Computation of Conformal Maps by Modified Schwarz-Christoffel Transformations

by

Louis Hill Howell

B.S. Mathematics, Massachusetts Institute of Technology (1985)

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 1990

Signature of Author ______________________________
Department of Mathematics
January 12, 1990

Certified by ______________________________
Lloyd N. Trefethen
Associate Professor of Applied Mathematics
Thesis Supervisor

Accepted by ______________________________
Daniel J. Kleitman
Chairman, Applied Mathematics Committee

Accepted by ______________________________
Sigurdur Helgason
Chairman, Departmental Graduate Committee

Computation of Conformal Maps by

Modified Schwarz-Christoffel Transformations

by

Louis Hill Howell

Submitted to the Department of Mathematics
on January 12, 1990, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Numerical conformal mapping techniques can be roughly divided into two categories. On one hand there are many methods which perform well on regions with smooth boundaries, but which either fail or require elaborate modifications when mapping regions with corners. On the other hand there is the Schwarz-Christoffel transformation, which is specifically designed for mapping polygons with any number of corners, but which does not, at least in its original form, extend to mapping domains with curved sides.

In this thesis we explore four issues regarding the use of the Schwarz-Christoffel transformation and related methods as robust conformal mapping techniques. The first of these, the crowding phenomenon, makes maps onto even mildly elongated regions so ill-conditioned they may be impossible to calculate. We overcome this difficulty by implementing a transformation, similar to the Schwarz-Christoffel formula, which uses an infinite strip as the standard domain. The mapping function itself is not original; it has previously been used by others to compute maps onto open channels for internal flow computations. The interpretation in terms of crowding is original, however, and we also present a new modification for mapping elongated rectangles onto closed polygons. By using robust and efficient software for numerical integration and solution of the parameter problem, we successfully map difficult regions with aspect ratios of up to thousands to one.

Second, we survey a number of techniques for numerically evaluating Schwarz-Christoffel integrals, including the compound Gauss-Jacobi algorithm, singularity removal methods, and several adaptive quadrature subroutines. The Gauss-Jacobi method is shown to be faster than its nearest competitor by roughly a factor of two; we recommend it for all applications unless strict error bounds are required. Experimental data are presented to support our conclusions.

Third, all Schwarz-Christoffel methods require the solution of a system of nonlinear algebraic equations. We examine two approaches to this problem that have appeared in the literature, Davis's iterative scheme and Trefethen's formulation as a nonlinear system. We show that Davis's method is difficult to generalize and is not always even locally convergent, and prove that although Trefethen's method never diverges locally, it too may diverge from distant starting points. We then discuss several approaches to the problem of designing a globally convergent algorithm.

Finally, a generalization of the Schwarz-Christoffel formula, dating back to Schwarz in 1869, describes conformal maps onto regions bounded by arbitrary arcs of circles as solutions to a nonlinear ordinary differential equation. The only attempt yet made to implement this equation numerically, by Bjørstad and Grosse in 1987, was not entirely successful. In the present work we examine several sources of error in their approach, most of which were previously unrecognized. We present improvements in methods for integration of the differential equation, normalization of approximate solutions, specification of the accessory parameters, formulation of the nonlinear system, and selection of appropriate starting points involving a minimum amount of distortion.

Thesis Supervisor: Lloyd N. Trefethen
Title: Associate Professor of Applied Mathematics

Acknowledgments

This thesis would quite simply not have been possible without the continuing support and encouragement of my advisor, Nick Trefethen. His influence has been instrumental in leading me to my present niche along the hazy boundary between mathematics and computing. I consider myself fortunate to have had such a caring mentor, who felt it much more important that he should be an asset to my career, than I to his.

Many of the ideas and techniques forming the foundations for Chapters 2 and 4 were originally Nick's. I am also indebted to Eric Grosse for a lengthy discussion of the properties of circular arc polygons, and to Frederick Dias and Tom DeLillo for sharing their general perspectives on conformal mapping. So many of the departmental riff-raff[1] have helped me struggle with the myriad details of this thesis that I doubt I can list them all: Many thanks to aed, alan, brasil, csn, dfine, engel, erandall, figa, jihad, jrs, jso, kerswell, lmsmith, mike, peter, prisia, reddy, rpt, santa, wally, walter, and ycchen, and my sincere apologies to those I have inevitably left out.

Though the 7.5 years I have spent at MIT have had a tremendous part in shaping my life, I must admit that the effects have been mostly due to cumulative, rather than individual, interactions. So many great minds, both professors and students, have touched me in some way during this time that it is hard to cite specific names without feeling that I am slighting many others. I therefore single out just one person, Alar Toomre, not so much for his many individual contributions, but because he comes closest in my mind to what I think of as the spirit of the Institute. I will always think of him as the canonical example of a good lecturer, and will surely recognize his handwriting until the day I die.

I would like to also acknowledge the support of some of the many people who have helped me reach this point without giving up or losing my mind. I thank Victor Chandler, Richard Coleman, Thomas Costello, George Fink, Chuck Geanangel, Richard Haggbloom, Bob Pielow, Renata Schemmer, and Breda Sheedy for leading me, pushing me, and occasionally cracking the whip when I needed a bit of extra motivation. I thank my parents for their unflagging support and encouragement and for standing behind me while letting me go my own way. My two daughters have been at times welcome, or unwelcome, diversions, and at times true sources of inspiration. Finally, Pamela, my wife, has put up with me and loved me through five years of Boston winters, slow buses, rude yankees, and night feedings. On the day we met my life took a definite turn for the better, and with any luck our future will be better still.

In loving memory of William K. Howell, Jr.

This work was supported by U.S. Air Force Grant AFOSR-87-0102, and also by a Presidential Young Investigator Award to Lloyd N. Trefethen.

[1] graduate students

Contents

Abstract **2**

Acknowledgments **4**

1 Introduction **12**

1.1 Conformal Mapping 12

1.2 The Schwarz-Christoffel Transformation 15

1.3 Numerical Methods 20

1.4 Contributions of this Thesis 21

2 Highly Elongated Polygons **24**

2.1 The Crowding Problem 25

2.2 The Strip Transformation 30

2.3 Numerical Computations 35

2.4 Mapping Rectangles to Closed Polygons 39

2.5 Variations 44

3 Evaluating Schwarz-Christoffel Integrals **48**

3.1 Paths of Integration 49

3.2 Compound Gauss-Jacobi Quadrature 50

3.3 Singularity Removal 52

3.4 Adaptive Quadrature 54

3.5 Conclusions 59

4 Solving the Parameter Problem **60**

4.1 Davis's Iteration: Local Divergence 62

4.2 Nonlinear Equations: Local Convergence, Global Divergence 67

4.3 Conditions for Global Convergence 74

5 Circular Arc Polygons **81**

5.1 The Schwarzian Differential Equation 81

5.2 Derivation 84

5.3 Numerical Integration 87

5.4 Ill-Conditioning and the Accessory Parameters 94

5.5 Formulation of the Nonlinear System 104

5.6 Mapping from a Strip 114

5.7 Summary 118

A Elliptic Functions **120**

Bibliography **124**

List of Figures

1-1 A conformal map from the unit disk to a simple polygon. 13

1-2 The elementary transformation $w = (z - z_j)^{\alpha_j - 1}$. 15

1-3 Two polygons with identical angles but different side lengths. 16

1-4 Only $n - 2$ of the sides of a polygon can be varied independently. . . . 17

1-5 Reflecting the image polygon through two different sides produces a new branch of the mapping function. This new branch is related to the original branch by a linear transformation. 18

2-1a Polygons which can be mapped by the standard Schwarz-Christoffel transformation. 26

2-1b Polygons which can be mapped by the strip transformation. 26

2-1c Polygons which cannot be mapped by either method. 26

2-2 Conformal map from a rectangle to the upper half-plane. 27

2-3 Conformal maps of the unit disk onto four rectangles. 28

2-4a Prevertices on the strip—solid dots show the correct values z_j^*, open circles show an incorrect set of values z_j. 32

2-4b Target polygon P^* defined by vertices w_j^*. 32

2-4c Polygon P defined by incorrect vertices w_j. 32

2-5 A single factor f_j, shown for $\alpha_j > 1$. 34

2-6 A moderately difficult elongated channel, Jacobian condition number approximately 35000. 37

2-7 Two channels with large n. 38

2-8 Conformal map from a rectangle to a strip. 40

2-9 Conformal map of a rectangle onto a distorted polygon ($\mu \approx 41.812465$). 42

2-10a Conformal map of a rectangle onto a star ($\mu \approx 163.28151$). Note that both ends of the polygon meet at the lower-left corner. 43

2-10b Conformal map of a rectangle onto a finite spiral ($\mu \approx 132.70454$). . . . 43

2-11 A polygon with 23 sides ($\mu \approx 156.6241139$). 44

2-12a A channel map with converging ends. 46

2-12b Enlargement of the barb in Figure 2-12a. 46

2-12c A rectangle map onto a polygonal conductor ($\mu \approx 49.436547$). 46

2-13 Conformal map onto an infinite spiral. 47

3-1 A simple polygon mapped from a rectangle using the strip transformation. 57

4-1 This conformal map from the upper half-plane to a polygon causes one version of Davis's iteration to diverge. 64

4-2 Davis's iteration is locally divergent when used with a strip transformation to this distorted channel. 66

4-3 No descent algorithm based entirely on side lengths can transform one of these polygons into the other. 73

4-4 A completely general algorithm should in principle be able to map even this polygon. 74

4-5 This type of self-intersection is permitted. However, if we examine sides sequentially starting at corner w_1, we encounter what *appears* to be an illegal intersection between w_4 and w_5. 76

4-6a This polygon does not intersect itself. 77

4-6b A self-intersecting version of the above polygon, this figure could easily be produced as an intermediate step while solving the parameter problem. 77

4-6c The lower boundary of this illegal polygon is identical to that of Figure 4-6b up to corner w_{14}, including the region where the intersection occurs. 77

5-1 Conformal map of the unit disk onto a circular arc polygon. The lines shown are images of radii and concentric circles in the disk. 83

5-2 This bounded region is more naturally mapped from the unit disk than from the upper half-plane. 95

5-3 An 8-sided circular arc polygon. No advantage has been taken of symmetry in computing this map. 101

5-4a Initial guess for the 7-sided polygon of Figure 5-2 using the standard parameters β_j. 102

5-4b Initial guess for the same polygon using the modified parameters β'_j. . . . 102

