \Omega$ with diameter at least ε.*

Proof   Suppose the claim is false, so that there exists $\varepsilon > 0$ and infinitely many lifts $\widehat{S}_n$, $n = 1, 2, \ldots$, of components of $S \cap \Omega$ such that the diameter of $\widehat{S}_n$ is at least ε for each n. These lifts accumulate on a non-degenerate continuum $K \subset \overline{\mathbb{D}}$.

If $K \cap \mathbb{D} \neq \varnothing$, then let $\widehat{z} \in K \cap \mathbb{D}$ and let $\widehat{V}$ be a neighborhood of $\widehat{z}$ which maps one-to-one under φ to a small round disk $V \subset \Omega$. Then $\widehat{V}$ meets infinitely many of the lifts $\widehat{S}_n$. On the other hand, because V is a round disk contained in Ω and S is a round circle or straight line, it follows that V can only meet one component of $S \cap \Omega$. This is a contradiction since φ is one-to-one on $\widehat{V}$.

Suppose then that $K \subset \partial\mathbb{D}$. Then there exists by Theorem 3 a set E of positive measure in the interior of K so that for each $\alpha \in E$, the radial limit $\lim_{r \to 1^-} \varphi(r\alpha)$ exists.

If there is a single component S' of $S \cap \Omega$ such that $\varphi(\widehat{S}_n) = S'$ for infinitely many n, then it is clear that the radial limit of φ at each $\alpha \in E$ must belong to $\overline{S'} \cap \partial\Omega$, which contains at most two points. But this contradicts Theorem 4.

Therefore, we may assume that the components $\varphi(\widehat{S}_n)$ are all distinct, which means their diameters must converge to 0. By passing to a subsequence if necessary, we may assume that the components $\varphi(\widehat{S}_n)$ converge to a single point $a \in S \cap \partial\Omega$. Then it is clear that the radial limit of φ at each $\alpha \in E$ must equal a, again a contradiction by Theorem 4.   ∎

Recall that the function $\overline{\varphi}$ is not necessarily continuous at points $\alpha \in \partial\mathbb{D}$ where it is defined. However, the next result shows that the restriction of $\overline{\varphi}$ to the region in between two lifted paths with the same endpoints is continuous.

Given a continuum X in $\mathbb{C}$, the *topological hull* of X, denoted $\mathrm{Hull}(X)$, is the smallest simply connected continuum in $\mathbb{C}$ containing X. Equivalently, $\mathrm{Hull}(X)$ is equal to $\mathbb{C} \setminus U$, where U is the unbounded component of $\mathbb{C} \setminus X$.

Lemma 10 *Suppose $\widehat{\lambda} : [0,1] \to \overline{\mathbb{D}}$ is a path such that $\overline{\varphi} \circ \widehat{\lambda}$ is a path in $\overline{\Omega}$ (i.e. is continuous). Let X be the union of $\widehat{\lambda}([0,1])$ with the two radial segments from the center 0 of $\mathbb{D}$ to $\widehat{p} = \widehat{\lambda}(0)$ and to $\widehat{q} = \widehat{\lambda}(1)$, and let $\Delta = \mathrm{Hull}(X)$. Then Δ is simply connected and locally connected, and the restriction $\overline{\varphi}\restriction_\Delta$ is continuous on Δ.*

Proof   Note that Δ is simply connected by definition, and it is straightforward to see that the topological hull of any locally connected continuum is locally connected. Since φ is continuous and $\overline{\varphi} = \varphi$ in $\mathbb{D}$, it remains to prove that the restriction of $\overline{\varphi}$ to Δ is continuous at each point of $\Delta \cap \partial\mathbb{D}$.

Let $\alpha \in \Delta \cap \partial\mathbb{D}$, and suppose for a contradiction that the restriction of $\overline{\varphi}$ to Δ is not continuous at α. Then there exists $\varepsilon > 0$ and a sequence of points $\langle \widehat{w}_n \rangle_{n=1}^\infty$ in Δ such that $\widehat{w}_n \to \alpha$, but $|\overline{\varphi}(\alpha) - \varphi(\widehat{w}_n)| \geq \varepsilon$ for all n.

Note that since $\Delta \cap \partial\mathbb{D} \subset \widehat{\lambda}([0,1])$ and $\overline{\varphi} \circ \widehat{\lambda}$ is continuous, we have that the restriction of $\overline{\varphi}$ to $\Delta \cap \partial\mathbb{D}$ is continuous. Hence, we may assume that $\widehat{w}_n \in \Delta \cap \mathbb{D}$ for each n. Moreover, since the restriction of $\overline{\varphi}$ to the radial segment from 0 to α is continuous, we may also assume that $\widehat{w}_n$ does not belong to this segment for all n.

For each n, there is a lift $\widehat{C}_n$ of a component of $\Omega \cap \partial B(\overline{\varphi}(\alpha), \varepsilon)$ which separates $\widehat{w}_n$ from α in $\overline{\mathbb{D}}$, where $B(\overline{\varphi}(\alpha), \varepsilon)$ is the open disk centered at $\overline{\varphi}(\alpha)$ of radius ε. By Theorem 8, this lift $\widehat{C}_n$ is either an arc with endpoints in $\partial\mathbb{D}$, or, in the case that $\overline{\varphi}(\alpha)$ is the only point of $\partial\Omega$ in the closed disk $\overline{B}(\overline{\varphi}(\alpha), \varepsilon)$, a circle with one point on $\partial\mathbb{D}$. By passing to a subsequence if necessary, we may assume that all of the $\widehat{C}_n$ are distinct, hence they are pairwise disjoint. According to Lemma 9, the diameters of the lifts $\widehat{C}_n$ converge to 0 as $n \to \infty$. Let X be as in the statement of the lemma, and for each n choose a point $\widehat{z}_n \in X \cap \widehat{C}_n$. Then $\widehat{z}_n \to \alpha$.

Suppose first that there is a subsequence $\widehat{z}_{n_k}$ of $\widehat{z}_n$ such that for each k, $\widehat{z}_{n_k}$ belongs to the radial segment from 0 to $\widehat{p}$. It follows that $\alpha = \widehat{p}$. Since $\overline{\varphi}$ is continuous on this radial segment, we have $\varphi(\widehat{z}_{n_k}) \to \overline{\varphi}(\alpha)$ as $k \to \infty$. But $\varphi(\widehat{z}_n) \in \partial B(\overline{\varphi}(\alpha), \varepsilon)$ for each n, so this is a contradiction. Likewise, we encounter a contradiction if infinitely many of the points $\widehat{z}_n$ belong to the radial segment from 0 to $\widehat{q}$.

Thus we may assume that all of the points $\widehat{z}_n$ belong to the set $\widehat{\lambda}([0,1])$. Let $s_n \in [0,1]$ such that $\widehat{z}_n = \widehat{\lambda}(s_n)$. By passing to a subsequence if necessary, we may assume that $s_n \to s_\infty$, and $\widehat{\lambda}(s_\infty) = \alpha$. Since $\overline{\varphi} \circ \widehat{\lambda}$ is continuous, it follows that $\overline{\varphi}(\widehat{z}_n) \to \overline{\varphi}(\alpha)$ as $n \to \infty$. But again this is a contradiction because $\varphi(\widehat{z}_n) \in \partial B(\overline{\varphi}(\alpha), \varepsilon)$ for each n. ∎

In the next result, we characterize paths in $\overline{[\gamma]}$ in terms of lifts.

