Bounded Distortion Harmonic Mappings in the Plane

Renjie Chen
University of North Carolina at Chapel Hill, USA

Ofir Weber
Bar Ilan University, Israel

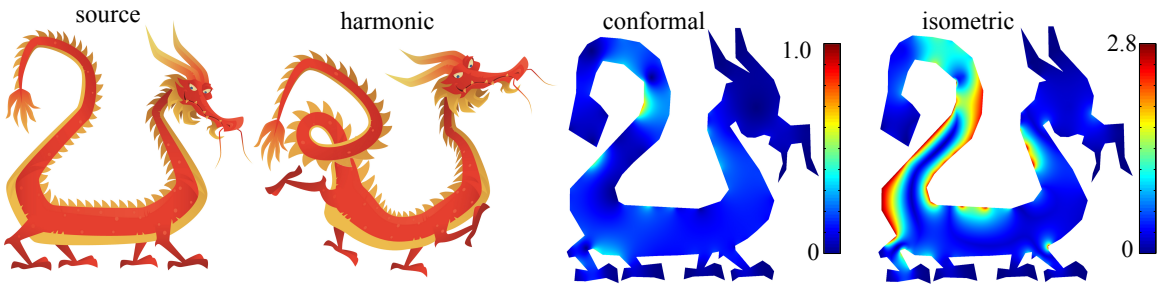

Figure 1: *Smooth bounded distortion harmonic mapping of the Dragon. The conformal distortion is bounded globally by 0.73 and the isometric distortion is bounded by 2.8.*

Abstract

We present a framework for the computation of harmonic and conformal mappings in the plane with mathematical guarantees that the computed mappings are C^∞, locally injective and satisfy strict bounds on the conformal and isometric distortion. Such mappings are very desirable in many computer graphics and geometry processing applications.

We establish the sufficient and necessary conditions for a harmonic planar mapping to have bounded distortion. Our key observation is that these conditions relate solely to the boundary behavior of the mapping. This leads to an efficient and accurate algorithm that supports handle-based interactive shape-and-image deformation and is demonstrated to outperform other state-of-the-art methods.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: injective mappings, conformal mappings, harmonic mappings, bounded distortion, planar deformation

1 Introduction

Computation of mappings between flat or curved surfaces is a fundamental task in computer graphics and geometry processing. Among the popular applications are shape-and-image deformation, shape interpolation-and-animation, surface parametrization, texture mapping, and finding dense correspondence between shapes for comparison and analysis. Harmonic mappings are extremely popular and were heavily used and investigated by researchers in the last two decades [Sheffer and de Sturler 2001; Lévy et al. 2002; Zayer et al. 2005; Kharevych et al. 2006; Weber et al. 2007; Joshi et al. 2007; Springborn et al. 2008; Ben-Chen et al. 2009; Weber and Gotsman 2010].

Harmonic mappings are considered very desirable for computer graphics applications due to their smoothness and the fact that they are mathematically and computationally easy to work with. For example, fixing the boundary values of the mapping uniquely dictates its values throughout the entire domain, implying that harmonic mappings form a low dimensional subspace. This subspace is also linear in the sense that linear combinations of harmonic functions remain harmonic.

While smoothness is one of the most desirable properties when it comes to graphics, it is generally not sufficient. An artist would often like to produce deformations that are locally injective and that preserve the orientation of local shape details rather than reverse it. Moreover, it is often desirable to control and restrict the amount of angle and isometric distortion induced by the mapping. Yet, the vast majority of existing methods do not explicitly enforce constraints that impose the desirable properties but rather try to obtain these properties indirectly. The main reason is the nonconvex nature of the constraints which leads to the inability to design efficient computational methods. Recently, new techniques to efficiently handle such nonconvex problems were introduced (e.g. [Lipman 2012]), leading to an increasing interest of the scientific community in solving such challenging problems.

Methods that are based on constrained minimization (hence provide strict guarantees) are almost exclusively mesh based. They represent the mappings as piecewise linear (PWL) approximations and are nonsmooth (C^0) by construction. Refining the resolution of the meshes typically improves the result visually (at the cost of longer computation times) but cannot change the fact that mathematically PWL mappings are nonsmooth.

In contrast, the recent algorithm of [Poranne and Lipman 2014] provides the first attempt to borrow some of the advanced techniques that were developed in the mesh based setting to the meshless one. Our algorithm improves upon [Poranne and Lipman 2014] in several different aspects, allowing us to compute smooth (C^∞) harmonic and conformal deformations with strict quality guarantees.

We develop an elegant mathematical theory that provides necessary and sufficient conditions for harmonic (and as a special case conformal) mappings to be locally injective and have bounded amount

ACM Reference Format
Chen, R., Weber, O. 2015. Bounded Distortion Harmonic Mappings in the Plane. ACM Trans. Graph. 34, 4, Article 73 (August 2015), 12 pages. DOI = 10.1145/2766989 http://doi.acm.org/10.1145/2766989.

SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.

DOI: http://doi.acm.org/10.1145/2766989

of conformal and isometric distortion. The remarkable fact is that our conditions are formulated solely in terms of boundary values. This leads to efficient and accurate constrained optimization procedure that produces (for the first time) smooth bounded distortion harmonic and conformal mappings.

2 Previous Work

Due to the abundance of methods that deals with computation of mappings, we only review the most closely related works.

Smooth barycentric coordinates Barycentric coordinates provide a simple and efficient way to compute smooth spatial deformations in two dimensions [Hormann and Floater 2006; Weber et al. 2009; Weber and Gotsman 2010; Weber et al. 2011; Weber et al. 2012b] and in three dimensions [Ju et al. 2005; Floater et al. 2005; Lipman et al. 2008; Joshi et al. 2007; Ben-Chen et al. 2009]. In the planar case, the deformation equation is a simple linear combination of real or complex-valued basis functions (the coordinates) with some complex coefficients. The efficiency stems from the fact that the basis functions are computed offline, before the interactive deformation session begins. Most research efforts were concentrated on the design of basis functions with good properties. However, barycentric mappings are in general not injective for an arbitrary choice of the coefficients [Jacobson 2013]. An exception is the special case of Wachspress coordinates where both source and target polygons are convex [Floater and Kosinka 2010]. The method of Schneider et al. [2013] computes a bijective mapping between two simple polygons by composing several mean value barycentric mappings, but there is no control over distortion. Moreover, the method only deals with global bijections and cannot produce locally injective mappings. [Weber and Gotsman 2010] presented a method to compute C^∞ conformal mappings based on the complex Hilbert barycentric coordinates but lacks the ability to handle positional constraints.

Variational methods Another popular choice to perform shape deformation is to solve a variational problem, either by using finite element discretization over a mesh [Igarashi et al. 2005; Sorkine and Alexa 2007; Liu et al. 2008] or by using smooth basis functions [Weber et al. 2009; Ben-Chen et al. 2009; Weber et al. 2012b]. However, none of these methods can guarantee that the produced mapping is injective or that it has bounded amount of distortion.

Bounded distortion The method of Weber et al. [2012a] is specifically designed to produce injective mappings with minimal amount of conformal distortion called extremal quasiconformal mappings. Based on the barrier method, [Schüller et al. 2013] produces a locally injective PWL mapping. The following works [Lipman 2012; Aigerman et al. 2014; Kovalsky et al. 2014] pose the bounded distortion mapping problem as a convexified constrained optimization. Recently, [Weber and Zorin 2014] provided an algorithm for mapping a mesh to a planar region with precise fixed boundary conditions that provides full guarantees that the produced mapping will be locally injective (in case such a mapping exists) at the price of a possible small amount of automatic refinement of the source mesh. [Levi and Zorin 2014] introduces the notion of strict minimizers to mapping computations based on the L_∞ norm. All the above mentioned algorithms produce PWL mappings which are nonsmooth by construction.

The algorithm of [Poranne and Lipman 2014] incorporates the convex optimization technique of [Lipman 2012] in a meshless setting and is able to produce injective bounded distortion mappings which are represented as linear combinations of basis functions. The authors provided the bounds analysis for various types of smooth basis functions, yet none of these functions are shape aware. As a result, they perform poorly on nonconvex shapes. As a remedy, the authors suggest replacing the Euclidean distance with the geodesic one which better captures the concavities of the shape. Albeit, smoothness of the mapping is lost and no analysis is provided for the bounds in this case. Instead, we use harmonic shape aware basis functions and provide a very tight analysis for the distortion bounds. Moreover, the entire analysis is restricted to the boundary of the domain, leading to much more efficient and accurate results.

3 Background

In this section we provide a short introduction to complex analysis and to the mathematical machinery of mappings in the plane. For further reading we refer to the book by Duren [2004].

A real-valued function $u(x, y)$ is harmonic if it satisfies the Laplace equation

$$\Delta u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0.$$

A mapping $[u(x, y), v(x, y)]$ from a region in the xy-plane to a region in the uv-plane is harmonic if both u and v are harmonic. Using complex notations $z = x + \mathrm{i}y$ we have $f(z) = u(z) + \mathrm{i}v(z)$. Thus a complex-valued harmonic function $f : \Omega \subset \mathbb{C} \to \mathbb{C}$ can be interpreted as a harmonic planar mapping. A complex-valued function $f = u + \mathrm{i}v$ is holomorphic (complex analytic) in a domain $\Omega \subset \mathbb{C}$ if it satisfies the Cauchy-Riemann equations at every point $z \in \Omega$

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}, \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}.$$

Holomorphic functions are the central objects of study in complex analysis. They are harmonic complex-valued functions (the converse is not true in general) and possess many strong and useful mathematical properties. For example, the sums, products and compositions of holomorphic functions are also holomorphic, and the quotient of two holomorphic functions is holomorphic wherever the denominator does not vanish. Holomorphic functions are differentiable and integrable (on a simply connected domain) infinite number of times and their derivative and anti-derivative are also holomorphic. An anti-holomorphic function is defined as the complex conjugate of a holomorphic function. For an arbitrary (not necessarily holomorphic) differentiable function $f(z)$, the complex derivatives are defined as $f_z = \frac{1}{2}(f_x - \mathrm{i}f_y)$ and $f_{\bar{z}} = \frac{1}{2}(f_x + \mathrm{i}f_y)$. A direct consequence of the Cauchy-Riemann equations is that holomorphic functions satisfy $f_{\bar{z}}(z) = 0$, while anti-holomorphic functions satisfy $f_z(z) = 0$. Hence, it is often convenient to denote the first complex derivative of a holomorphic function simply by $f'(z) = f_z(z)$.