5-5a Initial guess for the 8-sided polygon of Figure 5-3 using the standard parameters β_j. 103

5-5b Initial guess for the same polygon using the modified parameters β'_j. . . . 103

5-6 This region has exactly the same corners and angles as Figure 5-3. Only the orientations of the corners have been changed. 106

5-7 These two regions demonstrate that even specifying the corners *and* their orientations is not always sufficient to characterize a polygon. 107

5-8a This conformal map to a circular arc polygon does not send the center of the disk to the desired point, identified here by a black dot. 109

5-8b Changing the normalization to send the center to the correct point, however, inverts the polygon and reverses the orientation of the side from w_n to w_1. 109

5-9 The two solid dots, and the orientations of the lines at those points, are assumed to be determined by the remainder of the polygon (not shown). Holding the curvature of the left side fixed is not sufficient to make the length of that side, and thus the position of w_1, unique. 111

5-10 A simple example based on one computed by Bjørstad and Grosse. . . . 112

5-11 Bjørstad and Grosse's algorithm fails to converge to this region without the continuation method. Most of the Riemann sphere algorithms succeed, however. 113

List of Tables

3.1 Performance of three adaptive quadrature procedures, with various singularity removal options, on a 10-corner strip transformation problem. . . . 56

3.2 Performance of three variations on the routine QUANC8, with two different singularity removal options, on a 10-corner strip transformation problem. 58

Chapter 1

Introduction

1.1 Conformal Mapping

Given two open sets in the complex plane, a conformal map is an invertible analytic function sending one of these domains onto the other (Figure 1-1). Such maps are useful due to the relationship between complex analytic functions and solutions to Laplace's equation in two dimensions,

$$\nabla^2\phi \equiv \frac{d^2\phi}{dx^2} + \frac{d^2\phi}{dy^2} = 0. \tag{1.1}$$

If we let $z = x + iy$ be the independent complex variable and $f(z) = u(x, y) + iv(x, y)$ the mapping function, where x, y, u and v are all real, then $\nabla^2 u = 0$ and $\nabla^2 v = 0$ follow immediately from elementary complex analysis (cf. [2, ch. 2]). Moreover, if a real function $\phi_1(u, v)$ satisfies (1.1), then $\phi_2(x, y) = \phi_1[u(x, y), v(x, y)]$ does as well. A solution to Laplace's equation on a simple domain, e.g., the unit disk or a rectangle, can therefore be transformed into a solution on a more complicated domain given a conformal map between the two regions.

Laplace's equation itself is important in such diverse areas as heat conduction, electrostatics, and fluid mechanics. In a solid body of uniform composition, the temperature distribution approaches a solution to this equation as a steady state. The conductors

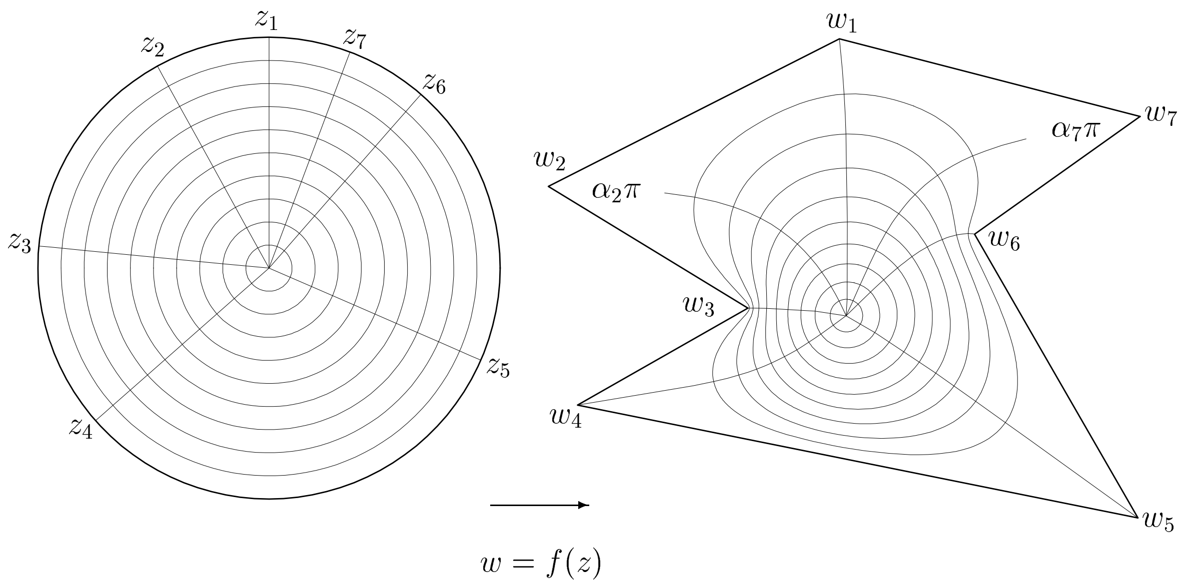

Figure 1-1: A conformal map from the unit disk to a simple polygon.

on integrated circuit chips are essentially 2-dimensional, so solutions obtained by conformal mapping are particularly appropriate for computation of such circuit properties as resistance [64] and capacitance [28], [41]. Other solutions describe certain inviscid fluid flows, and (1.1) even plays a role in the implementation of boundary conditions for vortex methods, which are designed to simulate unsteady flows at high Reynolds number [29]. (For vortex calculations of internal flows, for example, it would be natural to combine the conformal map from an infinite strip discussed in Chapter 2 with L. Greengard's recent algorithm for fast calculation of vortex interactions in such a strip [32].) Conformal mapping can also be applied directly to grid generation for other computational algorithms ([9], [21], [62]), though it provides less control over various important grid properties than do some other methods.

Given two topologically equivalent regions of the complex plane, the choice of conformal maps between these regions is quite limited. If S is any simply connected open set which is not the entire complex plane, the Riemann mapping theorem [35, pp. 324–328] states that there is always a unique conformal map from the unit disk to S that maps the origin to a given point of S and the positive real axis to a curve leaving that point in a given direction. (Conformal maps can also be found between some domains which are not simply connected, but we will not be concerned with such maps in this thesis.) It is not possible to specify arbitrary boundary conditions for the map between the two regions. At most three points of the boundary correspondence function may be specified; the rest must be computed along with the interior distribution.

A limited number of useful conformal maps can be derived from straightforward combinations of elementary functions and a few special functions (cf. Appendix A). For more general domains, however, numerical methods are essential. It is possible to calculate the mapping function directly by solving Laplace's equation on a region, but this approach is inherently two-dimensional and would often be as expensive as solving the original physical problem directly. A one-dimensional approach, where the boundary correspondence function is determined first as a stepping stone to the two-dimensional map, is much less costly.

A great variety of numerical methods can be used for generating conformal maps, a surprising number of which are currently being applied in practical algorithms. The diversity of these methods is in part due to the many connections between conformal mapping and other areas of mathematics.[1] Regions with smooth boundaries can be efficiently mapped by many different techniques, but corners often present special difficulties because they involve singularities in the mapping function. The Schwarz-Christoffel transformation and its relatives, on the other hand, explicitly take corners into account but have some difficulty representing arbitrary smooth curves. These methods are the subject of this thesis.

[1] For an overview of numerical conformal mapping techniques, see [35], [39] and [65]. Despite the apparent differences, many of these methods can be brought under a common mathematical framework—for more on this see the recent article by DeLillo [10], which builds upon earlier work by Gaier [26] and Gutknecht [33].

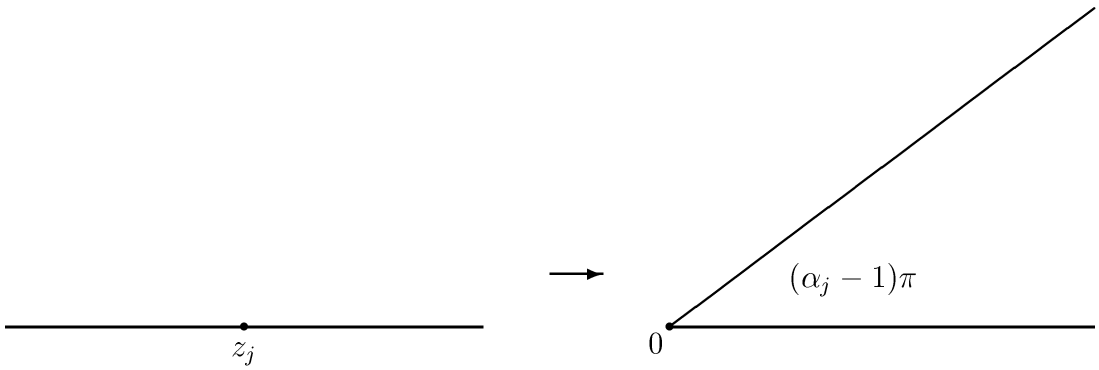

Figure 1-2: The elementary transformation $w = (z - z_j)^{\alpha_j - 1}$.