Theorem 11 *Let γ be a path in Ω (e.p.e.) joining p and q, and let $\widehat{\gamma}$ be a lift of γ with endpoints $\widehat{p}, \widehat{q} \in \overline{\mathbb{D}}$. If $\lambda \in \overline{[\gamma]}$, then there exists a lift $\widehat{\lambda}$ of λ (to $\overline{\mathbb{D}}$) with the same endpoints $\widehat{p}, \widehat{q}$. Conversely, if $\lambda : [0,1] \to \overline{\Omega}$ is a path joining p and q which has a lift $\widehat{\lambda} : [0,1] \to \overline{\mathbb{D}}$ with the same endpoints $\widehat{p}, \widehat{q}$, then $\lambda \in \overline{[\gamma]}$.*

Proof   Let $\lambda \in \overline{[\gamma]}$, and let h be a homotopy such that $h_0 = \gamma$, $h_1 = \lambda$, and h_t is a path in Ω (e.p.e.) joining p and q for each $t \in [0,1)$. For each $s \in [0,1]$, consider the path $t \mapsto h_t(s)$. Apply Theorem 6 to obtain a lift $\widehat{h}_t(s)$ such that $\widehat{h}_0(s) = \widehat{\gamma}(s)$. Define $\widehat{\lambda} : [0,1] \to \overline{\mathbb{D}}$ by $\widehat{\lambda}(s) = \widehat{h}_1(s)$. By Theorem 5, we have $\overline{\varphi} \circ \widehat{\lambda}(s) = \lambda(s)$ for all $s \in [0,1]$, and $\widehat{\lambda}(0) = \widehat{p}$ and $\widehat{\lambda}(1) = \widehat{q}$.

It remains to prove that $\widehat{\lambda}$ is continuous. Let $s \in [0,1]$. If $\widehat{\lambda}(s) \in \mathbb{D}$, then $\widehat{\lambda}$ is continuous at s by standard covering space theory. Suppose for a contradiction that $\widehat{\lambda}(s) \in \partial\mathbb{D}$ and $\widehat{\lambda}$ is not continuous at s. We proceed with an argument similar to the one given for Theorem 6. The sets $\widehat{h}([s - \frac{1}{n}, s + \frac{1}{n}] \times [1 - \frac{1}{n}, 1))$, $n = 1, 2, \ldots$, accumulate on a nondegenerate continuum $K \subset \partial\mathbb{D}$, which contains, by Theorem 3, a set E of positive measure such that the radial limit $\lim_{r \to 1^-} \varphi(r\alpha)$ exists for each $\alpha \in E$. We can choose, for each $\alpha \in E$, a sequence (s_n, t_n) converging to $(s, 1)$ such that $\widehat{h}(s_n, t_n) = r_n\alpha$, with $r_n \to 1$. It follows that the radial limit $\lim_{r \to 1^-} \varphi(r\alpha)$ is equal to

$$\lim_{n \to \infty} \varphi(r_n\alpha) = \lim_{n \to \infty} \varphi \circ \widehat{h}(s_n, t_n) = \lim_{n \to \infty} h(s_n, t_n) = \lambda(s),$$

a contradiction with Theorem 4. Therefore, $\widehat{\lambda}$ is a continuous lift of λ.

Conversely, suppose $\lambda : [0,1] \to \overline{\Omega}$ is a path joining p and q which has a lift $\widehat{\lambda} : [0,1] \to \overline{\mathbb{D}}$ with the same endpoints $\widehat{p}, \widehat{q}$ as $\widehat{\gamma}$. Let $\widehat{c}$ be the path in $\mathbb{D}$ (e.p.e.) such

that $\widehat{c}(0) = \widehat{p}$, $\widehat{c}(\frac{1}{2}) = 0$, $\widehat{c}(1) = \widehat{q}$, and $\widehat{c}$ linearly parameterizes the straight segments in between these points. Let $c = \overline{\varphi} \circ \widehat{c}$.

Let $X = \widehat{\lambda}([0,1]) \cup \widehat{c}([0,1])$, and let $\Delta = \mathrm{Hull}(X)$. Since Δ is simply connected, it follows that there is a homotopy $\widehat{h}$ between $\widehat{c}$ and $\widehat{\lambda}$ within Δ, such that $\widehat{h}_0 = \widehat{c}$, $\widehat{h}_1 = \widehat{\lambda}$, and $\widehat{h}_t$ is a path in $\mathbb{D}$ (e.p.e.) joining $\widehat{p}$ and $\widehat{q}$ for all $t \in [0,1)$. Since $\overline{\varphi}$ is continuous on Δ by Theorem 10, the composition $\overline{\varphi} \circ \widehat{h}$ is a homotopy between c and λ which establishes that $\lambda \in \overline{[c]}$. By the same reasoning, we can show that $\gamma \in [c]$. Therefore, $\lambda \in \overline{[\gamma]}$.   ∎

We now have the machinery in place to conclude that the class $\overline{[\lambda^{\gamma}_{[s_1,s_2]}]}$ described in the definition of a locally shortest path is well-defined.

Corollary 12   *Let γ be a path in Ω (e.p.e.) joining p and q, and let $\widehat{\gamma}$ be a lift of γ with endpoints $\widehat{p}, \widehat{q} \in \overline{\mathbb{D}}$. Let $\lambda \in \overline{[\gamma]}$, and let $\widehat{\lambda}$ be a lift of λ (to $\overline{\mathbb{D}}$) with the same endpoints $\widehat{p}, \widehat{q}$. Let $0 < s_1 < s_2 < 1$. A path ρ belongs to $\overline{[\lambda^{\gamma}_{[s_1,s_2]}]}$ if and only if there is a lift $\widehat{\rho}$ of ρ (to $\overline{\mathbb{D}}$) with endpoints $\widehat{\lambda}(s_1), \widehat{\lambda}(s_2)$.*

In particular, the definition of the class $\overline{[\lambda^{\gamma}_{[s_1,s_2]}]}$ is independent of the choice of homotopy h between γ and λ.

3   Proof of Theorem 1

Let $\Omega \subset \mathbb{C}$ be a connected open set, $p, q \in \overline{\Omega}$, and let γ be a path in Ω (e.p.e.) joining p and q. We may assume that either $p \neq q$ or γ is a non trivial loop, so that $\overline{[\gamma]}$ contains no constant path, since otherwise this is obviously the unique efficient path.

Let D_Ω be a large round disk in $\mathbb{C}$ which contains the entire path $\gamma([0,1])$. We may assume, without loss of generality, that Ω is contained in D_Ω, hence in particular is a bounded subset of $\mathbb{C}$. Indeed, any efficient path in $\overline{[\gamma]}$ with respect to Ω also belongs to $\overline{[\gamma]}$ with respect to $\Omega \cap D_\Omega$ and vice versa. To see this, suppose $h: [0,1] \times [0,1] \to \Omega$ is a homotopy between two paths, γ_1 and γ_2, in $\Omega \cap D_\Omega$. Consider another disk D'_Ω which is large enough to contain $\overline{D}_\Omega \cup h([0,1] \times [0,1])$, let $\varphi': \mathbb{D} \to \Omega \cap D'_\Omega$ be an analytic covering map onto the component of $\Omega \cap D'_\Omega$ containing γ_1 and γ_2, and lift h to a homotopy $\widehat{h}$ in $\mathbb{D}$ of paths joining $\widehat{p}$ and $\widehat{q}$. Clearly no lift of a component C of $D'_\Omega \cap \partial D_\Omega$ separates $\widehat{p}$ from $\widehat{q}$ is $\mathbb{D}$. It follows that one can retract the homotopy $\widehat{h}$ so that it does not intersect any of these lifts. Composing the retracted homotopy with φ' yields a homotopy from γ_1 to γ_2 in $\Omega \cap D_\Omega$. The same goes for the other notions of shortest path used in this paper. Hence, we assume $\Omega \subset D_\Omega$ for the remainder of this paper.