In a simply connected domain Ω, any complex-valued harmonic function f has a decomposition

$$f(z) = \Phi(z) + \overline{\Psi}(z), \tag{1}$$

where Φ and Ψ are holomorphic functions (see [Duren 2004] Section 1.2 for a proof). The above representation is unique up to an additive complex constant that can be chosen by setting e.g. $\Psi(z_0) = 0$ for an arbitrary point $z_0 \in \Omega$. This leads to a simple yet powerful observation that will be used extensively throughout the paper.

Corollary 1. *A complex-valued function f (not necessarily holomorphic) defined on a simply connected domain is harmonic if and only if the complex derivatives f_z and $f_{\bar{z}}$ are holomorphic and anti-holomorphic respectively.*

Proof. Since f is harmonic and the domain is simply connected, we have $f = \Phi + \overline{\Psi}$. By differentiating f with respect to z and

$\bar{z}$ we have $f_z = \Phi'$ and $f_{\bar{z}} = \overline{\Psi'}$. The result immediately follows from the fact that derivatives of holomorphic functions are also holomorphic. On the opposite direction, we have that f_z and $\overline{f_{\bar{z}}}$ are holomorphic. Since the domain is simply connected f_z and $\overline{f_{\bar{z}}}$ have holomorphic antiderivatives denoted as Φ and Ψ respectively. Finally, $\overline{\Psi}$ and Φ are both harmonic and so is their sum. □

A continuously differentiable mapping f (not necessarily holomorphic) is locally injective sense-preserving at a vicinity of a point x if $det(J_f(x)) > 0$, where J_f is the mapping's Jacobian matrix. For such a mapping f, the determinant of the Jacobian is given by

$$det(J_f) = |f_z|^2 - |f_{\bar{z}}|^2 .$$

As a consequence, f is locally injective sense-preserving mapping wherever

$$|f_z| > |f_{\bar{z}}| . \tag{2}$$

Note that it is essential for the inequality to be strict. To see that, consider the following simple example of the holomorphic function $g(z) = z^2$ defined on the unit disk, mapping the unit circle to itself with double covering. Since g is holomorphic, $g_{\bar{z}} = 0$ and $g_z = g' = 2z$, and we can write $|g_z| = 2|z| \geq 0 = |g_{\bar{z}}|$. Equation (2) holds for every point $z \neq 0$ inside the disk, but g fails to be locally injective at $z = 0$ where the (strict) inequality does not hold. It is easy to see from Equation (2) that a necessary condition for local injectivity is that $|f_z| > 0$ (also written $f_z \neq 0$). For the special case of a holomorphic function g, we have that g is locally injective in a domain Ω, if and only if its first complex derivative does not vanish at every point in the domain: $g'(z) \neq 0, \forall z \in \Omega$. In that case, g is called a conformal mapping, that is, a mapping that preserves angles between any two intersecting curves.

It is often desired to measure and control the amount of angle and metric distortion induced by a mapping. Most distortion measures can be formulated in terms of the smallest and largest singular values of the Jacobian matrix $0 \leq \sigma_2 \leq \sigma_1$. The singular values of a continuously differentiable complex-valued mapping f can be expressed by

$$\sigma_1 = |f_z| + |f_{\bar{z}}| , \quad \sigma_2 = \Big| \, |f_z| - |f_{\bar{z}}| \, \Big| . \tag{3}$$

If f is locally injective sense-preserving, the latter expression simplifies to $\sigma_2 = |f_z| - |f_{\bar{z}}|$ and clearly $\sigma_2 > 0$. The quantity $\mu_f = \frac{f_{\bar{z}}}{f_z}$ is called the *first complex dilatation* (a.k.a. the complex Beltrami coefficient) and its modulus $k = \frac{|f_{\bar{z}}|}{|f_z|}$ is the *little dilatation*. The little dilatation k is related to the ratio K of singular values of the Jacobian of f (*large dilatation*) by $K = \frac{\sigma_1}{\sigma_2} = \frac{1+k}{1-k}$. Both k and K measure the amount of conformal distortion. For locally injective sense-preserving mappings $0 \leq k(z) < 1$ and $k(z) = 0$ iff f is conformal. For harmonic complex-valued functions it is often more convenient to work with the *second complex dilatation* $\nu_f = \frac{\overline{f_{\bar{z}}}}{f_z}$ since if f_z does not vanish, ν_f is holomorphic. This follows directly by applying Corollary 1 and using the fact that the quotient of two holomorphic functions is holomorphic. Finally, we note that the first and second complex dilatation are related by $k = |\nu| = |\mu|$.

A continuous real-valued function ξ of a complex variable z defined on a domain $\Omega \subset \mathbb{C}$ is subharmonic if and only if for any closed disc $D(z, r) \subset \Omega$ centered at z with radius r, ξ satisfies the *mean value inequality*

$$\xi(z) \leq \frac{1}{2\pi} \int_0^{2\pi} \xi(z + re^{i\theta}) \, d\theta.$$

Intuitively, a subharmonic function is at any point not greater than the average of the values in a circle around it. This leads to a maximum principle for subharmonic functions. Namely, the maximum of a subharmonic function is always attained at the boundary of the domain (the minimum principle does not hold in general). Subharmonic functions form a convex cone, that is, linear combinations of subharmonic functions with positive coefficients is also subharmonic (in particular a sum of subharmonic functions is subharmonic). A useful property of holomorphic functions is that their modulus is subharmonic.

4 Bounded Distortion Mappings

The main contribution of this paper is the development of an underlying mathematical theory and algorithms for computing locally injective harmonic and conformal planar mappings which are guaranteed to have bounded amount of distortion.

Definition 2. *A continuously differentiable planar mapping $f : \Omega \subset \mathbb{C} \to \mathbb{C}$ is a (k, σ_1, σ_2) bounded distortion mapping if it satisfies the following conditions*

$$0 \leq k(z) \leq k < 1 \qquad \forall z \in \Omega, \tag{4a}$$
$$\sigma_1(z) \leq \sigma_1 < \infty \qquad \forall z \in \Omega, \tag{4b}$$
$$0 < \sigma_2 \leq \sigma_2(z) \qquad \forall z \in \Omega, \tag{4c}$$

where $k(z)$ is the dilatation of f, $\sigma_1(z), \sigma_2(z)$ are the singular values of the Jacobian of f, and k, σ_1, σ_2 are real constants.

Equation (4a) asserts that f has bounded amount of conformal distortion (f is called a quasiconformal mapping [Ahlfors 1966]). Equations (4b) and (4c) bounds the maximal and minimal amount of local stretch from above and below respectively. Together they can be used to bound various measures of isometric distortion such as $\max\{\sigma_1(z), 1/\sigma_2(z)\}$.

Observation 3. *A (k, σ_1, σ_2) bounded distortion mapping f is locally injective sense-preserving.*

Proof. From Equation (4c) we have that $\sigma_2(z) > 0$, hence $|f_z| \neq |f_{\bar{z}}|$, so at every point in the domain f is either sense-preserving ($|f_z| > |f_{\bar{z}}|$) or sense-reversing ($|f_z| < |f_{\bar{z}}|$). (4a) rules out $|f_z| < |f_{\bar{z}}|$. □

The conditions in Definition 2 involve every point in the domain of f. A natural question to raise is whether it is possible to reformulate these conditions solely in terms of the boundary behavior of the mapping, in the special case that f is harmonic? The next theorem addresses this question and provides alternative necessary and sufficient conditions for a *harmonic* mapping to be a (k, σ_1, σ_2) bounded distortion mapping.

Theorem 4. *A complex-valued harmonic mapping f defined on a simply connected domain Ω is (k, σ_1, σ_2) bounded distortion if and only if*

$$\oint_{\partial\Omega} \frac{f_z'(z)}{f_z(z)} dz = 0, \tag{5a}$$
$$0 \leq k(w) \leq k < 1 \qquad \forall w \in \partial\Omega, \tag{5b}$$
$$\sigma_1(w) \leq \sigma_1 < \infty \qquad \forall w \in \partial\Omega, \tag{5c}$$
$$0 < \sigma_2 \leq \sigma_2(w) \qquad \forall w \in \partial\Omega, \tag{5d}$$

Proof. If f is harmonic (k, σ_1, σ_2) bounded distortion on a simply connected domain, f_z is holomorphic (Corollary 1). Observation 3 asserts that $|f_z| > |f_{\bar{z}}|$ which imply that f_z does not vanish inside

the domain ($f_z \neq 0$). By applying Cauchy's argument principle to f_z we get

$$\oint_{\partial\Omega} \frac{f_z'(z)}{f_z(z)} dz = 2\pi \mathrm{i} N, \tag{6}$$

where N is the number of zeros of f_z. Since $N = 0$, Equation (5a) holds. Equations (5b),(5c), (5d) are satisfied simply because they are the restriction of Equations (4a),(4b),(4c) from Definition 2 to the boundary of the domain. The other direction is more complicated to prove and will rely on the following three small lemmas.

Lemma 5. *Conditions* (5a) *and* (5b) *imply condition* (4a).

Proof. Condition (5a) implies that $f_z \neq 0$ and so the second complex dilatation $\nu_f = \frac{\overline{f_{\bar{z}}}}{f_z}$ is holomorphic (Section 3). Since the modulus of a holomorphic function is a subharmonic function, $|\nu_f| = k(z)$ attains its maximum on $\partial\Omega$. Hence, bounding $k(z)$ from above by a constant k on the boundary (condition (5b)) implies that the same bound holds at every point in the domain (condition (4a)). □

Lemma 6. *Condition* (5c) *imply condition* (4b).