1.2 The Schwarz-Christoffel Transformation

Any conformal map from the upper half-plane to a simply connected polygon, with or without corners at infinity, can be expressed in the form

$$f(z) = A \int^z \prod_{j=1}^{n} (z' - z_j)^{\alpha_j - 1} dz' + B, \tag{1.2}$$

known as the Schwarz-Christoffel formula. In this equation z is the independent complex variable in the upper half-plane, the z_j are the "prevertices", i.e., points on the real axis which are preimages of the corners $w_j = f(z_j)$ of the polygon, the α_j are $1/\pi$ times the interior angles at those corners, and A and B are two complex constants. The same formula also applies to maps from the unit disk, with the prevertices z_j then falling on the unit circle. To convert between these two standard domains, it is merely necessary to substitute $h(z)$ and $h(z_j)$ for z and z_j, respectively, in equation (1.2)—the resulting expression can then be simplified back to the original form. Here $h(z)$ is a conformal map from the unit disk to the upper half plane:

$$h(z) = i \left(\frac{1 - z}{1 + z} \right). \tag{1.3}$$

The behavior of the mapping function (1.2) can be understood in two essentially different ways, each of which suggests its own set of possible variations to the basic for-

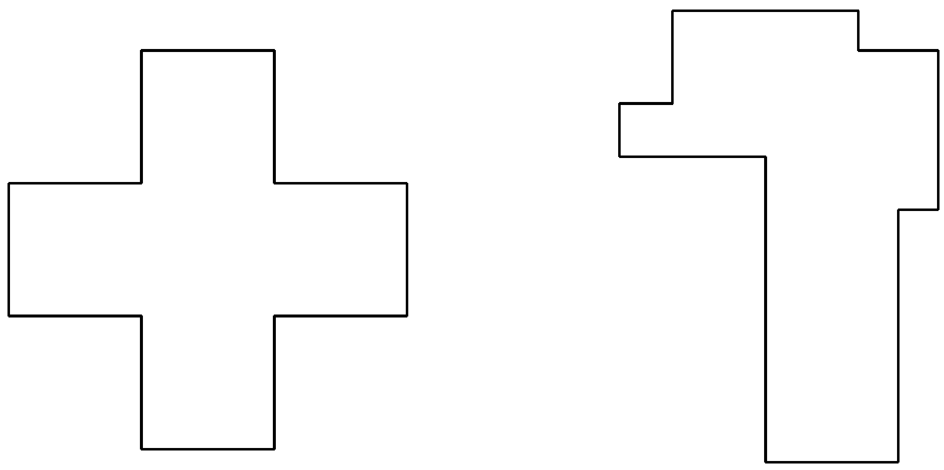

Figure 1-3: Two polygons with identical angles but different side lengths.

mula. The first approach is constructive, in that the Schwarz-Christoffel transformation is generated by combining a number of more elementary conformal maps. The action of the map $w = (z - z_j)^{\alpha_j - 1}$ on the upper half-plane is illustrated in Figure 1-2; note that the two halves of the real axis are mapped onto two lines of constant argument. The product of n of these elementary factors thus maps the real axis onto $n+1$ lines of constant argument, which upon integration yield the straight sides of the image polygon. Provided that the turning angles $(1 - \alpha_j)\pi$ at the corners of the polygon sum to exactly 2π, the mapping function will be analytic at infinity and the polygon will have exactly n sides. (Otherwise, an additional corner will appear and the polygon will have $n+1$ sides. Some versions of the formula (1.2) have only $n-1$ factors for this reason, with the nth prevertex implicitly falling at infinity.)

The Schwarz-Christoffel transformation therefore maps the upper half plane onto a polygonal region with interior angles $\alpha_j \pi$. Unfortunately, there is no simple way to predict the lengths of the sides of this region, and in general they will be different from the side lengths of the desired polygon (Figure 1-3). Consider the remaining free parameters in (1.2) and the degrees of freedom available to the image polygon. The Riemann

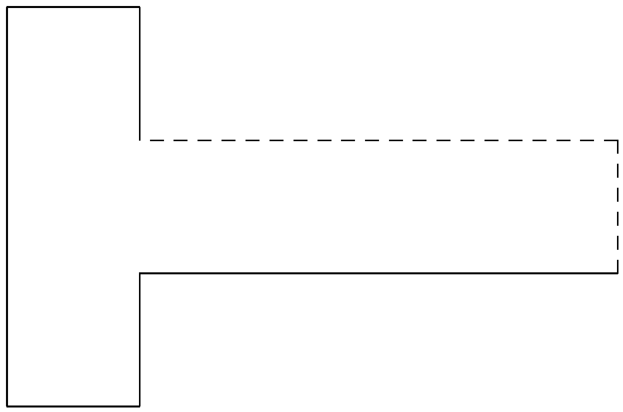

Figure 1-4: Only $n-2$ of the sides of a polygon can be varied independently.

mapping theorem permits three of the prevertices z_j to be specified arbitrarily, leaving $n-3$ real parameters to be determined as part of the solution procedure, plus the two complex constants A and B. If we use A and B to fix the positions of two corners of the image polygon, say w_n and w_1, then exactly $n-3$ of the remaining sides can be varied independently. The last two sides, provided that they are not parallel, are then uniquely determined (Figure 1-4). This correspondence between the number of free parameters in the mapping function and the number of degrees of freedom in the image polygon suggests, but without further elaboration does not prove, that any desired polygon can be mapped given appropriate values for the parameters z_j.

Several variations of the Schwarz-Christoffel formula can be developed by this approach using different elementary factors. Davis [9] considers the limit of a large number of sides with small turning angles in his derivation of a formula producing curved sides. Däppen [7,8] implements a not so well-known variation for maps onto doubly-connected polygons (see also [35, pp. 478–487], [37]). Elcrat, Trefethen and Dias [19,16] present a formula mapping onto a class of polygonal regions where one or two sides correspond to free streamlines in inviscid fluid flow. Finally, we will see in Chapter 2 how the map from

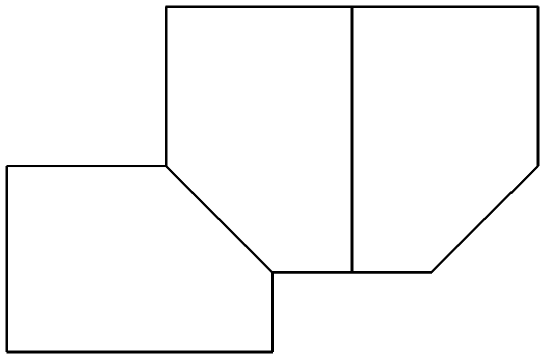

Figure 1-5: Reflecting the image polygon through two different sides produces a new branch of the mapping function. This new branch is related to the original branch by a linear transformation.

an infinite strip, originally derived by other methods, can be developed from the same principle. For more information on these and other similar Schwarz-Christoffel variations, see [23] and [67].

The other approach to the Schwarz-Christoffel transformation is more analytic: Given the existence of the conformal map via the Riemann mapping theorem, deduce the form of the mapping function. This method has the advantage that it demonstrates the validity of (1.2) for the particular desired polygon, not just for a family of related polygons. A detailed derivation of (1.2) using this approach is given by Nehari in [46]—for the sake of brevity we will only sketch the key ideas here.

Consider a map from the upper half-plane to the polygon shown in Figure 1-5. The mapping function f is analytic on the line segment joining any two adjacent prevertices, so by the Schwarz reflection principle we can find an analytic continuation of the map into the lower half plane by reflecting the image polygon through its corresponding side. Reflecting again through one of the other sides of the new polygon yields a continuation across another segment and back into the upper half-plane. This doubly reflected polygon represents a new branch of the original mapping function, and is simply a rotated and translated version of the original image polygon. Now consider the function

$$F(z) = f''(z)/f'(z). \tag{1.4}$$

F is invariant under any linear transformation of f, i.e., any operation sending f to $Af + B$, where A and B are complex constants. These linear transformations include all rotations and translations, so $F(z)$ is the same on every branch of the mapping function. Except for isolated singularities at the prevertices z_j, F is a single-valued analytic function defined at every point of the complex plane. Moreover, if all of the prevertices are finite, then f and F are analytic at infinity as well.

The next step in the derivation is to classify and remove the isolated singularity at each prevertex. Nehari shows that the mapping function near z_j can be written in the form $f(z) = (z - z_j)^{\alpha_j} g(z)$, where g is analytic and nonzero in a neighborhood of z_j. Substituting this expression into equation (1.4), we find that F has a simple pole of the form $(\alpha_j - 1)(z - z_j)^{-1}$ near z_j. If we subtract off all of these poles from the function F, we obtain an expression

$$\frac{f''(z)}{f'(z)} - \sum_{j=1}^{n} \frac{\alpha_j - 1}{z - z_j} \tag{1.5}$$

which is analytic in the entire complex plane.

The expression (1.5) is a bounded entire function and therefore reduces to a constant. This constant is, in fact, zero, as can be seen by substituting a Laurent series for $f(z)$ near infinity into (1.4). Writing f''/f' in a slightly different form, we now have a differential equation,

$$\frac{d}{dz} \log[f'(z)] = \sum_{j=1}^{n} \frac{\alpha_j - 1}{z - z_j}, \tag{1.6}$$

which can easily be integrated to yield equation (1.2), the familiar form of the Schwarz-Christoffel transformation.

We are aware of only one important Schwarz-Christoffel variation which is obtained through an extension of this procedure, that being the Schwarzian differential equation for mapping the upper half-plane onto a region bounded by arbitrary arcs of circles. The idea is to use a more general differential operator,

$$\{f, z\} = \left[\frac{f''(z)}{f'(z)}\right]' - \frac{1}{2}\left[\frac{f''(z)}{f'(z)}\right]^2, \tag{1.7}$$

which is invariant under not only linear transformations but under the reciprocal operation $f \to 1/f$ as well. The computational properties of the resulting differential equation will be the subject of Chapter 5.

1.3 Numerical Methods

Any conformal map from the upper half-plane to a simply connected polygon can be expressed as a Schwarz-Christoffel transformation. Two primary obstacles must be overcome, however, before (1.2) can be used to actually compute the map onto a given region. First, though any three of the prevertices z_j can be placed arbitrarily, there is no known analytic procedure for determining the positions of the others except in a few cases with a high degree of symmetry. Second, even when values for all of the z_j are known, only a few analytic solutions are known for the integral itself. When the image polygon is a triangle there is always a solution in terms of hypergeometric functions, and for rectangles (1.2) becomes an elliptic integral, but for virtually all more general polygons some kind of numerical approximation must be used.

In the century since the development of (1.2) many different numerical quadrature methods have been applied to the integration problem, with varying degrees of success. Difficulties arise due to the singular behavior of the integrand, with singularities both on and off an interval reducing the accuracy of many common schemes. We will discuss these problems at length in Chapter 3—it suffices to say at this point that a combination of adaptive mesh refinement techniques with some kind of analytic treatment for endpoint singularities permits the evaluation of essentially all Schwarz-Christoffel integrals to a high degree of precision.

As for the so-called parameter problem—the problem of determining the prevertex positions—the only general procedure is to use a series of successive approximations designed to converge to the desired solution. No methods are known, short of computing the entire conformal map by another method, for choosing a starting point which is guaranteed to be close to the solution. Most existing algorithms choose a starting point arbitrarily, or nearly so, for instance by making the prevertices equally spaced. It is therefore important that the procedure for choosing new approximations be as robust as possible, so as to make convergence likely even from poor starting points. Methods are now available which converge for most cases within the reach of practical computations. A detailed discussion of the issues involved in the development of such methods will be the subject of Chapter 4.

1.4 Contributions of this Thesis

In this thesis our emphasis is on developing robust algorithms for the actual computation of conformal maps via the Schwarz-Christoffel transformation and its relatives, rather than on results of purely theoretical interest.