Let $\varphi: \mathbb{D} \to \Omega$ be an analytic covering map, and let $\overline{\varphi}$ denote the extension of φ to those points in $\partial \mathbb{D}$ where the radial limit is defined, as in Section 2. Choose any lift $\widehat{\gamma}: [0,1] \to \overline{\mathbb{D}}$ of γ under $\overline{\varphi}$ (see Theorem 6(1)). So $\widehat{\gamma}$ is a path in $\mathbb{D}$ (e.p.e.) and $\overline{\varphi} \circ \widehat{\gamma} = \gamma$. Let $\widehat{p} = \widehat{\gamma}(0) \in \overline{\mathbb{D}}$ and $\widehat{q} = \widehat{\gamma}(1) \in \overline{\mathbb{D}}$. Since $\overline{[\gamma]}$ does not contain a constant path, we have $\widehat{p} \neq \widehat{q}$ (cf. Theorem 11).

In Section 3.2 and Section 3.3 below, we will establish the existence of an efficient path in $\overline{[\gamma]}$ via a recursive construction in which we repeatedly replace subpaths by

straight line segments, when doing so does not change the homotopy class. We begin with some preliminary results in Section 3.1.

3.1 Lifts of lines

Throughout this paper, when we use the word *line*, we mean straight line in $\mathbb{C}$. In this subsection, we consider lines L in $\mathbb{C}$ which intersect Ω and lifts of closures of components of $L \cap \Omega$ under $\overline{\varphi}$. By abuse of terminology, any such lift will be called a *lift of L*. For a given line L, $L \cap \Omega$ has at most countably many components, and each of these components has at most countably many lifts, each of which is, by Corollary 7, an arc in $\overline{\mathbb{D}}$ whose (distinct) endpoints are in $\partial \mathbb{D}$, and which is otherwise contained in $\mathbb{D}$. Observe that if $\widehat{L}_1$ and $\widehat{L}_2$ are distinct lifts of lines, then $\widehat{L}_1 \cap \widehat{L}_2$ contains at most one point; moreover, if $\widehat{L}_1 \cap \widehat{L}_2 = \{\widehat{z}\}$ for some $\widehat{z} \in \mathbb{D}$, then $\widehat{L}_1$ and $\widehat{L}_2$ cross transversally at $\widehat{z}$.

Our construction later in this section of an efficient path is based on the following reformulation of the definition of an efficient path.

Proposition 13 *Let $\lambda \in \overline{[\gamma]}$, and let $\widehat{\lambda}$ be a lift of λ (under $\overline{\varphi}$) joining $\widehat{p}$ to $\widehat{q}$. Then λ is an efficient path in $\overline{[\gamma]}$ if and only if for any lift $\widehat{L}$ of a line intersecting Ω, the set $\widehat{\lambda}^{-1}(\widehat{L})$ is connected (possibly empty).*

We next make a detailed study of lifts of lines, focusing on whether they separate $\widehat{p}$ from $\widehat{q}$ in $\overline{\mathbb{D}}$ or not.

Definition Let L be a line which intersects Ω and which does not contain p or q, and let $\widehat{L}$ be a lift of L.

- We call $\widehat{L}$ a *separating lift* if it separates $\widehat{p}$ from $\widehat{q}$ in $\overline{\mathbb{D}}$. The component of $\mathbb{D} \backslash \widehat{L}$ whose closure contains $\widehat{p}$ is called the *$\widehat{p}$-side* of $\widehat{L}$, and the component of $\mathbb{D} \backslash \widehat{L}$ whose closure contains $\widehat{q}$ is called the *$\widehat{q}$-side* of $\widehat{L}$.
- We call $\widehat{L}$ a *nonseparating lift* if it does not separate $\widehat{p}$ from $\widehat{q}$ in $\overline{\mathbb{D}}$. The component of $\mathbb{D} \backslash \widehat{L}$ whose closure does not contain $\widehat{p}, \widehat{q}$ is called the *shadow* of $\widehat{L}$, denoted $\mathrm{Sh}(\widehat{L})$.

Given two lines L_1, L_2 which intersect Ω, the *distance* between L_1 and L_2 is the Hausdorff distance between $L_1 \cap \overline{D_\Omega}$ and $L_2 \cap \overline{D_\Omega}$ (recall that D_Ω is a fixed large disk in $\mathbb{C}$ containing Ω); that is, the infimum of all $\delta > 0$ such that each point of $L_1 \cap \overline{D_\Omega}$ is within δ of a point in $L_2 \cap \overline{D_\Omega}$, and vice versa.

Definition

- Let $\widehat{V} \subset \mathbb{D}$ be an open set which maps one-to-one under φ to an open set $V \subset \Omega$, and let L be a line which intersects V. Let $\varepsilon > 0$, and assume ε is small enough so that every line L' which is within distance ε of L must also intersect V. The family of all lifts $\widehat{L}'$, of such lines L', which intersect $\widehat{V}$ will be called a *basic open set of lifts of lines* , and denoted $\widehat{\mathcal{N}}(\widehat{L}, \widehat{V}, \varepsilon)$.
- Let $\widehat{z} \in \mathbb{D}$. We say $\widehat{z}$ is a *stable point* if there exists a basic open set of lifts of lines $\widehat{\mathcal{N}}(\widehat{L}, \widehat{V}, \varepsilon)$, each element of which is nonseparating and contains $\widehat{z}$ in its shadow.
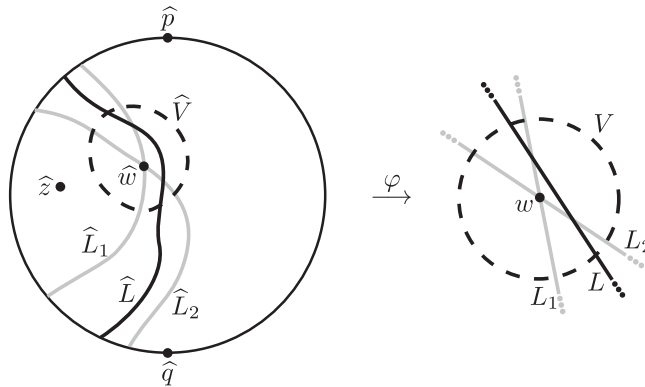
Figure 2: The configuration of lines described in the proof of Lemma 14.

Lemma 14 *A point $\widehat{z} \in \mathbb{D}$ is a stable point if and only if there are two intersecting nonseparating lifts $\widehat{L}_1, \widehat{L}_2$ of distinct lines L_1, L_2 such that $\widehat{z}$ is in the shadow of both $\widehat{L}_1$ and $\widehat{L}_2$.*

Furthermore, whenever $\widehat{L}_1, \widehat{L}_2$ are intersecting nonseparating lifts of distinct lines, the point of intersection of $\widehat{L}_1$ and $\widehat{L}_2$ is also a stable point.

Proof That any stable point has this property is immediate, since any basic open set of lifts of lines clearly contains pairs of distinct intersecting lifts.

Conversely, suppose $\widehat{L}_1, \widehat{L}_2$ are intersecting nonseparating lifts of distinct lines $L_1 \supset \varphi(\widehat{L}_1), L_2 \supset \varphi(\widehat{L}_2)$, and let $\widehat{z}$ be any point in the shadow of both $\widehat{L}_1$ and $\widehat{L}_2$. Let $\widehat{w}$ be the point of intersection of $\widehat{L}_1$ and $\widehat{L}_2$, and let $\widehat{V}$ be a neighborhood of $\widehat{w}$ which maps one-to-one under φ to a round disk V centered at $w = \varphi(\widehat{w})$.