Proof. Using Corollary 1, we know that since f is harmonic, f_z and $\overline{f_{\bar{z}}}$ are holomorphic. $|f_z|$ and $\left|\overline{f_{\bar{z}}}\right| = |f_{\bar{z}}|$ are subharmonic and so is their sum $\sigma_1(z) = |f_z| + |f_{\bar{z}}|$ (Equation (3) left). Since $\sigma_1(z)$ is subharmonic, its maximum is attained on the boundary and it is sufficient to bound it on the boundary from above. □

Lemma 7. *Conditions* (5a)*,* (5b)*, and* (5d) *together imply condition* (4c).

Proof. By using the same argument applied in the proof of Observation 3, we know that, on the boundary, $|f_z| > |f_{\bar{z}}|$. Hence, the expression for $\sigma_2(w)$ on the boundary simplifies to $|f_z(w)| - |f_{\bar{z}}(w)|$ (Equation (3) right). Next we define an auxiliary function $\varsigma(z) = \sigma_2 / |f_z(z)| + k(z)$, which is well defined on Ω since $f_z \neq 0$. From Lemma 5 we have that $k(z)$ is subharmonic. Since $f_z(z)$ is holomorphic, so is $1/f_z(z)$. It follows that $1/|f_z(z)|$ is subharmonic and since σ_2 is a positive constant, $\varsigma(z)$ is also subharmonic. Equation (5d) can be written as $\sigma_2 \leq |f_z(w)| - |f_{\bar{z}}(w)|$. Dividing both sides of the inequality by $|f_z(w)|$ and further manipulating we get

$$\frac{\sigma_2}{|f_z(w)|} + \frac{|f_{\bar{z}}(w)|}{|f_z(w)|} \leq 1, \tag{7}$$

where the left hand side is simply $\varsigma(w)$. The subharmonicity of ς is then used to show that $\varsigma(z)$ is bounded by 1 in the *entire* domain. By reversing the manipulations that led to Equation (7) we obtain $\sigma_2 \leq |f_z(z)| - |f_{\bar{z}}(z)| = \sigma_2(z)$ which concludes the proof of the lemma. □

Finally, since Lemmas 5, 6, and 7 show that conditions (4a), (4b), and (4c) are satisfied, the harmonic mapping f is (k, σ_1, σ_2) bounded distortion and the theorem is proved. □

Equipped with Theorem 4, we are now ready to design an algorithm for computing (k, σ_1, σ_2) bounded distortion harmonic mappings.

5 Optimization

In order to perform interactive shape deformation, we use the point-to-point (P2P) user metaphor. The user controls the deformation by manipulating the target position of a small amount of points located inside or on the boundary of the domain. In addition, the user prescribes the bounds $0 \leq k < 1, \quad 0 < \sigma_2 \leq \sigma_1$ and a numerical optimization problem is solved in order to find an optimal (k, σ_1, σ_2) bounded distortion harmonic mapping.

5.1 Discretization

To discretize the space of harmonic mappings, it will be convenient to represent f as $\Phi + \overline{\Psi}$, where Φ and Ψ are both holomorphic (Equation (1)). We represent holomorphic functions by using the Cauchy complex barycentric coordinates [Weber et al. 2009] which are derived by discretizing the boundary of the domain using a polygonal shape (a so-called cage) and computing the Cauchy transform [Bell 1992] of a piecewise linear trial function. Let $\hat{\mathrm{P}} = \{z_1, z_2, ..., z_n\}, \quad z_j \in \mathbb{C}$ be the vertices of a simply connected planar polygon (the cage), oriented counterclockwise (see Figure 2 for notations).

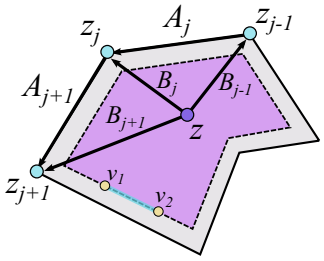

Figure 2: *Notations for Cauchy's coordinates. The purple area is the domain Ω bounded by a polygon* P. *The solid black polygon is an outward offset cage* $\hat{\mathrm{P}}$ *with vertices z_j oriented counterclockwise. z is any point in Ω. $A_j = z_j - z_{j-1}$, $B_j(z) = z_j - z$, and (v_1, v_2) is a segment on* P.

Using the discrete Cauchy transform [Weber et al. 2009], Φ and Ψ are represented as

$$\Phi(z) = \sum_{j=1}^{n} C_j(z)\varphi_j, \qquad \Psi(z) = \sum_{j=1}^{n} C_j(z)\psi_j, \tag{8}$$

where $C_j(z)$ is the j^{th} holomorphic Cauchy barycentric coordinate associated with vertex z_j and φ_j, ψ_j are complex coefficients. $C_j(z)$ and its first and second complex derivatives, $C_j'(z)$, $C_j''(z)$, possess a rather simple closed-form expressions (see Appendix A). For any choice of the coefficients φ_j, ψ_j, the mapping f is harmonic and its derivatives

$$\begin{aligned} f_z(z) &= \Phi'(z) = \sum_{j=1}^{n} C_j'(z)\varphi_j, \\ f_{\bar{z}}(z) &= \overline{\Psi'(z)} = \sum_{j=1}^{n} \overline{C_j'(z)\psi_j}, \end{aligned} \tag{9}$$

are holomorphic and anti-holomorphic respectively, and can be easily evaluated at any point z inside the domain using a straightforward formula.

5.2 Convexification

Harmonicity of f is a built-in property of the subspace we chose. Based on Theorem 4, we design a numerical optimization procedure that requires setting inequality constraints *only at the boundary* of the domain. However, both the conditions of Definition 2 and Theorem 4 pose a challenge for numerical optimization since they are nonconvex. To alleviate this difficulty, we adapt the effective convexification approach of [Lipman 2012] to the complex setting. This is done by substituting the *nonconvex* constraints of Theorem 4 with (maximal) *convex* constraints that imply the nonconvex ones. The convexified constraints are then enforced efficiently by using a Second Order Cone Program (SOCP) solver.

Let us start by expressing condition (5c) explicitly

$$|f_z(w)| + |f_{\bar{z}}(w)| \leq \sigma_1. \tag{10}$$

Fortunately, this constraint is already convex. The convexification of the rest of the constraints is done by introducing an auxiliary function θ. Let $\theta(w) : \partial\Omega \to \mathbb{R}$ be a continuous real-valued function defined on the boundary of the domain. For each selection of $\theta(w)$ a different maximal convex subspace of the full nonconvex space is chosen. By maximality of the subspace we mean that there is no other convex subspace that strictly contains it. For now, assume that θ is chosen arbitrarily. The geometric interpretation of θ and a strategy on how to choose it is given in Section 5.3.

Condition (5d) is expressed explicitly by

$$\sigma_2 \leq |f_z(w)| - |f_{\bar{z}}(w)|, \tag{11}$$

and is replaced by the following second order convex cone constraint

$$|f_{\bar{z}}(w)| \leq \mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right) - \sigma_2. \tag{12}$$

The above constraint defines a maximal convex subset that is contained in the nonconvex set defined by Equation (11). The fact that Equation (12) implies Equation (11) follows from

$$\mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right) \leq \left|f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right| = |f_z(w)|, \tag{13}$$

where the left inequality is due to the fact that the modulus of any complex number is no less than its real (and imaginary) part. The maximality follows by using the same arguments used in [Lipman 2012] or [Poranne and Lipman 2014] and is omitted here for brevity.

Next, we handle condition (5b) and rearrange it to take the following explicit form

$$|f_{\bar{z}}(w)| \leq k\,|f_z(w)|. \tag{14}$$

The second order convex cone constraint substitute for Equation (14) is given by

$$|f_{\bar{z}}(w)| \leq k\mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right), \tag{15}$$

which also defines a maximal convex subset. Containment is easily proved by multiplying both sides of the inequality (13) by the (nonnegative) constant k such that we have

$$|f_{\bar{z}}(w)| \leq k\mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right) \leq k\,|f_z(w)|.$$

As a regularization term we use the ARAP energy [Liu et al. 2008]

$$\mathrm{E}_{\mathrm{ARAP}} = \frac{1}{2}\int \left((\sigma_1 - 1)^2 + (\sigma_2 - 1)^2\right) da. \tag{16}$$

A simple calculation shows that if f is locally injective sense-preserving, the ARAP energy can be expressed by

$$\mathrm{E}_{\mathrm{ARAP}} = \int \left((|f_z| - 1)^2 + |f_{\bar{z}}|^2\right) da. \tag{17}$$

As this energy is nonconvex, we approximate it using the following quadratic functional in the spirit of [Poranne and Lipman 2014]

$$\int \left(\left|f_z\mathrm{e}^{\mathrm{i}\theta} - 1\right|^2 + |f_{\bar{z}}|^2\right) da. \tag{18}$$

It remains to deal with the nonlinear condition of Equation (5a). Without this condition, Theorem 4 cannot be applied. This is explained in detail in Section 6.4.

5.3 The Algorithm

In order to deform a domain Ω, bounded by a polygonal shape P, we compute an outward offset polygon, denoted as the cage $\hat{\mathrm{P}}$. This is done in order to avoid the singularities of the Cauchy's coordinates on the boundary. The barycentric coordinates and their derivatives are computed with respect to $\hat{\mathrm{P}}$ but are evaluated only on or inside P (during a preprocessing step). We typically use an offset of 0.1% of the overall length of P.

We then sample the boundary P uniformly with various densities to obtain the sets $\mathcal{M}$, $\mathcal{A}$, and $\mathcal{B}$. The set $\mathcal{M}$ is used in order to approximate the ARAP energy. $\mathcal{A}$ is a set of samples where (potentially) the convex constraints are being enforced. $\mathcal{B}$ is used in the process of evaluating the global bounds on distortion (Section 6). Finally, the user defines the set $\mathcal{P}$ that contains the deformation handles, by selecting small amount of points, on or inside P.