- **Highly Elongated Polygons**

  Chapter 2 is devoted to the crowding phenomenon in conformal mapping and its effect on Schwarz-Christoffel maps. A polygon with an aspect ratio as small as 15 to 1 may be impossible to map in finite precision arithmetic. We show that in many cases this problem may be reduced or eliminated by mapping from an infinite strip instead of the unit disk or the upper half-plane. Our discussion of this transformation in terms of crowding is new, though the formula for the mapping function has previously appeared as a method for computing internal flows in channels. We also present a new derivation which clarifies the role played by each term in the formula. Finally, we derive a modification of the strip transformation for mapping rectangles to closed polygons. By using robust software for evaluating integrals and solving the parameter problem, we successfully map highly distorted regions with aspect ratios of up to thousands to one.

- **Evaluating Schwarz-Christoffel Integrals**

  Chapter 3 describes an experimental study of methods for numerically evaluating Schwarz-Christoffel integrals. We investigate the compound Gauss-Jacobi algorithm, the removal of singularities by subtraction and by a change in the variable of integration, and five different adaptive quadrature routines—two quite different routines from QUADPACK and three variations on the routine QUANC8. Our main conclusion, backed by experimental evidence, is that the compound Gauss-Jacobi method is consistently the fastest approach, by a factor of about 2, and is sufficiently reliable to be used whenever strict error bounds are not required. Among the other procedures, the best performance was obtained by combining QUANC8 with a change of variables.

- **Solving the Parameter Problem**

  The only known methods for solving the parameter problem involve highly nonlinear iterative processes which in some cases may fail to converge. In Chapter 4 we examine two approaches which have each been adopted by a number of researchers. A method developed by Davis using geometric insight, aside from being difficult to generalize, is shown through specific counterexamples to sometimes be locally divergent. An alternative is to reformulate the problem as a set of nonlinear equations, since many sophisticated algorithms already exist for solving such systems. Past experience (cf. Reppe [56], Trefethen [62]) has shown that this approach succeeds in nearly all experiments not dominated by crowding. In the present work we apply the theory of iterative solutions for nonlinear systems to the Schwarz-Christoffel parameter problem, and show that with the systems which have been used in practice a robust algorithm will always converge locally, but may not converge globally. A specific example is given for which Trefethen's algorithm diverges from some distant starting points. Finally, we discuss several approaches which may in time yield a globally convergent algorithm.

- **Circular Arc Polygons**

  The final topic we examine, in Chapter 5, is the conformal mapping of circular arc polygons via the Schwarzian differential equation. Though this equation has been known since 1869, there has been only one previous attempt to implement it computationally, that of Bjørstad and Grosse in 1987. Difficulties due to nonlinearity and ill-conditioning are much more severe with this problem than with the Schwarz-Christoffel integral; Bjørstad and Grosse report that their code is generally slow to converge and sometimes fails even on relatively simple polygons. While we have not been able to entirely overcome these difficulties, we have made a number of improvements and identified several sources of error which were previously unrecognized. These improvements include a new method for integrating the differential equation, a new set of accessory parameters which improves the conditioning of the nonlinear system and also tends to yield better starting points for the iteration, new approaches to the formulation of the nonlinear system, including the problem of normalization, which avoid some of pitfalls in Bjørstad and Grosse's approach, and formulas for circular arc maps from the unit disk and from an infinite strip. Our eventual goal, not yet realized, is to apply the techniques discussed in Chapter 4 to derive a nonlinear system with which convergence can be guaranteed even from very distant starting points.

Chapter 2

Highly Elongated Polygons[1]

Many algorithms currently exist which can dependably compute Schwarz-Christoffel transformations for a wide variety of polygons (see, for example, [9], [55], [56], [62]). These methods fail, however, on some seemingly very simple examples. They cannot, for example, map a rectangle with an aspect ratio of only 20 to 1, or most other regions with a similar degree of elongation. The reason for this is an intrinsic property of conformal maps that goes by the name of the "crowding phenomenon" in the literature of numerical conformal mapping (see [11], [18], [27], [35], [42], [65], [69]). Whenever a disk or half-plane is mapped to an elongated region, some of the prevertices are located exponentially close together. For aspect ratios beyond 10 or 20 some groups of prevertices are likely to merge together in finite precision arithmetic, making the evaluation of (1.2) effectively impossible.

Fortunately, many of the distorted regions that come up in applications are highly elongated in only one direction. Indeed, the goal in such problems is often to map the region onto a channel or rectangle for purposes of grid generation or to obtain an exact or simplified solution; a disk or half-plane is introduced only as an intermediate step. The purpose of this chapter is to show that in such cases, the problem of crowding can be largely eliminated by dispensing with the intermediate domain and mapping directly

[1]The material in this chapter is adapted from [38].

from an infinite strip, which can be easily transformed to a rectangle if desired. In the language of numerical analysis, constructing the conformal map from a strip to an elongated polygon is often a well-conditioned problem, but conventional algorithms for it are unstable because they depend upon the solution of an ill-conditioned subproblem. Our algorithm is stable because it avoids the subproblem.

The capabilities described in this chapter are summarized in Figure 2-1. Many polygons, those without extreme elongation in any direction, can be mapped by the standard Schwarz-Christoffel methods embodied in SCPACK [63] (Figure 2-1a). Polygons that are highly elongated in one direction (as well as those that are not elongated) can be mapped by the modified Schwarz-Christoffel methods described here (Figure 2-1b). For polygons elongated in several directions, different methods will be required (Figure 2-1c).

The basis of our algorithm is a formula similar to (1.2) for mapping an infinite strip onto an arbitrary polygonal channel, and a variation of this formula for mapping a rectangle onto a closed polygon. These formulas are not essentially new; the first one, at least, was described by Woods in [68] and has been used by Davis and Sridhar [9,59,60,61] and by Floryan [21,22] for generating computational grids for internal flow problems. The present work differs, however, in emphasizing the crowding phenomenon and in considering the mapping of rectangles to closed polygons as well as other variations. Our solution method is also more robust, and permits the mapping of more complicated regions than those attempted previously.

2.1 The Crowding Problem

The phenomenon of crowding began to be widely recognized as an obstacle to numerical conformal mapping around 1980; the term "crowding" itself is due to Menikoff and Zemach [42]. To illustrate this phenomenon, we will examine a simple example which can be treated analytically. The Jacobian elliptic function $\mathrm{sn}(z|m)$ maps the rectangle with corners $-K$, K, $K+iK'$, and $-K+iK'$ to the upper half-plane, with the images

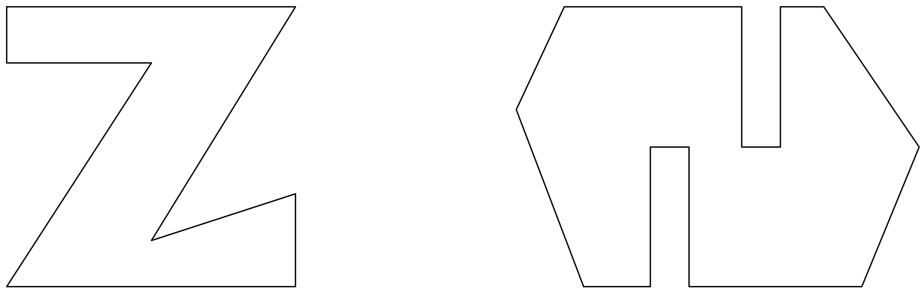

Figure 2-1a: Polygons which can be mapped by the standard Schwarz-Christoffel transformation.

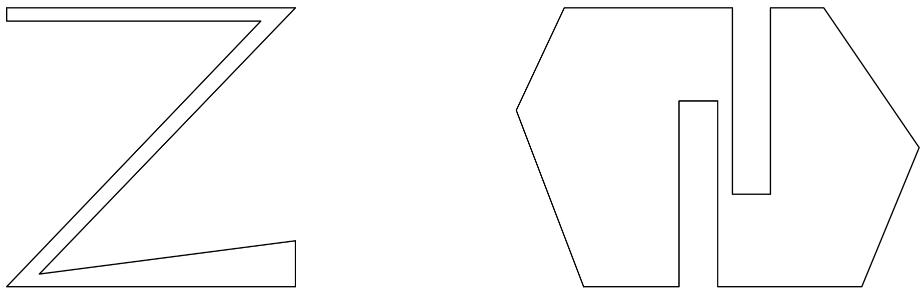

Figure 2-1b: Polygons which can be mapped by the strip transformation.

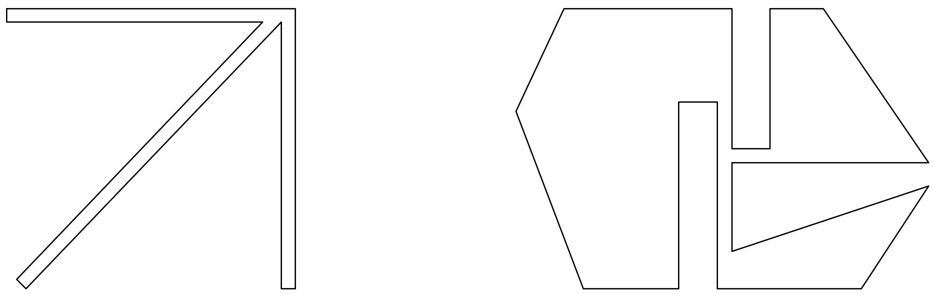

Figure 2-1c: Polygons which cannot be mapped by either method.