Consider a line L not containing w which intersects $L_1 \cap V$ and $L_2 \cap V$, and such that the lift $\widehat{L}$ of L which intersects $\widehat{V}$ does not contain any point in the intersection of the shadows of $\widehat{L}_1$ and of $\widehat{L}_2$. Let $\varepsilon > 0$ be small enough so that any line L' within distance ε of L has these same properties. Let $\widehat{L}'$ be an arbitrary element of the basic open set of lifts of lines $\widehat{\mathcal{N}}(\widehat{L}, \widehat{V}, \varepsilon)$; that is, $\widehat{L}'$ is the lift of a line L' within distance ε of L such that $\widehat{L} \cap \widehat{V} \neq \varnothing$. Since $\widehat{L}'$ intersects $\widehat{L}_1$ and $\widehat{L}_2$ inside $\widehat{V}$, it cannot cross them again, and so one endpoint of $\widehat{L}'$ is in the closure of $\mathrm{Sh}(\widehat{L}_1) \backslash \mathrm{Sh}(\widehat{L}_2)$ and the other is in the closure of $\mathrm{Sh}(\widehat{L}_2) \backslash \mathrm{Sh}(\widehat{L}_1)$ (see Figure 2). It follows that $\widehat{L}'$ is a nonseparating lift and $\mathrm{Sh}(\widehat{L}')$ contains $\mathrm{Sh}(\widehat{L}_1) \cap \mathrm{Sh}(\widehat{L}_2)$ (in particular, it contains the point $\widehat{z}$), as well as the point $\widehat{w}$. Thus $\widehat{z}$ and $\widehat{w}$ are both stable points.   ∎

Lemma 15 *The set of stable points is a dense open subset of $\mathbb{D}$.*

Moreover, except for a countable set of lines, every line L has the property that for each lift $\widehat{L}$ of L, the set of stable points in $\widehat{L}$ is a dense open subset of $\widehat{L}$.

Proof It follows immediately from Lemma 14 that the set of stable points is open in $\mathbb{D}$.
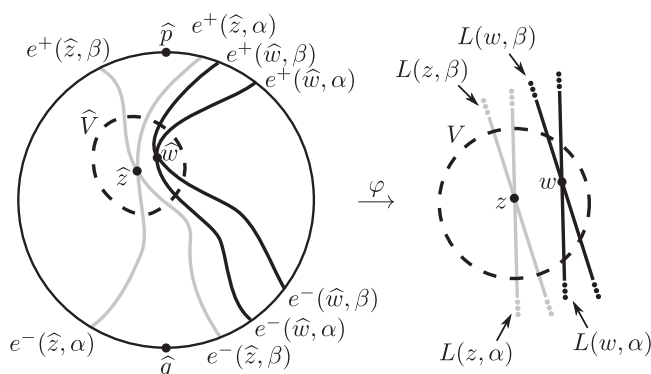
Figure 3: The configuration of lines described in the proof of Lemma 15.

Now let $\widehat{z} \in \mathbb{D}$, and let $\widehat{V}$ be a neighborhood of $\widehat{z}$. Suppose $\widehat{z}$ is not a stable point. By shrinking $\widehat{V}$, we may assume that $\widehat{V}$ maps one-to-one under φ to a disk V centered at $z = \varphi(\widehat{z})$.

For a given $\theta \in \mathbb{R}$, let $L(z, \theta)$ denote the straight line through z making angle θ with the positive real axis. Let $\widehat{L}(\widehat{z}, \theta)$ be the lift containing $\widehat{z}$ of the component of $L(z, \theta) \cap \Omega$ containing z. Denote the endpoints of $\widehat{L}(\widehat{z}, \theta)$ by $e^+(\widehat{z}, \theta)$ and $e^-(\widehat{z}, \theta)$, where $e^+(\widehat{z}, \theta)$ corresponds to following the line $L(z, \theta)$ to $\partial\Omega$ in the direction θ from z, and $e^-(\widehat{z}, \theta)$ corresponds to following the line $L(z, \theta)$ to $\partial\Omega$ in the direction $\theta + \pi$ from z. As θ increases, the line $L(z, \theta)$ revolves about the point z. Correspondingly, the lift $\widehat{L}(\widehat{z}, \theta)$ "revolves" about $\widehat{z}$. The endpoints of $\widehat{L}(\widehat{z}, \theta)$ move monotonically in the circle $\partial\mathbb{D}$, but not necessarily continuously, as θ increases.

Since $\widehat{z}$ is not stable, by Lemma 14, there can be at most one nonseparating lift of a line which contains $\widehat{z}$. It follows that if we consider the line $L(z, \theta)$ and increase θ to revolve the line about z, at some moment, the endpoint $e^+(\widehat{z}, \theta)$ must cross or "jump over" $\widehat{p}$, and at that same moment $e^-(\widehat{z}, \theta)$ must cross or "jump over" $\widehat{q}$. That is, there exists θ_0 such that for all α, β sufficiently close to θ_0 with $\alpha < \theta_0 < \beta$, we have that $\widehat{p}$ is on the "left" side of the arc $\widehat{L}(\widehat{z}, \alpha)$ (thinking of this arc as oriented from $e^-(\widehat{z}, \alpha)$ to $e^+(\widehat{z}, \alpha)$) and $\widehat{q}$ is on the "right" side, and $\widehat{p}$ is on the "right" side of $\widehat{L}(\widehat{z}, \beta)$ and $\widehat{q}$ is on the "left" side (see Figure 3).

Let $\widehat{w}$ be any point in $\widehat{V} \backslash \widehat{L}(\widehat{z}, \theta_0)$ and let $w = \varphi(\widehat{w})$. Let α, β be as above and sufficiently close to θ_0 so that the lines $L(w, \alpha)$ and $L(w, \beta)$ do not intersect either of the lines $L(z, \alpha)$ or $L(z, \beta)$ inside $\overline{D_\Omega}$ (recall that D_Ω is a fixed large disk in $\mathbb{C}$ containing Ω). This means that the lifts $\widehat{L}(\widehat{w}, \alpha)$ and $\widehat{L}(\widehat{w}, \beta)$ do not cross either of the lifts $\widehat{L}(\widehat{z}, \alpha)$ or $\widehat{L}(\widehat{z}, \beta)$. Because of the locations of $\widehat{p}$ and $\widehat{q}$ with respect to the lines $\widehat{L}(\widehat{z}, \alpha)$ and $\widehat{L}(\widehat{z}, \beta)$, it follows that $\widehat{L}(\widehat{w}, \alpha)$ and $\widehat{L}(\widehat{w}, \beta)$ are both nonseparating lifts (see Figure 3). Hence, by Lemma 14, $\widehat{w}$ is a stable point. Thus, the set of stable points is dense in $\mathbb{D}$.

For the second statement, we may apply the above argument at each $\widehat{z} \in \mathbb{D}$ to obtain a countable cover $\{\widehat{V}_i : i = 1, 2, \ldots\}$ of $\mathbb{D}$ by open sets with the property that for each i there is at most one lift of a line, $\widehat{L}_i$, such that every point of $\widehat{V}_i \backslash \widehat{L}_i$ is stable. Consider

the countable family of lines $\{\varphi(\widehat{L}_i) : i = 1, 2, \ldots\}$. Let L be any line not in this family, let $\widehat{L}$ be any lift of L, and let $\widehat{z} \in \widehat{L}$. Choose i so that $\widehat{z} \in V_i$. If $\widehat{z}$ is not stable, it must be the (unique) point of intersection of $\widehat{L}$ and $\widehat{L}_i$, hence every other point in $\widehat{L} \cap V_i$ is stable. Therefore the set of stable points in $\widehat{L}$ is a dense open subset of $\widehat{L}$.   ∎

We remark that since in the proof of Lemma 15, the point $\widehat{w}$ could be chosen on either side of $\widehat{L}(\widehat{z}, \theta_0)$, it follows from Proposition 13 and the proof of Lemma 15 that if $\widehat{\gamma}_0$ is a lift of an efficient path in $\overline{[\gamma]}$ with endpoints $\widehat{p}$ and $\widehat{q}$, then $\widehat{\gamma}_0([0,1]) \cap \mathbb{D}$ coincides with the set of nonstable points. This observation will not be used in the proof of Theorem 1.