The algorithm also maintains three active sets: $\mathcal{A}^k$, $\mathcal{A}^{\sigma_1}$, and $\mathcal{A}^{\sigma_2}$ that are initialized with the vertices of P. θ is initialized to 0. The user sets the parameters σ_1, σ_2, k that bounds the distortion and the interaction begins. In order to deform the shape, the user relocates the points in $\mathcal{P}$ from their original positions r_j to their target positions q_j and the following SOCP is solved

$$\begin{aligned}
\min_{\varphi,\psi}\quad & \mathrm{E}_{\mathrm{ARAP}} + \lambda\,\mathrm{E}_{\mathrm{P2P}} \\
s.t.\quad & \psi_1 = 0, \\
& \forall p \in \mathcal{A}^k \quad |f_{\bar{z}}(p)| \leq k\mathrm{Re}\left(f_z(p)\mathrm{e}^{\mathrm{i}\theta(p)}\right), \\
& \forall p \in \mathcal{A}^{\sigma_1} \quad |f_z(p)| + |f_{\bar{z}}(p)| \leq \sigma_1, \\
& \forall p \in \mathcal{A}^{\sigma_2} \quad |f_{\bar{z}}(p)| \leq \mathrm{Re}\left(f_z(p)\mathrm{e}^{\mathrm{i}\theta(p)}\right) - \sigma_2,
\end{aligned} \tag{19}$$

where the equality constraint is needed in order to nail down the constant degree of freedom of the representation $f = \Phi + \overline{\Psi}$. $\mathrm{E}_{\mathrm{ARAP}}$ is obtained by approximating the integral in Equation (18) with the following quadratic function

$$\mathrm{E}_{\mathrm{ARAP}} = \frac{1}{|\mathcal{M}|}\sum_{j=1}^{|\mathcal{M}|}\left(\left|f_z(p_j)\mathrm{e}^{\mathrm{i}\theta(p_j)} - 1\right|^2 + |f_{\bar{z}}(p_j)|^2\right), \tag{20}$$

where the samples $p_j \in \mathcal{M}$ are taken on the boundary alone. Adding samples to the interior is also possible but we have not noticed any improvement in the results. Positional constraints are enforced softly using the following energy term

$$\mathrm{E}_{\mathrm{P2P}} = \sum_{j=1}^{|\mathcal{P}|} |f(r_j) - q_j|^2. \tag{21}$$

Switching to hard positional constraints is easily done by replacing $\mathrm{E}_{\mathrm{P2P}}$ with linear equality constraints $f(r_j) = q_j$, but is less preferable as it may hamper the feasibility of the constrained optimization.

While the convex optimization in equation (19) is always feasible, the mapping it produces is only guaranteed to have bounded amount of distortion at a finite number of points (where the inequality constraints are enforced). Hence, the obtained solution must be validated by computing *global* distortion bounds (the validation process is explained in detail in Section 6). Once the solution is validated, we update the function $\theta(w)$ (this defines a different convex subspace) and the convex optimization is re-solved. This is repeated till the energy cannot be further reduced. Convergence typically requires 1-3 iterations. In each iteration, $\theta(w)$ is set to $-\arg\tilde{f}_z(w)$ where $\tilde{f}_z$ is our best estimate for the unknown f_z. We then have

$e^{i\theta} = |\tilde{f}_z|/\tilde{f}_z$ and the expression $\mathrm{Re}\left(f_z e^{i\theta}\right)$, that repetitively appears in our convexified constraints, becomes a good approximation to the expression $|f_z|$. This is desirable since the only difference between the full nonconvex constraints (Equations (11) and (14)) and the convex ones (Equations (12) and (15)) is the substitution of $\mathrm{Re}\left(f_z e^{i\theta}\right)$ with $|f_z|$. In practice, we start with $\tilde{f}_z = 0$ (corresponds to the identity mapping) and update this estimate in further iterations by using the solution f_z obtained in the previous iteration. This also ensures that the optimization at each iteration is feasible.

In order to reduce the computation time of the optimization (19), we use an active set approach where the main idea is to add inequality constraints only where they are actually needed. In contrast to [Poranne and Lipman 2014] that uses a single active set, we noticed that it is quite rare for more than one type of constraint to be violated simultaneously and that maintaining three active sets leads to significant reduction in the total number of active constraints. Before a new iteration begins, we evaluate $k(p)$, $\sigma_1(p)$, and $\sigma_2(p)$ on all points $p \in \mathcal{A}$ and add any point that violates the user bounds k, σ_1, and σ_2 to its corresponding active set $\mathcal{A}^k$, $\mathcal{A}^{\sigma_1}$, and $\mathcal{A}^{\sigma_2}$. In addition, we find the *local* extrema points (maxima for $k(p)$ and $\sigma_1(p)$, and minima for σ_2(p)) and add points which are relatively close to violate the bounds. For example, if $k(p_i)$ is a local maxima (i.e. $k(p_i) > k(p_{i-1})$ and $k(p_i) > k(p_{i+1})$) and $k(p_i) > 0.95k$, we add p_i to the active set $\mathcal{A}^k$ (and similarly for $\sigma_1(p)$). If $\sigma_2(p_i)$ is a local minima and $\sigma_2(p_i) < 1.15\sigma_2$, we add p_i to the active set $\mathcal{A}^{\sigma_2}$. On the other hand, any point p in the active sets for which the distortion goes sufficiently low is removed (we use thresholds of $0.945k$, $0.945\sigma_1$, and $1.2\sigma_2$).

5.4 Conformal Mappings

Since any conformal mapping is a harmonic one, we can easily specialize our optimization to the conformal case. For conformal mappings $f_{\bar{z}} = 0$, therefore we can remove Ψ of Equation (8) and eliminate the variables ψ. The expression for the singular values simplies to $\sigma_1(z) = \sigma_2(z) = |f_z(z)|$ and the convex constraints are given by

$$\begin{aligned} |f_z(w)| &\le \sigma_1, \\ \mathrm{Re}\left(f_z(w) e^{i\theta(w)}\right) &\ge \sigma_2, \end{aligned} \tag{22}$$

where the upper bound on $k(w)$ is not needed. The quadratic ARAP energy (Equation (20)) also becomes simpler by eliminating the term $|f_{\bar{z}}|^2$.

6 Global Bounds

In the previous section, a numerical optimization procedure was developed for bounding the distortion of a harmonic mapping at a finite number of boundary points. Since we are dealing with smooth mappings rather than discrete (PWL) mappings, we would like the distortion to be bounded at *every* point in our domain. This is challenging since clearly there are infinite number of points in the domain. The strategy we develop to address this challenge is based on some recent useful ideas by [Poranne and Lipman 2014], where the notion of modulus of continuity was successfully employed to derive global distortion bounds of planar mappings. In contrast to [Poranne and Lipman 2014] we provide a much tighter analysis for bounding the distortion. Moreover, the entire analysis is carried on the boundary alone, leading to a dramatic improvement in computational efficiency and accuracy. Given a harmonic mapping $f = \Phi + \overline{\Psi}$, (where Φ and Ψ are discretized as in Equation (8)), our goal in this section is to answer the following question: Is f a (k, σ_1, σ_2) bounded distortion mapping?

The solution to the optimization problem of Section 5 provides the optimal values of the variables φ and ψ. Given these, we compute global bounds on the distortion, providing a certificate for the mapping to be $(\tilde{k}, \tilde{\sigma}_1, \tilde{\sigma}_2)$ bounded distortion (and in particular locally injective). For obvious reasons, the obtained bounds $\tilde{k}$, $\tilde{\sigma}_1$, $\tilde{\sigma}_2$ will be slightly worse than the user specified bounds k, σ_1, σ_2 (which are enforced only at a finite number of points). We allow the user to set the maximal deviation from the user specified bounds and in case these are exceeded, we reject the result and perform a simple line search to obtain a valid result. More specifically, let $\varphi^{i-1}, \psi^{i-1}$ be the (valid) solution from the previous iteration and φ^i, ψ^i the solution to the optimization of the current iteration that we would like to validate. Starting with $t = 1$, we obtain

$$\begin{aligned} \varphi(t) &= (1-t)\varphi^{i-1} + t\varphi^i, \\ \psi(t) &= (1-t)\psi^{i-1} + t\psi^i, \end{aligned}$$

and compute the global distortion bounds corresponding to $\varphi(t)$ and $\psi(t)$. If the user threshold for the bounds is violated, we reduce t to $t/2$ and recompute the bounds until a bounded distortion mapping is obtained or a maximal number of steps is performed. We used default thresholds of $1.2k$, $1.3\sigma_1$, and $0.7\sigma_2$ and 10 bisection steps in all our experiments. In some scenarios a solution with specified distortion bounds that satisfies the positional constraints simply does not exist. Yet, the more common scenario for which the line search fails to provide sufficient progress is due to insufficient amount of active constraints. When this happens the active set approach automatically kicks in, adding more constraints to the optimization in subsequent iterations, which eventually leads to fulfillment of the bounds.

6.1 Bounding the Distortion

We now explain how to compute the global distortion bounds for the mapping f. To this end, we use the set $\mathcal{B}$ of samples computed in the preprocessing step and treat each two consecutive samples as a segment. Our key idea for obtaining tight bounds is to bound the distortion separately and independently on each segment. From this point on, the entire analysis is carried out on a *single* segment. Once a separate bound is obtained on each segment, the global bounds are computed as the maximum/minimum over all segments.

Denote by $|f_z|_{\min}$ and $|f_z|_{\max}$ the lower and upper bounds on $|f_z|$ and let $|f_{\bar{z}}|_{\max}$ be the upper bound on $|f_{\bar{z}}|$. On each segment, the bounds of the functions $k(w)$, $\sigma_1(w)$, and $\sigma_2(w)$ can be formulated in terms of $|f_z|_{\min}$, $|f_z|_{\max}$, and $|f_{\bar{z}}|_{\max}$ as follows

$$\begin{aligned} \sigma_1(w) &= |f_z(w)| + |f_{\bar{z}}(w)| \;\le\; |f_z|_{\max} + |f_{\bar{z}}|_{\max}, \\ \sigma_2(w) &= |f_z(w)| - |f_{\bar{z}}(w)| \;\ge\; |f_z|_{\min} - |f_{\bar{z}}|_{\max}, \\ k(w) &= \frac{|f_{\bar{z}}(w)|}{|f_z(w)|} \;\le\; \frac{|f_{\bar{z}}|_{\max}}{|f_z|_{\min}}. \end{aligned} \tag{23}$$

Next, we explain how to compute $|f_z|_{\min}$, $|f_z|_{\max}$, and $|f_{\bar{z}}|_{\max}$.