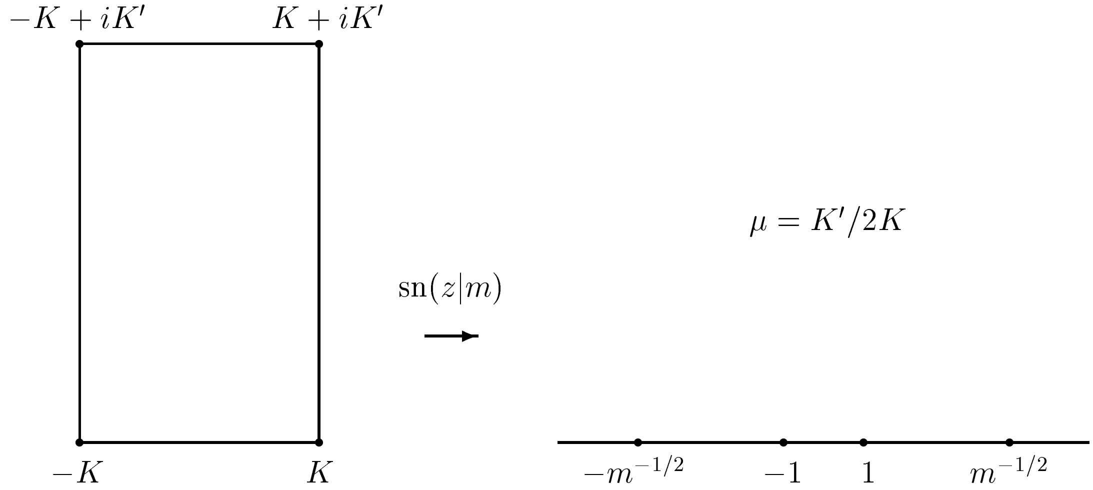

Figure 2-2: Conformal map from a rectangle to the upper half-plane.

of the corners being ± 1 and $\pm m^{-1/2}$ (Figure 2-2). The constants K and K' are complete elliptic integrals with parameters m and $m_1 = 1 - m$, respectively, so only one of K, K', and m can be specified independently. A summary of the properties of elliptic functions and elliptic integrals, including all of the material used in this thesis, can be found in [1].

The conformal modulus μ of the rectangle is $K'/2K$, and $m^{1/2}$ can be used as a measure of the crowding effect since it is the ratio of the smallest to the largest length scales in the upper half-plane. When μ is large the following asymptotic relationships hold:

$$K \sim \frac{\pi}{2}, \quad K' \sim \pi\mu, \quad m^{1/2} \sim 4e^{-\pi\mu}. \tag{2.1}$$

Thus for a modulus of only $\mu = 10$, for instance, the important length scales in the upper half-plane vary by factors on the order of 10^{13}.

It might seem that these difficulties are surmountable, since in this example all of the crowding occurs near the origin. In floating-point arithmetic the four numbers ± 1 and $\pm m^{-1/2}$ can all be represented to full accuracy, even though the the first two may be many orders of magnitude smaller than the last two. There are several reasons, though, why this is an unsatisfactory approach. First, many computers have a range of

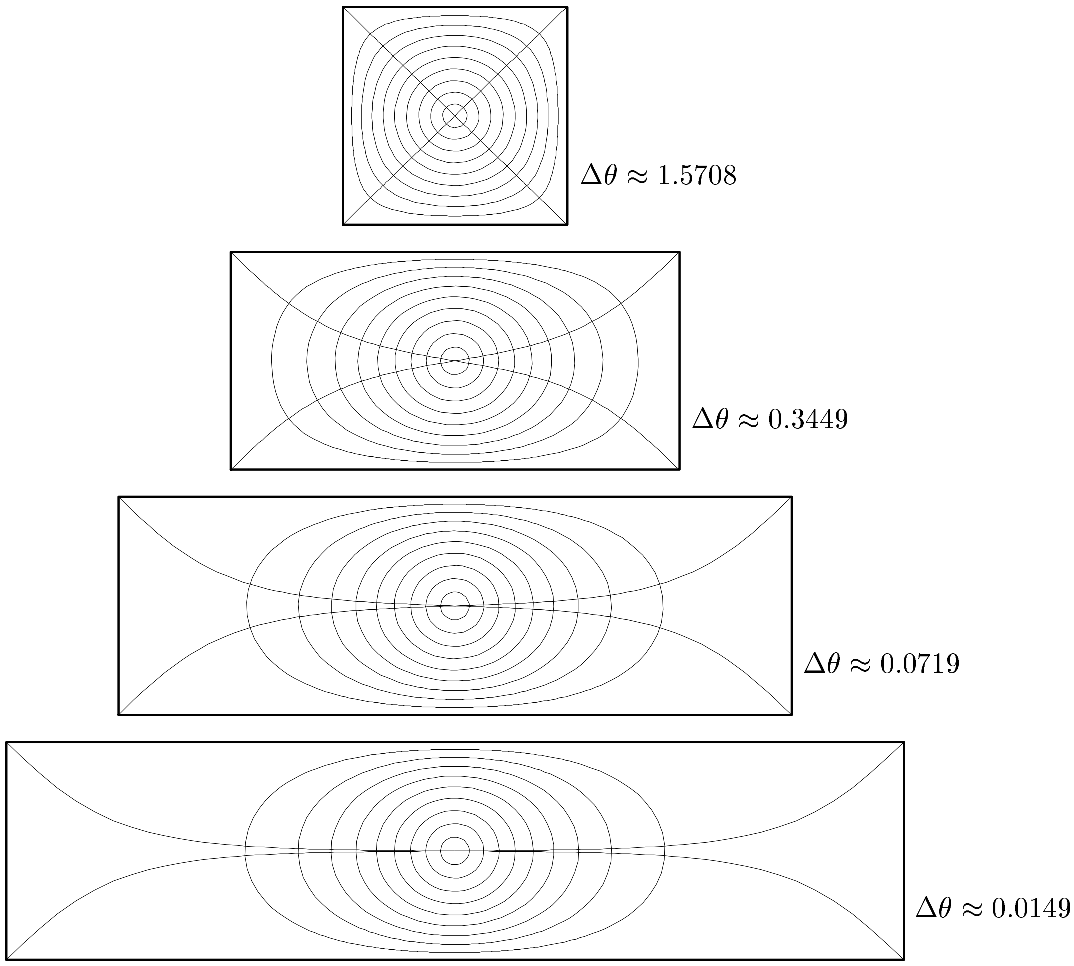

Figure 2-3: Conformal maps of the unit disk onto four rectangles.

permissible exponents too restricted to deal with conformal moduli greater than 25 or 50. Second, such a highly distorted mapping would be useless in many applications. Finally, the direct numerical evaluation of the integral (1.2) in such a situation would require a highly specialized quadrature algorithm. A natural first step in such an algorithm would be to change the variable of integration to $\log(z')$—which is in fact equivalent to using the strip transformation.

Figure 2-3 shows the crowding effect in the conformal map from the unit disk to a mildly elongated region. The four rectangles have moduli 1, 2, 3, and 4, and the internal lines shown are the images of radii to the four prevertices and of equally spaced concentric circles in the unit disk. Due to the conformal nature of each mapping, the innermost circle in each plot is nearly similar to the original disk, so in effect these figures show both the domain and range of each transformation. In particular, we can see the relative positions of the prevertices, and their angular separations are listed in the figure. In the bottom plot, with a conformal modulus of only 4, each of the two pairs of prevertices appears to the eye as a single point. For moduli three or four times larger than this the pairs fuse together in floating-point arithmetic, and the computation by standard methods becomes impossible.

It should be emphasized that crowding occurs when any portion of a domain is elongated. If the rectangles in Figure 2-3, for instance, were all extensions of a larger region to the left, then each one would still experience a crowding effect like that shown in the figure. A strongly acute outward-pointing corner can cause a similar problem (for a mild example, see Figure 2-12b, below). The methods considered in this chapter do not eliminate these secondary crowding effects, which degrade the local accuracy of the mapping and may in extreme cases cause the solution method to fail. In many cases, however, high accuracy can still be obtained in the remainder of the domain.

Crowding is a useful concept for guiding intuition, but it is difficult to give a precise definition which makes sense in all cases. The most straightforward approach is to characterize crowding by large variations in the derivative of the mapping function throughout the domain. In this sense, though, any outward-pointing corner, however obtuse, must be considered crowded, in that the derivative becomes infinite at such points. In practice, the singularities present at most isolated corners are not strong enough to cause serious problems in a computation. Difficulties only arise when a corner is extremely acute, or when, as in the rectangles of Figure 2-1, two or more nearby corners combine to yield a composite effect similar to that of a strong isolated singularity.

2.2 The Strip Transformation

In an intuitive sense, crowding is likely whenever the two regions being mapped are too dissimilar in their relative proportions. With polygons which, like those in Figure 2-1b, are highly elongated in only one direction, it therefore seems reasonable to try to match the geometry of the problem through the use of an elongated standard domain. Two possibilities immediately present themselves: A rectangle seems the most obvious choice for mapping closed polygons, while an infinite strip makes more sense for mapping the open-ended channels which often appear in fluid mechanics problems. In this section and the next we will discuss the derivation and computational implementation of a modified Schwarz-Christoffel formula for directly mapping an infinite strip to an arbitrary polygonal channel. Section 2.4 concerns an extension of this method to maps between rectangles and closed polygons.

All previous derivations of the strip transformation that we are aware of, such as those in [68] and [9], proceed by a direct change of variables in the standard Schwarz-Christoffel formula (1.2). The exponential function $e^{\pi z}$ maps an infinite strip of width 1 onto the upper half-plane, so to directly express the mapping function from the strip to an arbitrary polygon we simply compose the exponential with equation (1.2) to obtain

$$f(z) = A\int^{z} \prod_{j=1}^{n} (e^{\pi z'} - e^{\pi z_j})^{\alpha_j - 1} d(e^{\pi z'}) + B. \tag{2.2}$$

Before simplifying (2.2), let us first introduce one minor complication. The geometry of an infinite strip is clearly analogous to that of many other internal flow configurations in fluid mechanics. The most straightforward application is therefore that of finding conformal maps to infinite channels, i.e., elongated polygons with both endpoints at infinity. It makes sense to map the ends of the infinite strip to the ends of such a channel, so two prevertices are constrained at the outset and should be given special treatment in the mapping function. The positive (outflow) end of the strip maps to infinity in the upper half-plane and thus can be omitted from the formula altogether, as we discussed in the last chapter. The negative (inflow) end maps to zero—let us call this prevertex number 0 to distinguish it from the n prevertices which map to finite corners. Figure 2-4 illustrates the key features of this type of map: The two ends of the channel are implicitly constrained, so only one prevertex, z_1, can be held fixed. The remaining $n-1$ prevertices, on both sides of the strip, must be determined as part of the solution process.[2]

Moving the inflow factor out of the main product, we have

$$f(z) = A \int^{z} \exp[\pi(\alpha_0 - 1)z'] \prod_{j=1}^{n} (e^{\pi z'} - e^{\pi z_j})^{\alpha_j - 1} d(e^{\pi z'}) + B, \tag{2.3}$$

which, with a bit more rearranging of factors, becomes

$$\begin{aligned} f(z) &= A \int^{z} \pi \exp\left\{\frac{\pi}{2}\sum_{j=1}^{n}(\alpha_j - 1)z_j\right\} \exp\left\{\pi\left[\alpha_0 + \frac{1}{2}\sum_{j=1}^{n}(\alpha_j - 1)\right] z'\right\} \\ & \prod_{j=1}^{n} 2^{\alpha_j - 1}\left\{\frac{\exp\left[\frac{\pi}{2}(z' - z_j)\right] - \exp\left[-\frac{\pi}{2}(z' - z_j)\right]}{2}\right\}^{\alpha_j - 1} dz' + B. \end{aligned} \tag{2.4}$$

This hardly seems to be an improvement, but several of the factors are now redundant and can be removed. All constant factors, including the entire first exponential function, can be incorporated into the complex constant A. Using the fact that the turning angles of the polygon, including those at infinity, must sum to 2π, the second exponential can be simplified to $\exp[(\theta_+ - \theta_-)z'/2]$, where $\theta_- = -\pi\alpha_0$ is the divergence angle at the left end of the channel and θ_+ is the corresponding divergence angle at the right end. Finally, replacing each difference of exponentials with a hyperbolic sine, we have

$$f(z) = A \int^{z} \exp\left[\frac{1}{2}(\theta_+ - \theta_-)z'\right] \prod_{j=1}^{n} \left\{\sinh\left[\frac{\pi}{2}(z' - z_j)\right]\right\}^{\alpha_j - 1} dz' + B, \tag{2.5}$$

the simplest form of the strip transformation.