3.2   Sequence of approximations of the efficient path

Let $\mathcal{L}$ be a countable dense (in the sense of Hausdorff distanct) family of distinct straight lines which intersect Ω and do not contain p or q, and such that for each lift $\widehat{L}$ of a line $L \in \mathcal{L}$, the set of stable points in $\widehat{L}$ is a dense open subset of $\widehat{L}$ (this is possible by Lemma 15). Let $\widehat{\mathcal{L}}$ denote the (countable) set of all lifts of lines in $\mathcal{L}$. We enumerate the elements of this set: $\widehat{\mathcal{L}} = \langle \widehat{L}_i \rangle_{i=1}^{\infty}$.

We construct a sequence of paths γ_i, $i \geq 1$, by recursion. To begin, let $\gamma_1 = \gamma$ and $\widehat{\gamma}_1 = \widehat{\gamma}$. Having defined γ_i and its lift $\widehat{\gamma}_i$, we define γ_{i+1} and $\widehat{\gamma}_{i+1}$ as follows:

- If $\widehat{\gamma}_i^{-1}(\widehat{L}_i)$ has cardinality ≤ 1, then put $\gamma_{i+1} = \gamma_i$ and $\widehat{\gamma}_{i+1} = \widehat{\gamma}_i$. Otherwise, let s_1 and s_2 be the smallest and largest (respectively) $s \in [0,1]$ such that $\widehat{\gamma}_i(s) \in \widehat{L}_i$. Let $\widehat{\gamma}_{i+1}$ be the path in $\mathbb{D}$ (e.p.e.) defined by $\widehat{\gamma}_{i+1}(s) = \widehat{\gamma}_i(s)$ for $s \notin [s_1, s_2]$, and $\widehat{\gamma}_{i+1}\restriction_{[s_1,s_2]}$ parameterizes the subarc of $\widehat{L}_i$ with endpoints $\widehat{\gamma}_i(s_1)$ and $\widehat{\gamma}_i(s_2)$ (or $\widehat{\gamma}_{i+1}\restriction_{[s_1,s_2]}$ is constantly equal to $\widehat{w}$ if $\widehat{\gamma}_i(s_1) = \widehat{\gamma}_i(s_2) = \widehat{w}$). Let $\gamma_{i+1} = \overline{\varphi} \circ \widehat{\gamma}_{i+1}$.

Lemma 16   *Let $\widehat{L}$ be a lift of a line L. If i is such that $\widehat{\gamma}_i^{-1}(\widehat{L}) \subset [0,1]$ is connected (respectively, empty), then $\widehat{\gamma}_j^{-1}(\widehat{L})$ is connected (respectively, empty) for all $j \geq i$.*

In particular, for all $j > i \geq 1$, $\widehat{\gamma}_j^{-1}(\widehat{L}_i)$ is connected (or empty).

Proof   Let $\widehat{L}$ be a lift of a line L, and fix i. We will argue that if $\widehat{\gamma}_i^{-1}(\widehat{L}) \subset [0,1]$ is connected (respectively, empty), then $\widehat{\gamma}_{i+1}^{-1}(\widehat{L})$ is connected (respectively, empty). The claim then follows by induction. To this end, given the definition of $\widehat{\gamma}_{i+1}$, clearly we may assume $\widehat{L} \neq \widehat{L}_i$, so that either $\widehat{L} \cap \widehat{L}_i \cap \mathbb{D} = \varnothing$ or $\widehat{L}$ and $\widehat{L}_i$ cross transversally in $\mathbb{D}$.

Assume first that $\widehat{\gamma}_i([0,1]) \cap \widehat{L} = \varnothing$. This means in particular that $\widehat{p}, \widehat{q} \notin \widehat{L}$, $\widehat{L}$ is a nonseparating lift, and $\widehat{\gamma}_i([0,1])$ is disjoint from the shadow of $\widehat{L}$. If $\widehat{L}_i$ does not meet $\widehat{L}$ in $\mathbb{D}$, then clearly $\widehat{\gamma}_{i+1}([0,1]) \cap \widehat{L} = \varnothing$ as well. If $\widehat{L}_i$ meets $\widehat{L}$ in $\mathbb{D}$, then these two lifts cross transversally. If $\widehat{\gamma}_{i+1}([0,1])$ meets $\widehat{L}$, we must have that $\widehat{\gamma}_{i+1}([s_1,s_2]) \cap \widehat{L} \neq \varnothing$ where $s_1 = \min\{s \in [0,1] : \widehat{\gamma}_i(s) \in \widehat{L}_i\}$ and $s_2 = \max\{s \in [0,1] : \widehat{\gamma}_i(s) \in \widehat{L}_i\}$, since $\widehat{\gamma}_{i+1} = \widehat{\gamma}_i$ outside of $[s_1,s_2]$. Since $\widehat{\gamma}_{i+1}([s_1,s_2])$ parameterizes a subarc of $\widehat{L}_i$, which crosses $\widehat{L}$ transversally, it follows that either $\widehat{\gamma}_{i+1}(s_1)$ or $\widehat{\gamma}_{i+1}(s_2)$ is in the shadow of $\widehat{L}$. But $\widehat{\gamma}_{i+1}(s_1) = \widehat{\gamma}_i(s_1)$ and $\widehat{\gamma}_{i+1}(s_2) = \widehat{\gamma}_i(s_2)$, so this contradicts the assumption that $\widehat{\gamma}_i([0,1]) \cap \widehat{L} = \varnothing$. Thus $\widehat{\gamma}_{i+1}([0,1]) \cap \widehat{L} = \varnothing$, as desired.

Now assume that $\widehat{\gamma}_i([0,1]) \cap \widehat{L} \neq \varnothing$, and $\widehat{\gamma}_i^{-1}(\widehat{L})$ is an interval $[a,b]$ (here we allow the possibility that $a = b$, which means $[a,b] = \{a\}$). Suppose $\widehat{\gamma}_{i+1}$ is obtained from

$\widehat{\gamma}_i$ by replacing the subpath of $\widehat{\gamma}_i$ from s_1 to s_2 ($0 < s_1 < s_2 < 1$) with a segment of $\widehat{L}_i$. We leave it to the reader to confirm the following claims:

- If $s_2 \le a$, then $\widehat{\gamma}_{i+1}^{-1}(\widehat{L}) = [a, b]$;
- If $s_1 \le a \le s_2 \le b$, then $\widehat{\gamma}_{i+1}^{-1}(\widehat{L}) = [s_2, b]$;
- If $s_1 \le a \le b \le s_2$, then $\widehat{\gamma}_{i+1}^{-1}(\widehat{L})$ is either empty (if the section $\widehat{\gamma}_{i+1}([s_1, s_2])$ of $\widehat{L}_i$ does not meet $\widehat{L}$), or consists of the single point where the section $\widehat{\gamma}_{i+1}([s_1, s_2])$ of $\widehat{L}_i$ crosses $\widehat{L}$;
- It is not possible that $a \le s_1 < s_2 \le b$, since $\widehat{L} \ne \widehat{L}_i$ (unless $\widehat{\gamma}_i$ is constant on $[s_1, s_2]$, in which case $\widehat{\gamma}_{i+1} = \widehat{\gamma}_i$);
- If $a \le s_1 \le b \le s_2$, then $\widehat{\gamma}_{i+1}^{-1}(\widehat{L}) = [a, s_1]$; and
- If $b \le s_1$, then $\widehat{\gamma}_{i+1}^{-1}(\widehat{L}) = [a, b]$.