6.2 Lipschitz Continuous Complex Functions

Definition 8. *A complex-valued function $g : \Omega \subset \mathbb{C} \to \mathbb{C}$ is called L-Lipschitz continuous if there exists a real constant $L \ge 0$ such that, for all z and w in Ω*

$$|g(z) - g(w)| \;\le\; L\,|z - w|\,.$$

Any such L is referred to as a Lipschitz constant (or modulus of uniform continuity) of g. Note that the above definition also includes real-valued functions (of a complex variable) as a special

case since any real-valued function is also complex-valued (with zero imaginary part). Intuitively, the Lipschitz constant limits the rate in which the function can change.

Our goal is to compute the Lipschitz constants of the real-valued functions $|f_z|$ and $|f_{\bar{z}}|$. However it turns out that computing the Lipschitz constants of the complex-valued functions f_z and $f_{\bar{z}}$ is much easier. For that the following proposition becomes handy (see Appendix B for a proof).

Proposition 9. *If g is L-Lipschitz continuous then $|g|$, $\mathrm{Re}\,(g)$, $\mathrm{Im}\,(g)$, and $\overline{g}$ are also L-Lipschitz continuous.*

Given a segment (v_1, v_2) and a real-valued L-Lipschitz continuous function g, we can bound the restriction of g to the segment from above and from below by

$$\frac{(g^1 + g^2 - Ll)}{2} \leq g(t) \leq \frac{(g^1 + g^2 + Ll)}{2} \quad \forall t \in [0,1], \quad (24)$$

where $l = |v_1 - v_2|$ is the segment length, $g(t) = g((1-t)v_1 + tv_2)$ is the value of g at a point on the segment parameterized by t, and $g^1 = g(0)$, $g^2 = g(1)$ are the function values at the endpoints.

Based on Equation (24) we obtain the following expressions which are then used to bound the distortion via Equation (23).

$$|f_z|_{\min} = \frac{\left|f_z^1\right| + \left|f_z^2\right| - L_{f_z} l}{2}, \quad (25)$$

$$|f_z|_{\max} = \frac{\left|f_z^1\right| + \left|f_z^2\right| + L_{f_z} l}{2}, \quad (26)$$

$$|f_{\bar{z}}|_{\max} = \frac{\left|f_{\bar{z}}^1\right| + \left|f_{\bar{z}}^2\right| + L_{f_{\bar{z}}} l}{2}, \quad (27)$$

where the superscript indicates the function values at the segment endpoints and L_{f_z} and $L_{f_{\bar{z}}}$ are the Lipschitz constants of f_z and $f_{\bar{z}}$ respectively on the segment.

6.3 Computing the Lipschitz Constants

We now explain how to compute the Lipschitz constants L_{f_z} and $L_{f_{\bar{z}}}$. A property of Lipschitz continuous complex functions is that if g and q are L_g and L_q Lipschitz continuous respectively then any linear combination $ag + bq$ with complex coefficients a and b is Lipschitz continuous with constant $|a|\, L_g + |b|\, L_q$.

Therefore, based on Equation (9), valid Lipschitz constants (not necessarily the smallest) can be computed using

$$L_{f_z} = \sum_{j=1}^{n} L_{C_j'} \, |\varphi_j|, \quad (28)$$

$$L_{f_{\bar{z}}} = L_{\overline{f_{\bar{z}}}} = \sum_{j=1}^{n} L_{C_j'} \, |\psi_j|, \quad (29)$$

where the only missing piece is $L_{C_j'}$, which is the Lipschitz constant of the derivative of the j^{th} Cauchy coordinate (associated with vertex z_j).

Proposition 10. *The function $C_j'(z)$ is Lipschitz at any point in the domain excluding the cage vertices and its Lipschitz constant is given by the following formula*

$$L_{C_j'} = \frac{|z_{j+1} - z_{j-1}|}{2\pi \, d(z_{j-1}) d(z_j) d(z_{j+1})}, \quad (30)$$

where z_{j-1} and z_{j+1} are the vertices adjacent to z_j on the cage $\hat{\mathrm{P}}$ and $d(z_j)$ is a function that computes the distance between the segment (v_1, v_2) and the point z_j. See Figure 2 for notations.

We prove this proposition in Appendix C.

During the preprocessing step, we compute a $|\mathcal{B}| \times n$ matrix of Lipschitz constants, where $|\mathcal{B}|$ is the (typically large) number of segments on which we want to compute the bounds and n is the number of vertices in the cage. Then at runtime, as φ and ψ change constantly, we need to evaluate Equations (28) and (29).

As can be seen from Equation (24), the tightness of the bounds depends on the magnitude of L and l. Reducing l is achieved by increasing the number of samples in $|\mathcal{B}|$ which decreases the length of each segment. However, it turns out that we can also obtain smaller L_{f_z} and $L_{f_{\bar{z}}}$ than those obtained by Equations (28) and (29). For simplicity, we only describe the strategy to reduce the magnitude of L_{f_z}. Reducing $L_{f_{\bar{z}}}$ is done similarly. The idea is to substitute the term $|\varphi_j|$ in Equation (28) by the more general term

$$|\varphi_j + az_j + b|, \quad (31)$$

where a and b are complex constants. To see why this substitution makes sense, let us first express f_z in a slightly different way

$$f_z(z) = \sum_{j=1}^{n} C_j'(z)(\varphi_j + b) = -a + \sum_{j=1}^{n} C_j'(z)(\varphi_j + b + az_j),$$

which is possible since $C_j(z)$ are (complex) barycentric coordinates. The first equality is due to $\sum C_j'(z) = 0$ (constant precision) and the second one is due to $\sum C_j'(z) z_j = 1$ (linear precision). Then the Lipschitz constant is given by

$$L_{f_z} = \sum_{j=1}^{n} L_{C_j'} \, |\varphi_j + b + az_j|, \quad (32)$$

where the Lipschitz constant of $-a$ is zero hence omitted.

It is important to realize that the constants a and b can be chosen differently on each segment. Ideally, for each segment we should choose a and b that minimize L_{f_z} in Equation (32). This boils down to solving a linear program with two complex variables. However, solving $|\mathcal{B}|$ linear programs interactively is computationally challenging. Instead, we suggest a heuristic for approximating the optimal solution, that turned out to be fast and effective in practice. For each segment, we find the largest element $L_{C_r'}$ in $\{L_{C_j'}\}_{j=1}^n$. We choose another index q to be $r+1$ if $L_{C_{r-1}'} < L_{C_{r+1}'}$. Otherwise we set $q = r-1$. Finally a and b are chosen to satisfy the following two linear equations

$$\begin{aligned} \varphi_r + b + az_r &= 0, \\ \varphi_q + b + az_q &= 0, \end{aligned} \quad (33)$$

which guarantee that the largest element in $\{L_{C_j'}\}_{j=1}^n$ and its largest neighbor are contributing nothing to the sum of Equation (32). Figure 3 shows a comparison of the obtained Lipschitz constants on a typical model with $n = 78$ cage vertices and $|\mathcal{B}| = 340$ boundary segments on which we evaluate L.

6.4 Nonvanishing Derivative

The bounds that were computed in the previous section were obtained solely on the boundary of the domain. To certify the harmonic mapping f as bounded distortion mapping using Theorem 4, we also need to satisfy condition (5a). As explained in Section 4, the argument principle can be used to show that Equation (5a) holds iff f_z does not vanish at every point inside the domain (note that $f_z \neq 0$ on the boundary does not necessarily mean that $f_z \neq 0$ inside). Using our specific discretization, the integrand that appears

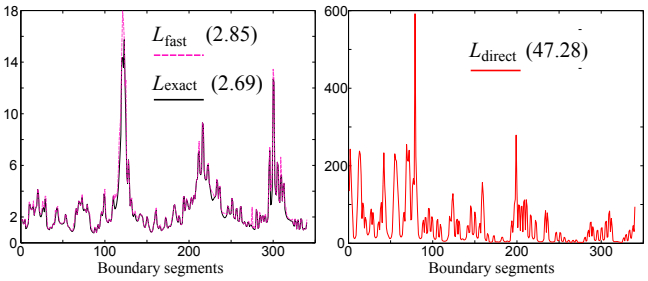

Figure 3: *Comparison of Lipschitz constants for f_z on 340 boundary segments. L_{fast} is our fast approximation procedure which took 2.1 milliseconds to compute. L_{exact} is the optimal solution obtained by solving 340 linear programs which took 25 seconds. L_{direct} is the naïve direct evaluation of Equation* (28) *which is plotted separately due to the large scale differences. Note how L_{fast} approximates closely L_{exact} and the improvement over L_{direct}. The mean of L for each method is denoted in parenthesis.*

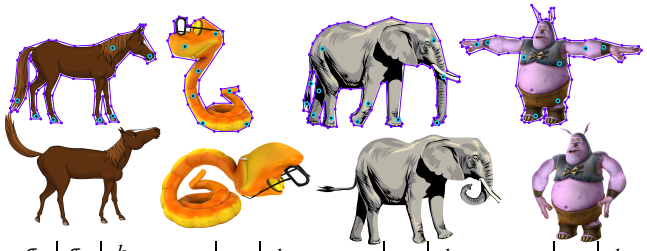

σ_2	σ_1	k
0.65	1.65	0.12
0.64	1.75	0.13

σ_2	σ_1	k
0.60	1.50	0.20
0.59	1.51	0.20

σ_2	σ_1	k
0.50	1.60	0.25
0.40	1.70	0.27

σ_2	σ_1	k
0.60	1.50	0.10
0.58	1.52	0.11

Figure 4: *Distortion bounds of harmonic mappings. The tables show the user bounds (top row) and the global bounds (bottom row) obtained using $|\mathcal{B}| = 15,000$ segments (which took approximately 25ms to compute). Cages are visualized as purple polygons.*

in Equation (5a) has an explicit simple formula. However, we were unable to obtain a closed-form expression for its antiderivative. One simple, direct and accurate way to obtain a sharp answer to whether f_z vanishes is to evaluate the integral numerically. However, we have developed an alternative which is several orders of magnitude faster to evaluate. For that, we will need to rely on the following theorem (proved in Appendix D).