[2] A word about notation: Throughout the remainder of this thesis, whenever it is important to distinguish between the correct target polygon with its corresponding (unknown) prevertices, and an incorrect image polygon mapped from a guessed set of prevertices, we identify the correct values with asterisks: z_j^*, w_j^*, etc. Also, in our discussion of the strip transformation we number the prevertices counterclockwise around the strip starting from the lower left, and use n' to denote the number of prevertices on the lower side of the strip.

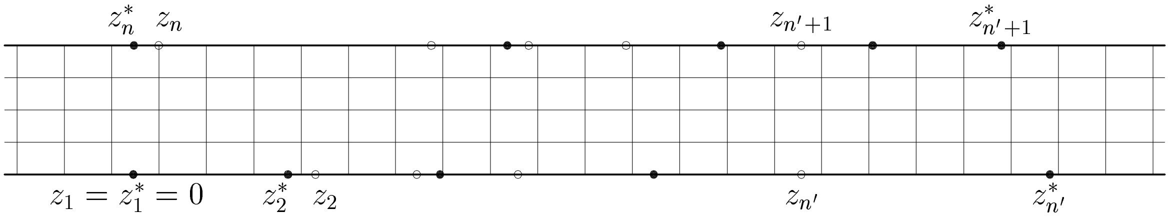

Figure 2-4a: Prevertices on the strip—solid dots show the correct values z_j^*, open circles show an incorrect set of values z_j.

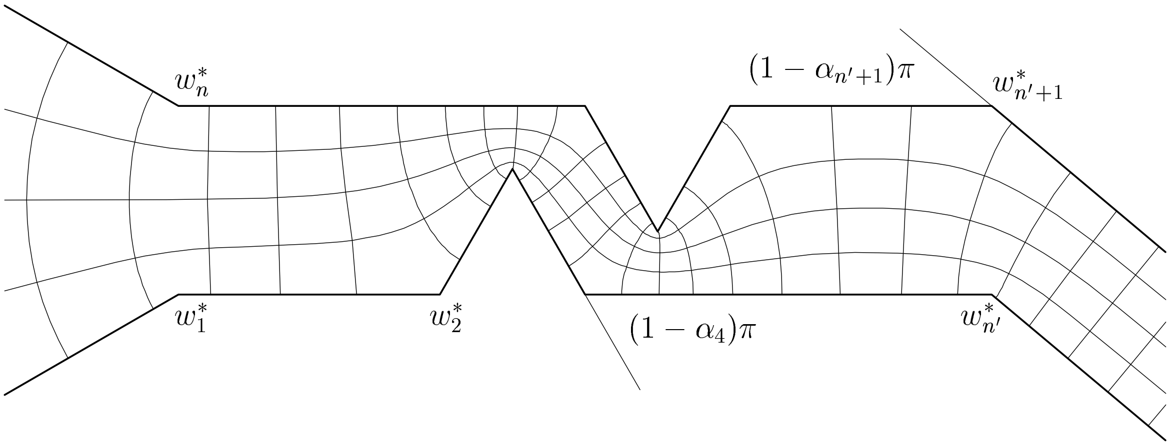

Figure 2-4b: Target polygon P^* defined by vertices w_j^*.

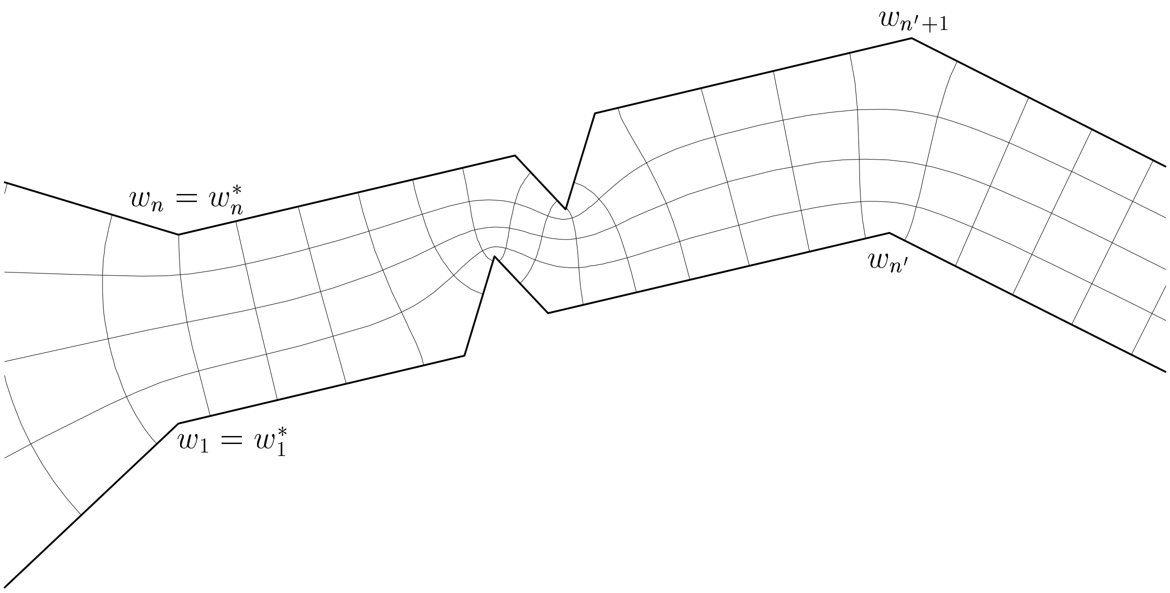

Figure 2-4c: Polygon P defined by incorrect vertices w_j.

This derivation by direct substitution has the advantage of not requiring any sort of additional proof—the strip formula is valid simply because the Schwarz-Christoffel formula is valid. There is no explanation, however, for why the mapping function takes the particular form (2.5), so similar in many ways to the standard formula (1.2). We therefore present an alternate derivation from first principles which, though less rigorous, provides a better heuristic interpretation of each factor in (2.5).

One of the basic ideas we discussed in Chapter 1 was the interpretation of the Schwarz-Christoffel transformation in terms of its elementary factors. In order for a mapping function of the form

$$f(z) = A \int^{z} \prod_{j=1}^{n} f_j(z')dz' + B \tag{2.6}$$

to map any standard region onto a polygon, the integrand must map the boundary of the region in a piecewise fashion onto lines of constant argument. The most straightforward way to accomplish this is to have each elementary factor itself map the boundary onto such lines. With the upper half-plane as the standard domain the necessary angles can be introduced by a simple power function. With the strip, however, the side of the boundary opposite a given prevertex raises an additional complication. First we must find an intermediate map sending a prevertex z_j to the origin and folding the opposite boundary onto a single line of constant argument; only then, by raising this function to a power, can we create an angle in one side without disturbing the other. The appropriate intermediate map is, in fact, the hyperbolic sine function, as it appears in (2.5).

In practice we use two slightly modified versions of these elementary factors in order to conveniently handle branch cuts. For those prevertices z_j on the lower side of the strip,

$$f_j(z) = \left\{-i \sinh\left[\frac{\pi}{2}(z - z_j)\right]\right\}^{\alpha_j - 1}, \quad 1 \le j \le n', \tag{2.7}$$

and for z_j on the upper side,

$$f_j(z) = \left\{-i \sinh\left[-\frac{\pi}{2}(z - z_j)\right]\right\}^{\alpha_j - 1}, \quad n' + 1 \le j \le n. \tag{2.8}$$

The effect of one of these factors is shown in Figure 2-5. The changes are mathematically

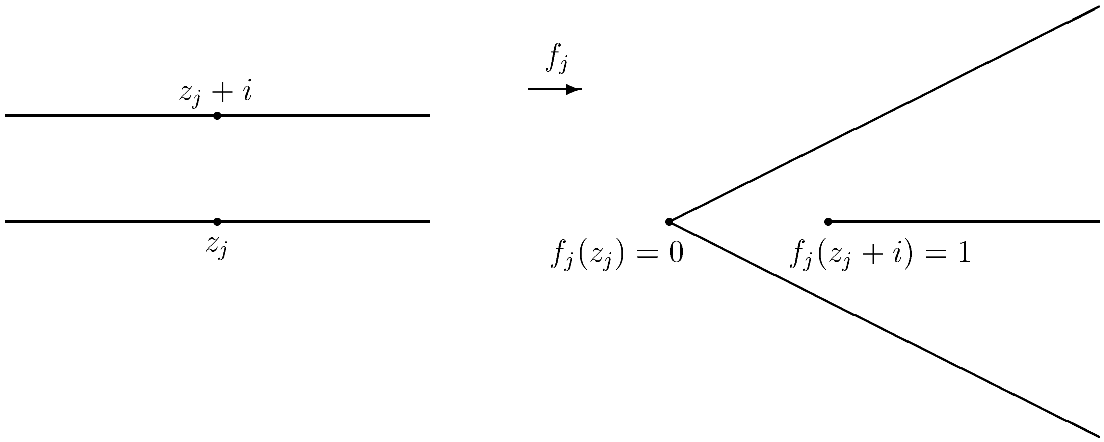

Figure 2-5: A single factor f_j, shown for $\alpha_j > 1$.

trivial since the constant factors can simply be absorbed into the complex constant A. The point is that in this form all of the branch cuts of the f_j are directed away from the strip. (Conventionally, and in Fortran, these branch cuts are located on the negative real axis.)