In all cases, $\widehat{\gamma}_{i+1}^{-1}(\widehat{L})$ is connected (possibly empty), as desired. ∎

3.3 Convergence to the efficient path

For each $n = 1, 2, \ldots$, we choose sets $\mathcal{G}_n$ of lines in $\mathcal{L}$ with the following properties:

- $\mathcal{G}_n$ is a finite subset of $\mathcal{L}$ for each n;
- $\mathcal{G}_n \subset \mathcal{G}_{n+1}$ for each n;
- each component of $\Omega \setminus \bigcup \mathcal{G}_n$ has diameter less than $\frac{1}{n}$; and
- if $G_1, G_2 \in \mathcal{G}_n$ are distinct lines and $\widehat{G}_1$ and $\widehat{G}_2$ are separating lifts of G_1 and G_2, respectively, then $\widehat{G}_1$ and $\widehat{G}_2$ are either disjoint or meet in a stable point.

Let $\widehat{\mathcal{G}}_n$ be the set of all lifts of lines in $\mathcal{G}_n$. By Lemma 9 only finitely many of the elements of $\widehat{\mathcal{G}}_n$ intersect the original lifted path $\widehat{\gamma}_1$, and among them are those that are separating lifts. We denote the set of all separating lifts in $\widehat{\mathcal{G}}_n$ by $\widehat{\mathcal{G}}_n^s$.

We define an order $<$ on $\widehat{\mathcal{G}}_n^s$ as follows. Let $\widehat{G}_1, \widehat{G}_2 \in \widehat{\mathcal{G}}_n^s$ be distinct separating lifts.

First suppose $\widehat{G}_1 \cap \widehat{G}_2 = \varnothing$. Then $\widehat{G}_1 < \widehat{G}_2$ if $\widehat{G}_2$ is on the $\widehat{q}$-side of $\widehat{G}_1$; otherwise $\widehat{G}_2 < \widehat{G}_1$.

Now suppose $\widehat{G}_1 \cap \widehat{G}_2 = \{\widehat{z}\}$ for some stable point $\widehat{z} \in \mathbb{D}$. Since $\widehat{z}$ is stable, it follows from density of the family of lines $\mathcal{L}$ that there exists $\widehat{W} \in \widehat{\mathcal{L}}$ such that $\widehat{W}$ is nonseparating and contains $\widehat{z}$ in its shadow. Then $\widehat{G}_1 < \widehat{G}_2$ if $\widehat{G}_2 \setminus \overline{\mathrm{Sh}(\widehat{W})}$ is on the $\widehat{q}$-side of $\widehat{G}_1$; otherwise $\widehat{G}_2 < \widehat{G}_1$.

It is straightforward to see that this relation $<$ is a well-defined linear order on $\widehat{\mathcal{G}}_n^s$.

For each $n = 1, 2, \ldots$, choose a finite set $\widehat{\mathcal{W}}_n$ of lifts of lines from $\widehat{\mathcal{L}}$ such that for each pair of distinct intersecting separating lifts $\widehat{G}_1, \widehat{G}_2 \in \widehat{\mathcal{G}}_n^s$, there is a lift $\widehat{W} \in \widehat{\mathcal{W}}_n$ witnessing that $\widehat{G}_1 < \widehat{G}_2$ or that $\widehat{G}_2 < \widehat{G}_1$ (i.e. such that $\mathrm{Sh}(\widehat{W})$ contains $\widehat{G}_1 \cap \widehat{G}_2$). Let $i(n)$ be large enough so that $\widehat{\mathcal{W}}_n \subset \{\widehat{L}_i : i = 1, \ldots, i(n) - 1\}$, and also all of the (finitely many) elements of $\widehat{\mathcal{G}}_n$ which intersect the original lifted path $\widehat{\gamma}_1$ are contained in $\{\widehat{L}_i : i = 1, \ldots, i(n) - 1\}$.

Take an arbitrary n, and let $\widehat{G}_1, \ldots, \widehat{G}_m$ enumerate the elements of $\widehat{\mathcal{G}}_n^s$, listed in $<$ order. For any $i \ge i(n)$, by Lemma 16, the lifted path $\widehat{\gamma}_i$ does not cross any of the nonseparating lifts in $\widehat{\mathcal{G}}_n \cup \widehat{\mathcal{W}}_n$, and has connected intersection with all the lifts that it meets, in particular with the elements of $\widehat{\mathcal{G}}_n^s$. Thus there are $m + 1$ components $\widehat{R}_1, \ldots, \widehat{R}_{m+1}$ of $\mathbb{D} \setminus \bigcup(\widehat{\mathcal{G}}_n \cup \widehat{\mathcal{W}}_n)$ such that for any $i \ge i(n)$, the lifted path $\widehat{\gamma}_i$ starts
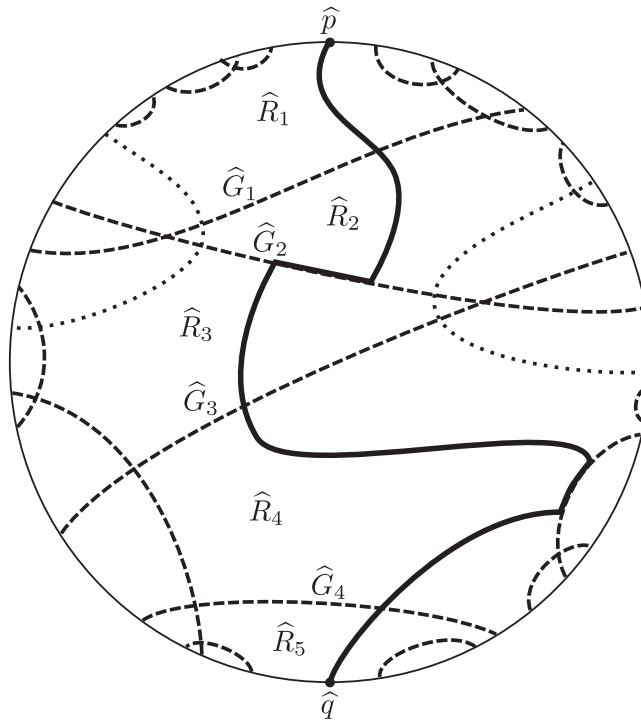
Figure 4: An illustration depicting a possible configuration of the lifts of lines $\widehat{G}_k$ and regions $\widehat{R}_k$ in between, as described in Section 3.3. The thick curve represents a path $\widehat{\gamma}_i$ (for $i \geq i(n)$), the dashed arcs represent elements of $\widehat{\mathcal{G}}_n$ (including the separating lifts $\widehat{G}_k$), and the dotted arcs represent elements of $\widehat{\mathcal{W}}_n$.

at $\widehat{p}$ then runs in $\widehat{R}_1$, until it crosses $\widehat{G}_1$ and enters $\widehat{R}_2$, then crosses $\widehat{G}_2$ and enters $\widehat{R}_3$, and so on, until it crosses $\widehat{G}_m$ to enter $\widehat{R}_{m+1}$ and ends at $\widehat{q}$ (see Figure 4).