Theorem 11. *Let f be a complex-valued harmonic function defined on a simply connected domain Ω. Let $\theta(w)$ be any real-valued continuous function defined on the boundary. We denote by γ the function*

$$\gamma(w) = \mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right).$$

If $\gamma(w) > 0$ at every point on the boundary *then f_z does not vanish* inside *the domain.*

In Appendix E we derive the Lipschitz constant L_γ and show how to use it in order to formulate the following sufficient condition for the positivity of γ on a specific segment

$$(2 + |\theta^2 - \theta^1|)\, l\, L_{f_z} < (2 - |\theta^2 - \theta^1|)\Big(\left|f_z^1\right| + \left|f_z^2\right|\Big), \quad (34)$$

where l is the length of the segment, $\left|f_z^1\right|, \left|f_z^2\right|$ are the values of $|f_z|$ at the endpoints, and θ^1, θ^2 are the values of θ at the endpoints. Satisfying this condition on all the segments in $\mathcal{B}$ guarantees that $f_z \neq 0$ throughout the entire domain. θ is chosen to be piecewise linear, satisfying $\theta(w^i) = -\arg f_z(w)\big|_{w^i}$ at the samples $w^i \in \mathcal{B}$.

Note that the argument function is multivalued and that off-the-shelf software implementations are designed to return the principal branch which may lead to discontinuities in θ. To alleviate that, we first compute the change in θ between each two consecutive boundary samples by using

$$d\theta^i = Arg(f_z^i / f_z^{i-1}).$$

Arg is the principal branch of arg. This avoids branching problems since f_z^i / f_z^{i-1} is typically very close to 1. We then compute the cumulative sum of all the differences $d\theta^i$ to obtain $\theta(w^i)$.

Figures 1 and 4 show harmonic deformations that were obtained using our algorithm. The global bounds were obtained by applying the procedure explained in this section using $15,000$ boundary segments. Evaluation of the distortion bounds took approximately 25 milliseconds.

7 Results

We have implemented our algorithm as a plugin to Autodesk Maya. The user interface and final OpenGL rendering are done in C++ and the algorithm itself is implemented in Matlab (through Matlab's engine). We used the interior point method implemented in the Mosek software to solve the SOCP. The results were obtained on an Intel i7-3770 machine with 16GB memory and an Nvidia Quadro K6000 graphics card.

The typical running time for a single solve with 100 basis functions and 100 active constraints is around 100ms which allows interactive editing even for our most complicated models (e.g. the Dragon). Less complicated models such as the Bar, Pants, and Snake are processed in approximately 25ms including rendering for which we used a textured mesh with 40K triangles. We observed that the running times of the Mosek SOCP solver roughly behave linearly on the number of basis functions and the number of active constraints.

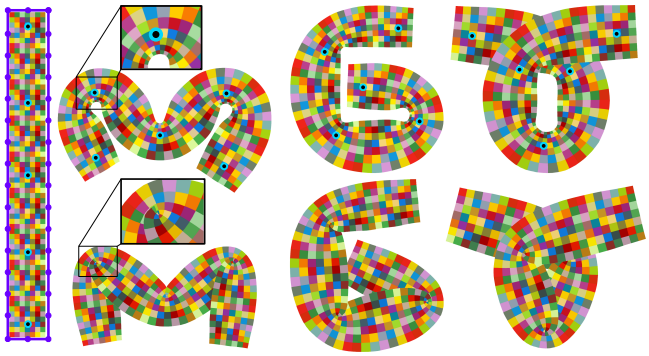

Figure 5: *Deformations of the Bar with five positional constraints. Comparison of our C^∞ certified conformal mappings (top row) with P2P-Cauchy coordinates - the unconstrained optimization of [Weber et al. 2009] (bottom row). Positional constraints are precisely satisfied. Note the singularities that P2P-Cauchy develops.*

Parameters selection We used the following default parameters for all our experiments. The number of energy samples $|\mathcal{M}|$ was set to $10,000$. For the active set approach we used $|\mathcal{A}| = 2000$ samples from which constraints can become active. The distortion was evaluated on $|\mathcal{B}| = 15,000$ segments and the default outward offset was chosen to be 0.1% of the overall length of P. We balance the two energy terms ($\mathrm{E_{ARAP}}$ and $\mathrm{E_{P2P}}$) using $\lambda = 100$. No special need to further tweak this parameter was required.

While it is not guaranteed that the positional constraints will always be satisfied, our algorithm is quite effective in doing so. We visual-

ize the user prescribed P2P target positions as cyan disks while the actual deformed points are visualized as smaller black dots. Black dot which is centered inside a cyan disk indicates that the positional constraint is being satisfied. As evident, all the results presented in the paper, as well as in the accompanying video, precisely satisfy the user constraints. Moreover, Table 1 provides the exact value of the E_{P2P} energy for the presented results.

	Horse	Snake	Elephant	Troll	Bar	Giraffe	Pants	Dragon
λE_{p2p}	3.59E-05	1.72E-06	1.35E-06	7.31E-05	2.95E-06	7.28E-06	3.08E-05	5.50E-07
n	64	34	75	69	34	42	56	139

Table 1: *Soft positional constraints. The table states the value of the positional constraints energy for all results.* λ *is set to* 100. n *is the number of cage vertices.*

Figure 5 shows a comparison of our conformal method (Section 5.4) with the P2P-Cauchy complex barycentric coordinates of [Weber et al. 2009]. While both methods use the same holomorphic subspace (Equation (8)), the unconstrained optimization of [Weber et al. 2009] develops singularities and injectivity is not maintained. In contrast, our mappings are *certified* as C^∞ conformal mappings.

In Figure 6 we show a deformation of the Giraffe using our method and compare it to Provably Good Planar Mappings (PGPM) [Poranne and Lipman 2014] and As-Rigid-As-Possible (ARAP) [Liu et al. 2008]. The Gaussians used in PGPM are agnostic to the shape of the source model, leading to smooth but highly distorted result. The ARAP method failed to produce a locally injective mapping.

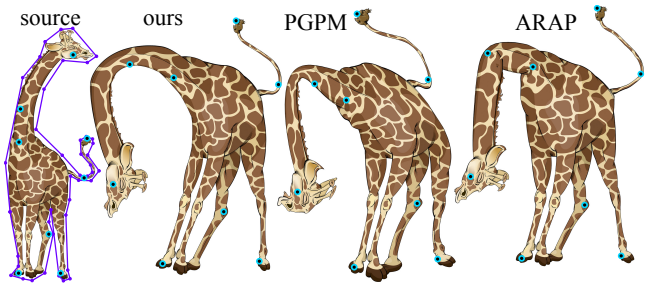

Figure 6: *Comparison of the Giraffe. Left to right. Source image. Our result with user specified parameters:* $(k, \sigma 1, \sigma 2) = (0.55, 3, 0.3)$. *This mapping is certified as a* $(0.554, 3.008, 0.296)$ *bounded distortion mapping using the procedure described in Section 6. PGPM with 40 Gaussians using Euclidian distance. The mapping is smooth but note the high distortion at the head, neck and tail. The ARAP result looks overall good but a closer look at the neck reveals that the mapping has collapsed and is not locally injective.*

In Figure 7, the Pants model is deformed to an extreme pose using three different methods. [Schüller et al. 2013] uses a *constrained* ARAP minimization that manages to produce locally injective mapping (away from the boundary) but similarly to unconstrained ARAP, fails to be smooth. We applied PGPM using smooth Gaussians but due to the strong concavities of the source shape, the amount of distortion was excessive and the positional constraints failed to be satisfied. We switched to shape aware Gaussians by replacing the inner distance from Euclidean to geodesic one as suggested by [Poranne and Lipman 2014], but the nonsmoothness of the geodesic paths near the reflex vertices of the cage manifested into highly distorted mapping. Our harmonic mapping was able to satisfy the constraints precisely and is visually and mathematically smooth. The global bound for the isometric distortion $\max\{\sigma_1, 1/\sigma_2\}$ was 2.8 and the conformal distortion was bounded by 0.61.

Figure 8 compares our smooth conformal mapping (Section 5.4) with the piecewise linear approximations obtained by [Lipman 2012] and [Levi and Zorin 2014].

8 Summary and Discussion

We have presented a framework for computing smooth harmonic and conformal planar mappings with a bounded amount of distortion. A novel underlying theory reveals the connection between the boundary distortion of a harmonic mappings and the distortion in the interior of the domain. An efficient and accurate algorithm is employed, which requires forcing inequality constraints only on the boundary. Moreover, the amount of samples needed for the analysis of the global bounds is relatively small, leading to tight bounds and the ability to perform the analysis in real time, providing a constant accurate feedback to the user.

Limitations Our framework has several limitations. While using harmonic mappings in most situations is perfectly adequate, in some scenarios it can fail to capture the intended result. For example, assume that the user wants to keep the boundary of the domain fixed while manipulating the interior. Such an effect cannot be captured with a harmonic mapping. Another limitation is the use of soft positional constraints which leads to the inability to guarantee precise satisfaction of positional constraints. This design decision is motivated by the fact that with *soft* positional constraints, the bounded distortion mapping problem is always feasible. Switching to *hard* constraints in our framework is easy but it may draw the problem infeasible. Moreover, in extreme cases, an improper choice of θ may lead to the infeasibility of the convex problem while the nonconvex problem is feasible.

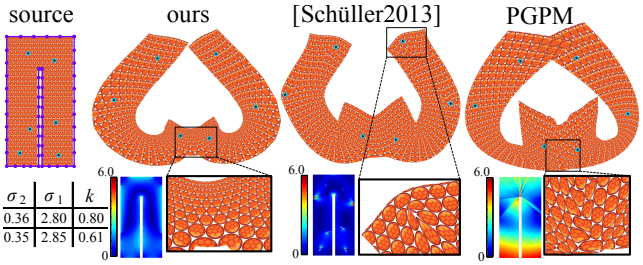

Figure 7: *Deformation of the Pants shape. Comparison of our harmonic mapping with [Schüller et al. 2013] and PGPM using shape aware Gaussians. The user specified bounds for our algorithm are given in the table (first row) along with the global bounds (second row). The color maps illustrate the isometric distortion visualized on the source domains with a common* $[1, 6]$ *range. Note how [Schüller et al. 2013] concentrates the distortion in the vicinity of the handles. PGPM distributes the distortion more uniformly but the result is highly nonsmooth (see the zoom in images).*

Future work We believe that the determination problem of whether the nonconvex (full space) bounded distortion mapping problem is feasible, is one of the most challenging open problems in computer graphics. Navigating inside this nonconvex space is done in our framework using the function θ. However, in contrast to other methods that apply convexification, our θ is merely one dimensional as it is defined solely on the boundary. This implies that the structure of the nonconvex space is simpler for the special case of harmonic mappings and a better characterization of this space is a possible direction for further research. Another interesting direction for future work would be to extend the theory we have developed to volumetric harmonic mappings. This may be challenging since the current theory relies on complex holomorphic functions that do not posses a direct analogue in higher dimensions.