The functions f_j defined here introduce the required angles at the vertices on both sides of the channel, but they always produce equal divergence angles at $\pm\infty$. As with the other derivation, the additional factor

$$f_0(z) = \exp\left[\frac{1}{2}(\theta_+ - \theta_-)z\right] \tag{2.9}$$

provides the necessary adjustment. In this case, however, we can see the actual role f_0 plays in the mapping function: It changes the argument of the integrand on the upper side of the strip, while leaving the lower side unaffected. (The final orientation of the image channel is of course determined by the constant A.) The full strip transformation is thus given by

$$f(z) = A\int^z \prod_{j=0}^{n} f_j(z')dz' + B, \tag{2.10}$$

where the individual functions f_j are defined by equations (2.7–2.9).

2.3 Numerical Computations

All implementations of the Schwarz-Christoffel transformation and related methods, including the strip transformation (2.10), must face two main numerical obstacles: evaluation of singular integrals and solution of the parameter problem. The integral in (2.10) has prevertex singularities that are essentially identical to those in the standard formula (1.2); as we will discuss the various merits of several numerical integration techniques in Chapter 3, there is no need for a separate explanation of the methods here. At a theoretical level the two parameter problems are also quite similar. It seems appropriate to describe briefly some implementational details of our algorithm, however, since much of this material is specific to the strip transformation. Justification for our procedure, and some alternate approaches to the parameter problem, will appear in much greater detail in Chapter 4.

We view the parameter problem as a general system of nonlinear equations $\mathbf{F}(\mathbf{x}) = \mathbf{0}$. Two of the map's three degrees of freedom are accounted for by the ends of the channel; the third we fix by placing z_1 at the origin (see Figure 2-4a). One obvious choice for the $n-1$ independent variables might be the real coordinates $\mathrm{Re}(z_2), \ldots, \mathrm{Re}(z_n)$ of the remaining prevertices. This choice leads to a constrained system, however, since the prevertices on each side of the strip must appear in the proper order. To remove the constraints, we use a change of variables similar to the one in SCPACK:

$$x_j = \begin{cases} \mathrm{Re}(z_n), & j = 1, \\ \log(z_j - z_{j-1}), & 2 \le j \le n', \\ \log(z_j - z_{j+1}), & n' + 1 \le j \le n - 1. \end{cases} \tag{2.11}$$

As for the dependent variables, we first compute the positions of $w_1, \ldots, w_n$ by integrating (2.10) and using the constants A and B to fix w_1 and w_n at their correct positions. The

$n-1$ functions to be set to zero are then given by

$$F_j = \begin{cases} \mathrm{Im}\left[\log\left(\dfrac{w_2 - w_1}{w_2^* - w_1^*}\right)\right], & j = 1, \\ \mathrm{Re}\left[\log\left(\dfrac{w_j - w_{j-1}}{w_j^* - w_{j-1}^*}\right)\right], & 2 \le j \le n', \\ \mathrm{Re}\left[\log\left(\dfrac{w_j - w_{j+1}}{w_j^* - w_{j+1}^*}\right)\right], & n'+1 \le j \le n-1, \end{cases} \tag{2.12}$$

where the w_j^* are the correct corner positions and the w_j are the current approximations. F_1 measures the orientation of the entire figure in space—this is not fixed by the constant A since the line from w_1 to w_n is not a side of the polygon. Each of the other F_j involves the logarithm of a side length. The use of logarithms improves the scaling of the problem when some sides are much longer than others, as often occurs with elongated regions.

We have experimented with three nonlinear equations packages for solving this system: Powell's subroutine NS01A [53], the Minpack routine HYBRD [44], and an implementation of Schnabel's pseudo-code from Dennis and Schnabel [14]. All three are based on a hybrid (dogleg) quasi-Newton algorithm with secant updates. On average we obtained slightly better results with HYBRD, and in addition one of our test problems caused Powell's code to fail. (This difficulty was apparently due to an overly strict stopping criterion rather than a fundamental failure of the algorithm.) Some of our test problems were quite difficult, involving extremely distorted polygons like those shown later in the chapter; such problems sometimes required several hundred evaluations of the functions (2.12). Note that a channel which has sharp bends and odd corners should not necessarily be considered highly distorted. A region like that in Figure 2-6, which is essentially a bent strip, is a bit poorly scaled but does not cause major difficulties. The only cases where either HYBRD or Schnabel's code failed eventually to find a solution involved severe crowding in regions which were elongated in more than one direction, as in Figure 2-1c. Though for individual problems there were sometimes wide variations in the number of iterations required by the different routines, all three gave fairly similar

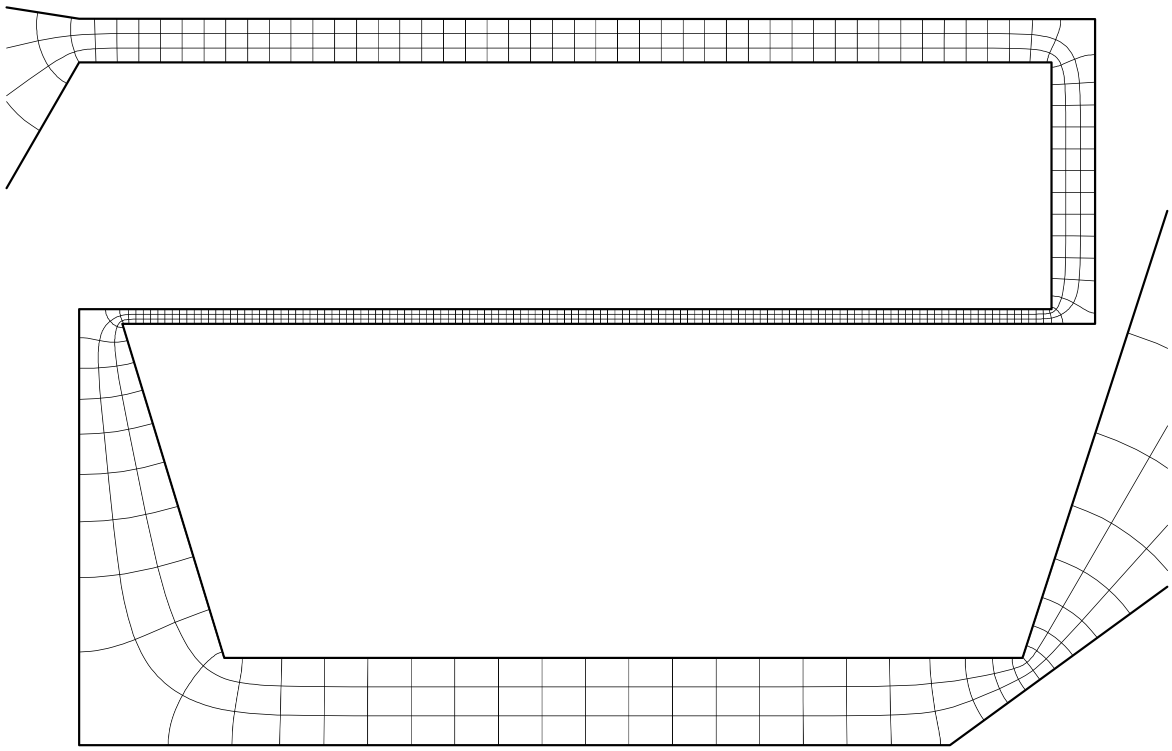

Figure 2-6: A moderately difficult elongated channel, Jacobian condition number approximately 35000.

performance when averaged over a number of different cases.

Dias [15] and Bjørstad and Grosse [6] have used other nonlinear equations packages for other Schwarz-Christoffel problems, with similar results.

The secant algorithms converge superlinearly once they are near the solution (see [14]), but for difficult problems may take a long time to get near it. Convergence times seem to be nearly independent of the starting point, a clear indication that we do not have a good algorithm for picking starting points. To examine the typical convergence rate as a function of n, at least two different types of problems should be considered, as shown in Figure 2-7. In the first example the geometry becomes progressively more complicated as n increases, while in the second the geometry is roughly constant, and increasing n merely improves the resolution of the curved part of the boundary. (There are better ways to approximate curved boundaries; see [21] and [61].) In practice, we find that for problems