For each of these regions $\widehat{R}_j$, $\varphi(\widehat{R}_j)$ is contained in a component of $\Omega \setminus \bigcup \mathcal{G}_n$, hence has diameter less than $\frac{1}{n}$. Therefore, provided we parameterize the path γ_i so that if $\widehat{\gamma}_{i(n)}(t) \in \widehat{R}_j$ then $\widehat{\gamma}_i(t) \in \widehat{R}_j$ for all $i \geq i(n)$, we have that $d_{\sup}(\gamma_{i_1}, \gamma_{i_2}) < \frac{1}{n}$ for all $i_1, i_2 \geq i(n)$; in fact there is a homotopy between γ_{i_1} and γ_{i_2} which moves no point more than $\frac{1}{n}$, obtained by moving $\widehat{\gamma}_{i_1}(t)$ to $\widehat{\gamma}_{i_2}(t)$ within the closure of a region $\widehat{R}_j$ which contains $\widehat{\gamma}_{i(n)}(t)$.

We remark that according to Proposition 13, if $\lambda \in \overline{[\gamma]}$ is any efficient path, it must follow this same pattern traversing through the (closures of the) regions $\widehat{R}_1, \ldots, \widehat{R}_{m+1}$ in the same monotone order. Therefore, any two such efficient paths, if parameterized appropriately, are within $\frac{1}{n}$ of each other. Since n is arbitrary, this establishes that there can be at most one efficient path in $\overline{[\gamma]}$ (up to reparameterization).

It remains to confirm that our construction above yields an efficient path in $\overline{[\gamma]}$. Since n was arbitrary, we now have that $\langle \gamma_i \rangle_{i=1}^{\infty}$ is a Cauchy sequence of paths, hence

it converges to a path γ_∞ from p to q. Moreover, by putting together the (smaller and smaller) homotopies between the paths in this sequence, we have that $\gamma_\infty \in \overline{[\gamma]}$.

Lemma 17 γ_∞ *is an efficient path in* $\overline{[\gamma]}$.

Proof   This follows from Proposition 13: if L is any line intersecting Ω which has a lift $\widehat{L}$ such that $\widehat{\gamma}_\infty([0,1]) \cap \widehat{L}$ is not connected, then by density of the family $\mathcal{L}$ it is easy to see there must be a line $L' \in \mathcal{L}$ close to L, with a lift $\widehat{L}'$ close to $\widehat{L}$, such that $\widehat{\gamma}_\infty([0,1]) \cap \widehat{L}'$ is also not connected. But if $\widehat{L}' = \widehat{L}_i$ in the enumeration of $\widehat{\mathcal{L}}$, then for all $j > i$, $\widehat{\gamma}_j([0,1]) \cap \widehat{L}'$ is connected according to Lemma 16. Since $\widehat{\gamma}_i \to \widehat{\gamma}_\infty$, the same must be true for $\widehat{\gamma}_\infty$, a contradiction. ∎

This completes the proof of Theorem 1.

4   Proof of Theorem 2

Having established the existence of a unique efficient path λ in $\overline{[\gamma]}$, we now prove Theorem 2 using Theorem 1 and its proof. In particular, we observe that in the construction of the efficient path in the proof of Theorem 1, we started with a path in $[\gamma]$ (such as γ itself) and constructed a sequence of paths converging to the efficient path, where each path in the sequence is obtained from the previous one by replacing a subpath with a straight segment. Such a replacement will always decrease the len length of the path and also, when finite, the Euclidean length.

It is clear that any path in $\overline{[\gamma]}$ of smallest len length (or Euclidean length if finite) must be efficient, because replacing a subpath with a straight segment decreases the length. Thus we have (3) $\Rightarrow$ (2) and (4) $\Rightarrow$ (2).

For (2) $\Rightarrow$ (3), suppose for a contradiction that there is a path $\rho \in \overline{[\gamma]}$ with strictly smaller len length than the efficient path $\lambda \in \overline{[\gamma]}$. By continuity of the function len, it follows that there exists $\rho' \in [\gamma]$ whose len length is also strictly smaller than λ. If we follow the construction in the proof of Theorem 1 starting with ρ' instead of γ, we would then obtain an efficient path in $\overline{[\gamma]}$ of len length strictly smaller than λ, contradicting the uniqueness of the efficient path.

For (2) $\Rightarrow$ (4), suppose for a contradiction that there is a path $\rho \in \overline{[\gamma]}$ with finite Euclidean length E which is strictly smaller than the Euclidean length of the efficient path $\lambda \in \overline{[\gamma]}$. This means there exists $\varepsilon > 0$ and values $0 = s_0 < s_1 < \cdots < s_N = 1$ such that $\sum_{i=1}^N |\lambda(s_i) - \lambda(s_{i-1})| > E + \varepsilon$. Let $\widehat{\gamma}$ be a lift of γ with endpoints $\widehat{p}, \widehat{q} \in \overline{\mathbb{D}}$, and let $\widehat{\lambda}$ and $\widehat{\rho}$ be lifts of λ and ρ with the same endpoints $\widehat{p}, \widehat{q}$.

We will replace sections of the path ρ with straight line segments to obtain a new path $\rho' \in \overline{[\gamma]}$ which goes within $\frac{\varepsilon}{2N}$ of each of the points $\lambda(s_i)$, in order. To obtain ρ', we proceed as follows. For each $i = 1, \ldots, N-1$, consider the point $\widehat{\lambda}(s_i) \in \overline{\mathbb{D}}$.

- If $\widehat{\lambda}(s_i) \in \mathbb{D}$, by Lemma 15, we can choose two stable points $\widehat{w}_1, \widehat{w}_2$ in a small neighborhood $\widehat{V}$ of $\widehat{\lambda}(s_i)$ which projects one-to-one under φ to a small disk centered at $\lambda(s_i)$ of radius less than $\frac{\varepsilon}{2N}$, so that $\widehat{w}_1$ and $\widehat{w}_2$ are on opposite sides of $\widehat{\lambda}([0,1]) \cap \widehat{V}$. Let $\widehat{L}_1$ and $\widehat{L}_2$ be nonseparating lifts of lines such that $\widehat{w}_1 \in \mathrm{Sh}(\widehat{L}_1)$

and $\widehat{w}_2 \in \mathrm{Sh}(\widehat{L}_2)$. By Proposition 13, $\widehat{\lambda}$ does not cross $\widehat{L}_1$ or $\widehat{L}_2$. It follows that if we replace the section of $\widehat{\rho}$ between its first and last intersections with $\widehat{L}_1$ with an arc in $\widehat{L}_1$, and likewise for $\widehat{L}_2$, then the resultant path must meet $\widehat{V}$.

- If $\widehat{\gamma}(s_i) \in \partial\mathbb{D}$, choose an arc $\widehat{A}$ such that $\overline{\varphi}(\widehat{A})$ has diameter less than $\frac{\varepsilon}{2N}$, with endpoints $\widehat{\lambda}(s_i)$ and a stable point $\widehat{w} \in \mathbb{D}$. Let $\widehat{L}$ be a lift of a line such that $\widehat{w} \in \mathrm{Sh}(\widehat{L})$. By Proposition 13, $\widehat{\lambda}$ does not cross $\widehat{L}$. It follows that if we replace the section of $\widehat{\rho}$ between its first and last intersections with $\widehat{L}$ with an arc in $\widehat{L}$, then the resultant path must meet $\widehat{A}$.

After making the finitely many replacements of subpaths of $\widehat{\rho}$ with segments of lifts of lines as described above, the resultant path $\widehat{\rho}'$ is such that $\rho' = \overline{\varphi} \circ \widehat{\rho}'$ is continuous, and $\rho' \in \overline{[\gamma]}$ by Theorem 11. Clearly, the Euclidean length of ρ' is not greater than E. On the other hand, by construction, there are values $0 = s_0' < s_1' < \cdots < s_N' = 1$ such that $|\rho'(s_i') - \lambda(s_i)| < \frac{\varepsilon}{2N}$ for each $i = 0, \ldots, N$. It follows that the Euclidean length of ρ' is at least

$$\sum_{i=1}^{N} |\rho'(s_i') - \rho'(s_{i-1}')| > \sum_{i=1}^{N} \left(|\lambda(s_i) - \lambda(s_{i-1})| - 2 \cdot \frac{\varepsilon}{2N} \right)$$

$$= \left(\sum_{i=1}^{N} |\lambda(s_i) - \lambda(s_{i-1})| \right) - \varepsilon$$

$$> (E + \varepsilon) - \varepsilon = E,$$

a contradiction. Therefore, λ has smallest Euclidean length among all paths in $\overline{[\gamma]}$.