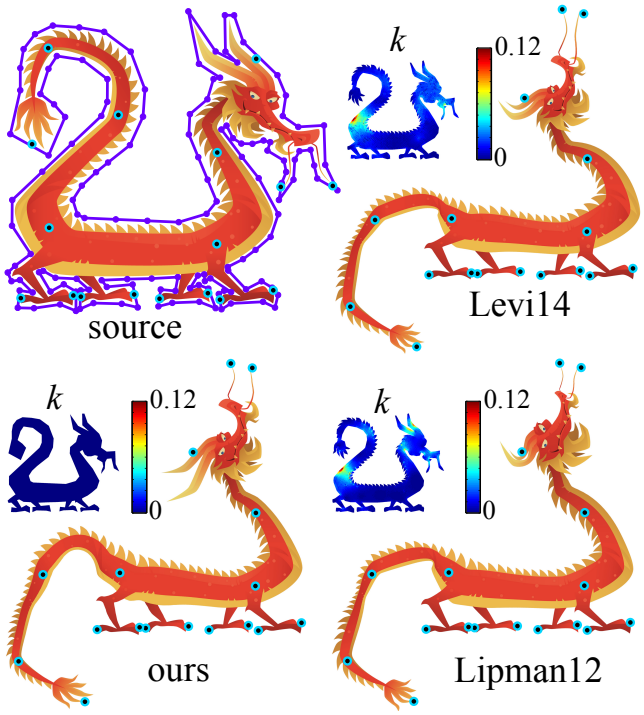

Figure 8: *Conformal deformations of the Dragon. On the bottom left is a smooth conformal mapping computed with our algorithm. The conformal distortion is strictly $k = 0$ while the upper bound on the isometric distortion is* 1.82*. The methods of [Levi and Zorin 2014] and [Lipman 2012] (with LSCM energy) produce discrete conformal approximations with maximal conformal distortion of $k = 0.12$ and maximal isometric distortion of* 6.54 *and* 5.96 *respectively.*

Acknowledgements

We thank Dror Rawitz for insightful discussions and to Yehonatan Wasserman and Alon Bright for their help with the derivation of the Lipschitz constants. This research was partially supported by the BeingThere Centre, collaboration between UNC Chapel Hill, ETH Zurich, and NTU Singapore and by grant number 2012264 from the United States-Israel Binational Science Foundation (BSF). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Quadro k6000 GPU.

References

AHLFORS, L. 1966. *Lectures on quasiconformal mappings*, vol. 38. Amer. Mathematical Society.

AIGERMAN, N., PORANNE, R., AND LIPMAN, Y. 2014. Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics (TOG) 33*, 4, 69.

BELL, S. R. 1992. *The Cauchy transform, potential theory and conformal mapping*, vol. 7. CRC press.

BEN-CHEN, M., WEBER, O., AND GOTSMAN, C. 2009. Variational harmonic maps for space deformation. *ACM Transactions on Graphics (TOG) 28*, 3, 34.

DUREN, P. 2004. *Harmonic mappings in the plane*. Cambridge University Press.

FLOATER, M. S., AND KOSINKA, J. 2010. On the injectivity of wachspress and mean value mappings between convex polygons. *Advances in Computational Mathematics 32*, 2, 163–174.

FLOATER, M. S., KÓS, G., AND REIMERS, M. 2005. Mean value coordinates in 3d. *Computer Aided Geometric Design 22*, 7, 623–631.

HORMANN, K., AND FLOATER, M. S. 2006. Mean value coordinates for arbitrary planar polygons. *ACM Transactions on Graphics (TOG) 25*, 4, 1424–1441.

IGARASHI, T., MOSCOVICH, T., AND HUGHES, J. F. 2005. As-rigid-as-possible shape manipulation. *ACM transactions on Graphics (TOG) 24*, 3, 1134–1141.

JACOBSON, A. 2013. Bijective mappings with generalized barycentric coordinates: a counterexample. *Journal of Graphics Tools 17*, 1-2, 1–4.

JOSHI, P., MEYER, M., DEROSE, T., AND GREEN, B. 2007. Harmonic coordinates for character articulation. In *ACM Transactions on Graphics (TOG)*, vol. 26, 71.

JU, T., SCHAEFER, S., AND WARREN, J. 2005. Mean value coordinates for closed triangular meshes. In *ACM Transactions on Graphics (TOG)*, vol. 24, 561–566.

KHAREVYCH, L., SPRINGBORN, B., AND SCHRÖDER, P. 2006. Discrete conformal mappings via circle patterns. *ACM Trans. Graph. 25*, 2, 412–438.

KOVALSKY, S. Z., AIGERMAN, N., BASRI, R., AND LIPMAN, Y. 2014. Controlling singular values with semidefinite programming. *ACM Transactions on Graphics (TOG)*, 4.

LEVI, Z., AND ZORIN, D. 2014. Strict minimizers for geometric optimization. *ACM Transactions on Graphics (TOG) 33*, 6, 185.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM Transactions on Graphics (TOG) 21*, 3, 362–371.

LIPMAN, Y., LEVIN, D., AND COHEN-OR, D. 2008. Green coordinates. In *ACM Transactions on Graphics (TOG)*, vol. 27, 78.

LIPMAN, Y. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM Transactions on Graphics (TOG) 31*, 4, 108.

LIU, L., ZHANG, L., XU, Y., GOTSMAN, C., AND GORTLER, S. 2008. A local/global approach to mesh parameterization. In *Computer Graphics Forum*, vol. 27, Wiley Online Library, 1495–1504.

PORANNE, R., AND LIPMAN, Y. 2014. Provably good planar mappings. *ACM Transactions on Graphics (TOG) 33*, 4, 76.

SCHNEIDER, T., HORMANN, K., AND FLOATER, M. S. 2013. Bijective composite mean value mappings. In *Computer Graphics Forum*, vol. 32, Wiley Online Library, 137–146.

SCHÜLLER, C., KAVAN, L., PANOZZO, D., AND SORKINE-HORNUNG, O. 2013. Locally injective mappings. In *Computer Graphics Forum*, vol. 32, Wiley Online Library, 125–135.

SHEFFER, A., AND DE STURLER, E. 2001. Parameterization of Faceted Surfaces for Meshing using Angle-Based Flattening. *Engineering with Computers 17*, 3, 326–337.

SORKINE, O., AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, 109–116.

SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM Transactions on Graphics (TOG) 27*, 3, 77.

WEBER, O., AND GOTSMAN, C. 2010. Controllable conformal maps for shape deformation and interpolation. *ACM Transactions on Graphics (TOG) 29*, 4, 78.

WEBER, O., AND ZORIN, D. 2014. Locally injective parametrization with arbitrary fixed boundaries. *ACM Transactions on Graphics (TOG) 33*, 4, 75.

WEBER, O., SORKINE, O., LIPMAN, Y., AND GOTSMAN, C. 2007. Context-aware skeletal shape deformation. In *Computer Graphics Forum*, vol. 26, Wiley Online Library, 265–274.

WEBER, O., BEN-CHEN, M., AND GOTSMAN, C. 2009. Complex barycentric coordinates with applications to planar shape deformation. In *Computer Graphics Forum*, vol. 28, Wiley Online Library, 587–597.

WEBER, O., BEN-CHEN, M., GOTSMAN, C., AND HORMANN, K. 2011. A complex view of barycentric mappings. In *Computer Graphics Forum*, vol. 30, Wiley Online Library, 1533–1542.

WEBER, O., MYLES, A., AND ZORIN, D. 2012. Computing extremal quasiconformal maps. In *Computer Graphics Forum*, vol. 31, Wiley Online Library, 1679–1689.

WEBER, O., PORANNE, R., AND GOTSMAN, C. 2012. Biharmonic coordinates. In *Computer Graphics Forum*, vol. 31, Wiley Online Library, 2409–2422.

WEBER, O. 2010. *Hybrid Methods for Interactive Shape Manipulation*. PhD thesis, Technion - Israel Institute of Technology.

ZAYER, R., RÖSSL, C., KARNI, Z., AND SEIDEL, H.-P. 2005. Harmonic guidance for surface deformation. In *Computer Graphics Forum*, vol. 24, Wiley Online Library, 601–609.

A Expressions for Cauchy Coordinates

For completeness we provide the expressions for the Cauchy coordinates as well as its first and second complex derivatives taken from appendices B,C, and D in [Weber 2010]. We use the notations of Figure 2.

A.1 The Cauchy Coordinates

$$C_j(z) = \frac{1}{2\pi \mathrm{i}} \left(\frac{B_{j+1}(z)}{A_{j+1}} \log \frac{B_{j+1}(z)}{B_j(z)} - \frac{B_{j-1}(z)}{A_j} \log \frac{B_j(z)}{B_{j-1}(z)} \right)$$

A.2 The Derivative of Cauchy Coordinates

$$C_j'(z) = \frac{1}{2\pi \mathrm{i}} \left(\frac{1}{A_{j+1}} \log \frac{B_j(z)}{B_{j+1}(z)} + \frac{1}{A_j} \log \frac{B_j(z)}{B_{j-1}(z)} \right)$$

A.3 The Second Derivative of Cauchy Coordinates

$$C_j''(z) = \frac{1}{2\pi \mathrm{i}} \left(\frac{z_{j+1} - z_{j-1}}{B_{j-1}(z) B_j(z) B_{j+1}(z)} \right)$$

B Proof of Proposition 9

We start by proving the following small lemma.

Lemma 12. *For any two complex numbers a and b the following inequality holds*

$$\Big| \, |a| - |b| \, \Big| \leq |a - b| \, .$$

Proof. Let us write

$$|a| = |(a - b) + b| \leq |a - b| + |b| \, ,$$

where the right inequality is due to the triangle inequality. By shifting $|b|$ to the left we get

$$|a| - |b| \leq |a - b| \, , \tag{35}$$

Similarly, we can write

$$|b| = |(b - a) + a| \leq |a - b| + |a| \, ,$$

and manipulating to get

$$- |a - b| \leq |a| - |b| \, . \tag{36}$$

Combining Equations (35) and (36) gives

$$- |a - b| \leq |a| - |b| \leq |a - b| \, ,$$

which concludes the proof. □

Recalling Proposition 9:

If g is L-Lipschitz continuous then $|g|$, $\mathrm{Re}\,(g)$, $\mathrm{Im}\,(g)$, and $\overline{g}$ are also L-Lipschitz continuous.

Proof. $|g|$ is L-Lipschitz continuous by the following inequality

$$\Big| \, |g(z)| - |g(w)| \, \Big| \leq |g(z) - g(w)| \leq L \, |z - w| \, ,$$

where the right inequality is simply since g is L-Lipschitz and the left inequality is the application of Lemma 12 to $g(z)$ and $g(w)$. We proceed by proving that $\mathrm{Re}\,(g)$ is L-Lipschitz continuous. This is shown by the following statement

$$|\mathrm{Re}\,(g(z)) - \mathrm{Re}\,(g(w))| = |\mathrm{Re}\,(g(z) - g(w))| \leq |g(z) - g(w)| \, ,$$

where the left equality follows from the linearity of the Re operator and the right inequality follows from the fact that the modulus of any complex number is no smaller than the absolute value of its real part. From the above equation we can deduce that

$$|\mathrm{Re}\,(g(z)) - \mathrm{Re}\,(g(w))| \leq L \, |z - w| \, .$$

Showing that $\mathrm{Im}\,(g)$ is L-Lipschitz continuous is done similarly.

Finally, showing that $\overline{g}$ is L-Lipschitz is done by using the equality

$$\left| \overline{g(z)} - \overline{g(w)} \right| = |g(z) - g(w)| \, .$$ □

C Lipschitz Constant for $C_j'(z)$

We begin the derivation by proving a simple lemma that is then used to compute the Lipschitz constant. While the real analogue of this lemma is a classic result, we could not find in the literature an analogue for holomorphic functions.

Lemma 13. *Let g be a holomorphic function defined on a convex domain $\Omega \in \mathbb{C}$ with a bounded first complex derivative, then $L = \|g'\|_\infty$ is a Lipschitz constant for g.*

Proof. We show this by the definition of Lipschitz continuity

$$|g(z) - g(w)| = \left| \int_c g'(\xi) d\xi \right| \leq \int_c \|g'\|_\infty ds = \|g'\|_\infty \, |z - w| \, ,$$

where c is any path between w and z fully contained in Ω. The first equality is due to Cauchy's integral theorem and the last equality is due to convexity of Ω that allows us to use the straight line connecting w to z as the shortest possible path. □

It is essential to realize that while our polygonal domain is not necessarily convex, we can still apply the above lemma since the entire analysis and computations of Lipschitz constants in our case are done on a single (convex) segment.

To compute a Lipschitz constant for $C_j'(z)$ on an arbitrary segment (v_1, v_2) strictly inside the cage, it is sufficient to bound $\left|C_j''(z)\right|$ from above. Specifically (see Appendix A.3) we need an upper bound for

$$\frac{|z_{j+1} - z_{j-1}|}{2\pi \left|B_{j-1}(z)\right| \left|B_j(z)\right| \left|B_{j+1}(z)\right|},$$

where $z \in (v_1, v_2)$. The geometric meaning of $|B_j| = |z_j - z|$ is simply the distance between the j^{th} cage vertex and a point on the segment z. To obtain an upper bound we substitute each $|B_j|$ in the denominator with the distance between z_j and the entire segment

$$L_{C_j'} = \frac{|z_{j+1} - z_{j-1}|}{2\pi\, d(z_{j-1}) d(z_j) d(z_{j+1})}.$$

Note that the first derivative of the Cauchy coordinates has logarithmic singularity on the vertices of the cage. Nonetheless, our segments are located on P, away from the offset polygon $\tilde{P}$, hence, d cannot vanish.

To compute the distance $d(p)$ between a point p and a segment we simply use the parametrization $z = (1-t)v_1 + tv_2$ and compute the minimum distance with respect to t. The closed form solution is given by

$$t = \frac{\mathrm{Re}\left((v_2 - v_1)(\overline{p} - \overline{v_1})\right)}{|v_2 - v_1|^2},$$

$$d(p) = \begin{cases} |p - (1-t)v_1 - tv_2| & 0 < t < 1 \\ |p - v_1| & t \le 0 \\ |p - v_2| & t \ge 1 \end{cases}$$

D Proof of Theorem 11

Recalling Theorem 11:

Let f be a complex-valued harmonic function defined on a simply connected domain Ω. Let $\theta(w)$ be any real-valued continuous function defined on the boundary. We denote by γ the function

$$\gamma(w) = \mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right).$$

If $\gamma(w) > 0$ at every point on the boundary *then f_z does not vanish* inside *the domain.*

Proof. Since $\theta(w)$ is continuous, it has a unique harmonic extension $\theta(z)$ to the interior of the domain (by existence and uniqueness of the solution to the Dirichlet problem). On a simply connected domain, a real-valued harmonic function $\theta(z)$ always admits a unique (up to a real constant) harmonic conjugate ϕ such that $\theta(z) + \mathrm{i}\phi(z)$ is holomorphic. The operator that takes θ to ϕ is called the Hilbert transform [Bell 1992]. Let us define the function

$$g(z) = \mathrm{e}^{\mathrm{i}(\theta(z) + \mathrm{i}\phi(z))} = \frac{\mathrm{e}^{\mathrm{i}\theta(z)}}{\mathrm{e}^{\phi(z)}}.$$

$g(z)$ is holomorphic since the complex exponential is holomorphic and composition of holomorphic functions is holomorphic. We proceed by defining another function

$$\eta(z) = \mathrm{Re}\left(f_z(z)g(z)\right) = \frac{1}{\mathrm{e}^{\phi(z)}}\mathrm{Re}\left(f_z(z)\mathrm{e}^{\mathrm{i}\theta(z)}\right).$$

Since the product of holomorphic functions is holomorphic, we have that $f_z(z)g(z)$ is holomorphic and its real part $\eta(z)$ is harmonic. It follows from the harmonicity of η that the minimum of η is attained on the boundary. η is positive on the boundary (since γ is), hence must be positive everywhere (and so is γ). Finally, since γ is positive everywhere, f_z cannot vanish. □

E The Positivity of γ

We begin by deriving the Lipschitz constant for

$$\gamma(w) = \mathrm{Re}\left(f_z(w)\mathrm{e}^{\mathrm{i}\theta(w)}\right).$$

Proposition 9 provides us with the comfort of computing the Lipschitz constant of $f_z\mathrm{e}^{\mathrm{i}\theta}$ instead, which is simpler. For that we extend a well known lemma for real-valued functions to complex-valued functions.

Lemma 14. *Let f and g be holomorphic with Lipschitz constants L_f and L_g respectively, then the pointwise product fg has Lipschitz constant*

$$L_{fg} = L_f \|g\|_\infty + L_g \|f\|_\infty.$$

Proof. Since the product fg is holomorphic we can use the product rule for holomorphic functions. Then the following is true at any point in the domain

$$\left|\frac{d}{dz}fg\right| = \left|f'g + g'f\right| \le \left|f'g\right| + \left|g'f\right| = \left|f'\right||g| + \left|g'\right||f|.$$

Since $|f'|$ is bounded from above by L_f and $|g'|$ is bounded from above by L_g (Lemma 13) we have

$$\left|f'\right||g| + \left|g'\right||f| \le L_f\|g\|_\infty + L_g\|f\|_\infty,$$

which concludes the proof. □

Using Lemma 14 and the fact that $\left|\mathrm{e}^{\mathrm{i}\theta}\right| = 1$ we have

$$L_{f_z\mathrm{e}^{\mathrm{i}\theta}} = L_{f_z} + L_{\mathrm{e}^{\mathrm{i}\theta}}\|f_z\|_\infty. \tag{37}$$

Since θ was chosen to be piecewise linear, we have that $L_{\mathrm{e}^{\mathrm{i}\theta}} = \left|\theta^2 - \theta^1\right|/l$. Next, we substitute $\|f_z\|_\infty$ in Equation (37) with the upper bound obtained in Equation (26) and obtain the following explicit expression for L_γ

$$L_\gamma = L_{f_z} + \frac{\left|\theta^2 - \theta^1\right|}{2\,l}\left(\left|f_z^1\right| + \left|f_z^2\right| + L_{f_z}l\right). \tag{38}$$

With that in hand, we can set a lower bound on γ

$$\gamma_{\min} = \frac{\gamma^1 + \gamma^2 - L_\gamma l}{2}. \tag{39}$$

Using $\theta(w^i) = -\arg f_z(w)\big|_{w^i}$ at the samples leads to the following expression for γ at the samples

$$\gamma(w^i) = \mathrm{Re}\left(f_z(w^i)\frac{\left|f_z(w^i)\right|}{f_z(w^i)}\right) = \left|f_z(w^i)\right|. \tag{40}$$

Finally, the positivity of γ on the segment can be guaranteed if $\gamma_{\min} > 0$ (Equation (39))

$$\gamma^1 + \gamma^2 - L_\gamma l > 0. \tag{41}$$

Substituting $\gamma^1 = \left|f_z^1\right|$, $\gamma^2 = \left|f_z^2\right|$, and L_γ (Equation (38)) in the above equation and further manipulating leads to the final condition on the positivity of γ

$$(2 + \left|\theta^2 - \theta^1\right|)\, l\, L_{f_z} < (2 - \left|\theta^2 - \theta^1\right|)\Big(\left|f_z^1\right| + \left|f_z^2\right|\Big).$$