The proof that (1) $\Rightarrow$ (2) is straightforward and left to the reader.

Finally, for (2) $\Rightarrow$ (1), let $0 < s_1 < s_2 < 1$ and consider the path $\lambda\restriction_{[s_1,s_2]}$. We claim that there exist s_1', s_2' such that $0 < s_1' \leq s_1 < s_2 \leq s_2' < 1$ and $\widehat{\lambda}(s_1')$ and $\widehat{\lambda}(s_2')$ can be joined by a path in $\overline{\mathbb{D}}$ which projects under $\overline{\varphi}$ to a path of finite Euclidean length. We rely on the fact that any two points $\widehat{w}_1, \widehat{w}_2 \in \mathbb{D}$ can be joined by a path in $\mathbb{D}$ which projects to a path of finite Euclidean length (e.g. a piecewise linear path). If $\widehat{\lambda}([0, s_1]) \cap \mathbb{D} \neq \varnothing$, then let $0 < s_1' \leq s_1$ be such that $\widehat{\lambda}(s_1') \in \mathbb{D}$, and let $\widehat{w}_1 = \widehat{\lambda}(s_1')$. If $\widehat{\lambda}([0, s_1]) \subset \partial\mathbb{D}$, we choose s_1' and $\widehat{w}_1$ as follows. Note that in this case, $\lambda(s_1)$ is not isolated in $\partial\Omega$, so we can apply Theorem 8 to obtain a component $\widehat{C}_1$ of $\varphi^{-1}(\partial B(\lambda(s_1), \varepsilon))$, for some sufficiently small $\varepsilon > 0$, whose closure is an arc which separates $\widehat{\lambda}(s_1)$ from $\widehat{p}$ in $\overline{\mathbb{D}}$. In particular, there exists $0 < s_1' < s_1$ such that $\widehat{\lambda}(s_1')$ is in the closure of $\widehat{C}_1$. Let $\widehat{w}_1$ be any point in $\widehat{C}_1$. Thus, we have established that there exists $0 < s_1' \leq s_1$ and $\widehat{w}_1 \in \mathbb{D}$ such that either $\widehat{\lambda}(s_1') = \widehat{w}_1$ or $\widehat{\lambda}(s_1')$ can be joined to $\widehat{w}_1$ by a path in $\overline{\mathbb{D}}$ which projects to a circular arc. By the same reasoning, there exists $s_2 \leq s_2' < 1$ and $\widehat{w}_2 \in \mathbb{D}$ such that either $\widehat{\lambda}(s_2') = \widehat{w}_2$ or $\widehat{\lambda}(s_2')$ can be joined to $\widehat{w}_2$ by a path in $\overline{\mathbb{D}}$ which projects to a circular arc. It follows that $\widehat{\lambda}(s_1')$ can be joined to $\widehat{\lambda}(s_2')$ by a path in $\overline{\mathbb{D}}$ which projects to a path of finite Euclidean length. Therefore, by Corollary 12, there exists a path of finite Euclidean length in $\overline{[\lambda_{[s_1',s_2']}^{\gamma}]}$.

Since λ is efficient, it follows from Proposition 13 that $\lambda\restriction_{[s_1',s_2']}$ is the efficient path in $\overline{[\lambda_{[s_1',s_2']}^{\gamma}]}$, hence has smallest Euclidean length in this class according to the implication (2) $\Rightarrow$ (4). In particular, $\lambda\restriction_{[s_1',s_2']}$ has finite Euclidean length in light of the previous paragraph. This in turn implies that the subpath $\lambda\restriction_{[s_1,s_2]}$ has finite Euclidean length.

Again by Proposition 13, we have that $\lambda\!\restriction_{[s_1,s_2]}$ is the efficient path in $\overline{\left[\lambda^\gamma_{[s_1,s_2]}\right]}$, hence has smallest Euclidean length in this class according to the implication $(2) \Rightarrow (4)$, as required. Therefore, λ is locally shortest.

References

[1] L. V. Ahlfors, *Conformal invariants: topics in geometric function theory*. McGraw-Hill Series in Higher Mathematics, McGraw-Hill Book Co., New York, Düsseldorf, Johannesburg, 1973. MR 0357743.

[2] R. D. Bourgin and P. L. Renz, *Shortest paths in simply connected regions in R^2*. Adv. Math. **76**(1989), no. 2, 260–295. MR 1013673.

[3] J. W. Cannon, G. R. Conner, and A. Zastrow, *One-dimensional sets and planar sets are aspherical*. Topol. Appl. **120**(2002), nos. 1–2, 23–45. In memory of T. Benny Rushing. MR 1895481.

[4] J. B. Conway, *Functions of one complex variable. II*. Graduate Texts in Mathematics, 159, Springer-Verlag, New York, NY, 1995. MR 1344449.

[5] P. Fatou, *Séries trigonométriques et séries de Taylor*. Acta Math. **30**(1906), no. 1, 335–400. MR 1555035.

[6] L. C. Hoehn, L. G. Oversteegen, and E. D. Tymchatyn, *Extension of isotopies in the plane*. Trans. Amer. Math. Soc. **372**(2019), no. 7, 4889–4915. MR 4009398.

[7] L. C. Hoehn, L. G. Oversteegen, and E. D. Tymchatyn, *A canonical parameterization of paths in $\mathbb{R}^n$*. Houston J. Math. **46**(2020), 465–489.

[8] C. Kent, *Homotopy type of planar continua*. Preprint, 2020, arXiv:1709.09211

[9] K. Kuratowski, *Topology, Vol. II*. New ed., revised and augmented. Translated from the French by A. Kirkor. Academic Press, New York, NY, 1968.

[10] E. Lindelöf, *Sur un principe générale de l'analyse et ses applications à la théorie de la représentation conforme*. Acta Soc. Sci. Fennicae **46**(1915), 1–35.

[11] M. Morse, *A special parameterization of curves*. Bull. Amer. Math. Soc. **42**(1936), no. 12, 915–922. MR 1563464.

[12] C. Pommerenke, *Boundary behaviour of conformal maps*. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 299, Springer-Verlag, Berlin, Germany, 1992. MR 1217706.

[13] F. Riesz, *Über die Randwerte einer analytischen Funktion*. Math. Z. **18**(1923), no. 1, 87–95. MR 1544621.

[14] F. Riesz and M. Riesz, *Über die Randwerte einer analytischen Function* 4. Cong. Scand. Math. Stockholm 1916, pp. 87–95.

[15] E. Silverman, *Equicontinuity and n-length*. Proc. Amer. Math. Soc. **20**(1969), 483–486. MR 0237719 (38 #6000).

Department of Computer Science & Mathematics, Nipissing University, 100 College Drive, Box 5002, North Bay, ON P1B 8L7, Canada
e-mail: loganh@nipissingu.ca

Department of Mathematics, University of Alabama at Birmingham, Birmingham, AL 35294, USA
e-mail: overstee@uab.edu

Department of Mathematics and Statistics, University of Saskatchewan, 106 Wiggins Road, Saskatoon, SK S7N 5E6, Canada
e-mail: tymchat@math.usask.ca