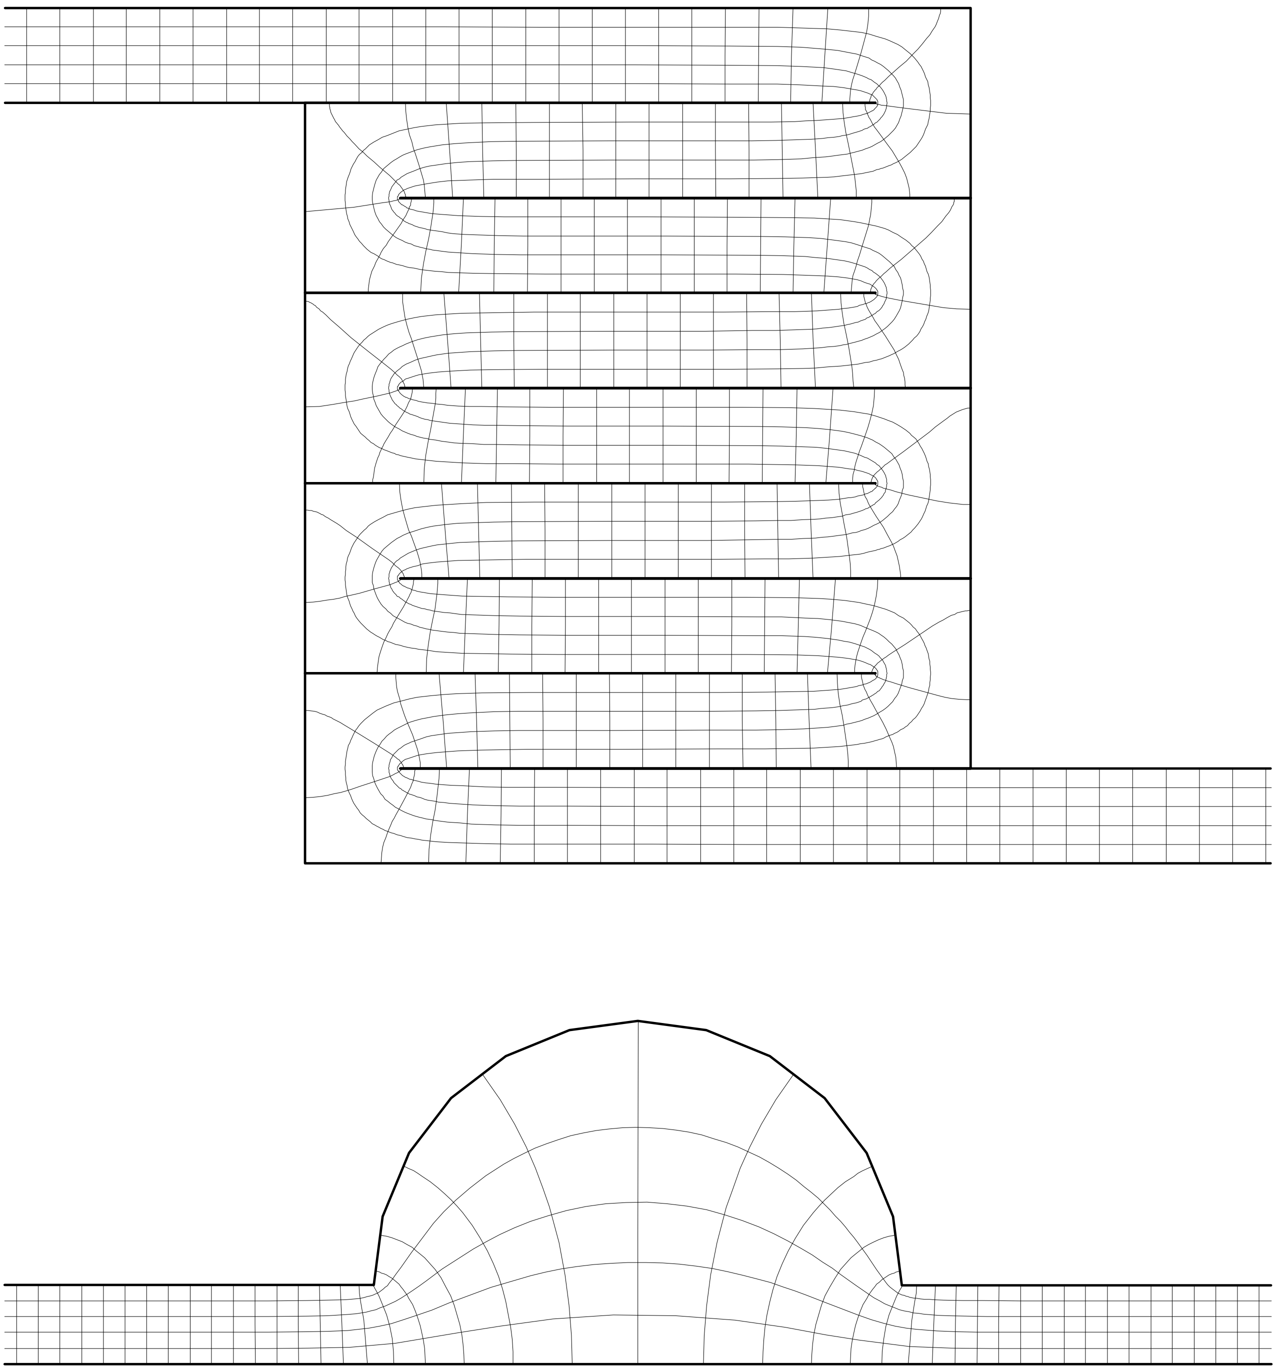

Figure 2-7: Two channels with large n.

of the first kind the number of iterations required is roughly $\mathrm{O}(n)$, while for problems of the second kind it is closer to $\mathrm{O}(1)$. Since each evaluation of (2.12) requires $\mathrm{O}(n^2)$ operations, and the evaluation of the initial Jacobian matrix by finite differences requires n evaluations, the total work required to solve the parameter problem is at least $\mathrm{O}(n^3)$ in both cases.[3] For n less than about 50 the Jacobian evaluation is not the dominant factor in the calculation, however, so problems with simple geometries typically display behavior closer to $\mathrm{O}(n^2)$.

2.4 Mapping Rectangles to Closed Polygons

When the target domain is a closed polygon rather than an infinite channel, it is often appropriate to take an elongated rectangle as the standard region rather than an infinite strip.[4] The aspect ratio of the rectangle will be equal to the conformal modulus ($\approx$ electrical resistance) of the original polygon with its four distinguished vertices. We calculate this conformal map by mapping first from the rectangle to the strip by means of an elliptic function, then from the strip to the polygon by the strip transformation.

The function $s(z) = \frac{1}{\pi} \log \operatorname{sn}(z|m)$ maps a rectangle onto a strip of width 1, sending the corners $-K$, K, $K + iK'$, and $-K + iK'$ of the rectangle to the points i, 0, L, and $L + i$, respectively, where K, K', and L are all functions of m (Figure 2-8). It seems reasonable to require each corner of the rectangle to map to a vertex of the polygon, and in many applications this is what is called for since the conformal modulus is required. However, the formulation of the parameter problem which we used in the previous section would not permit this since in general no two prevertices have the same real part. We can avoid this difficulty, though, since for a closed polygon we no longer need to specify the images of the ends of the strip. The two extra degrees of freedom thus obtained can

[3] The linear algebra required by the secant algorithm can be held to $\mathrm{O}(n^2)$ per iteration by using secant updates (see [14]), and is typically negligible compared with the cost of evaluating (2.12).

[4] For other approaches to conformal mapping onto rectangles, see the work by Gaier [27,28] and by Papamichael et al. [48,49,50,51].

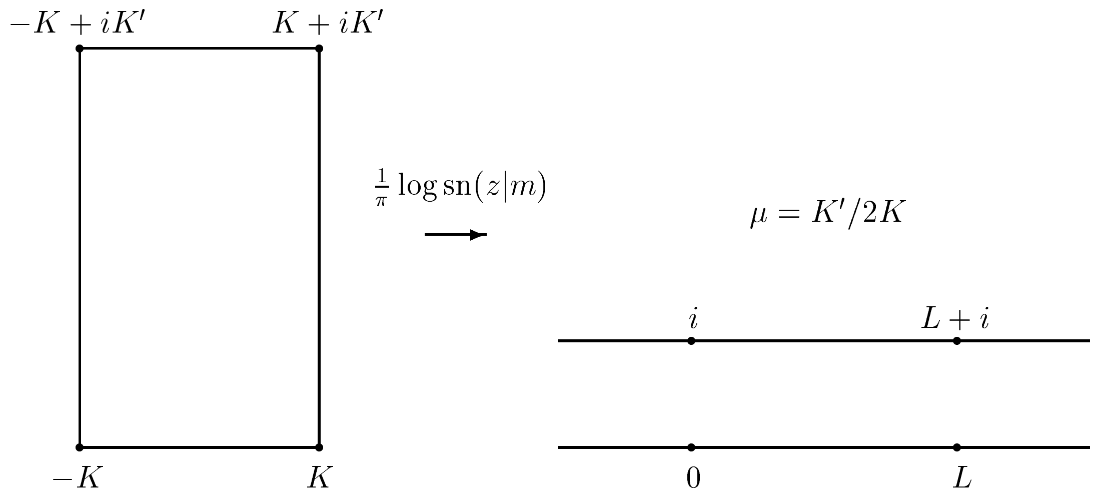

Figure 2-8: Conformal map from a rectangle to a strip.

be used instead to fix z_n at i and require that the two rightmost prevertices have the same real part. Solving the parameter problem with these restrictions thus gives us an appropriate value for L, from which m, K, and K' can be calculated using the known properties of elliptic functions. Formulas for all of our computations involving elliptic functions are discussed in the Appendix.

In formulating the parameter problem for this version of the strip transformation, we again use the constants A and B in (2.10) to send w_1 and w_n to their correct positions. A suitable set of unconstrained independent variables corresponding to the new normalization is

$$x_j = \begin{cases} \log(z_{j+1} - z_j), & 1 \le j \le n' - 2, \\ \frac{1}{2}[\log(z'_n - z_{n'-1}) + \log(z_{n'+1} - z_{n'+2})], & j = n' - 1, \\ \log(z_{j+2} - z_{j+3}), & n' \le j \le n - 3. \end{cases} \tag{2.13}$$

Note that there are only $n - 3$ independent variables, corresponding to the fact that now only $n - 3$ side length conditions are required to determine the shape of the closed polygon:

$$F_j = \begin{cases} \log \left| \dfrac{w_{j+1} - w_j}{w_{j+1}^* - w_j^*} \right|, & 1 \le j \le k-2, \\ \\ \log \left| \dfrac{w_{j+2} - w_{j+3}}{w_{j+2}^* - w_{j+3}^*} \right|, & k-1 \le j \le n-3. \end{cases} \tag{2.14}$$

Here w_k is the "omitted corner"—if $n-1$ corners of the polygon and all of the angles are known, then the position of the remaining corner is determined and cannot be specified separately. The side between w_1 and w_n is fixed by the constants A and B, and the two sides which intersect at w_k do not enter into the parameter problem, so exactly $n-3$ side lengths are sufficient to determine the shape uniquely.

The best choice of w_k is problem-dependent, and an improper choice can make the system of nonlinear equations much more difficult to solve. In Figure 2-9, for example, w_7 and w_8 would both be poor choices for w_k. Since the two sides which intersect at w_7 are collinear, it is not possible to determine the position of the corner by finding the intersection of the sides. To see why w_8 is bad, picture a slight distortion of this polygon whereby the distance between w_{11} and w_{12} is increased and the horizontal tube is correspondingly narrowed. This narrowing increases the aspect ratio of the tube dramatically, so z_8 and z_9 would be far to the right of their correct positions. This will only affect the dependent variables to a small degree, though, if the distance between w_8 and w_9 is not included in (2.14). The system of nonlinear equations will therefore be poorly scaled, and the algorithm will probably take much longer to converge. A much better choice in this case would be w_{12}, which does not introduce any scaling problems. Several other choices would be equally good, and in fact it is sometimes helpful to change w_k in the middle of the solution process if the nonlinear equations algorithm is making slow progress.[5]

At present, our code leaves the choice of w_k up to the user, but we may be able to automate this in the future. Figure 2-10 shows two examples of regions for which no choice

[5]The system of equations for the channel map can be altered in a similar manner, but the presence of the angle in (2.12) raises complications. One must be careful not to create a set of nonlinear equations with more than one solution.

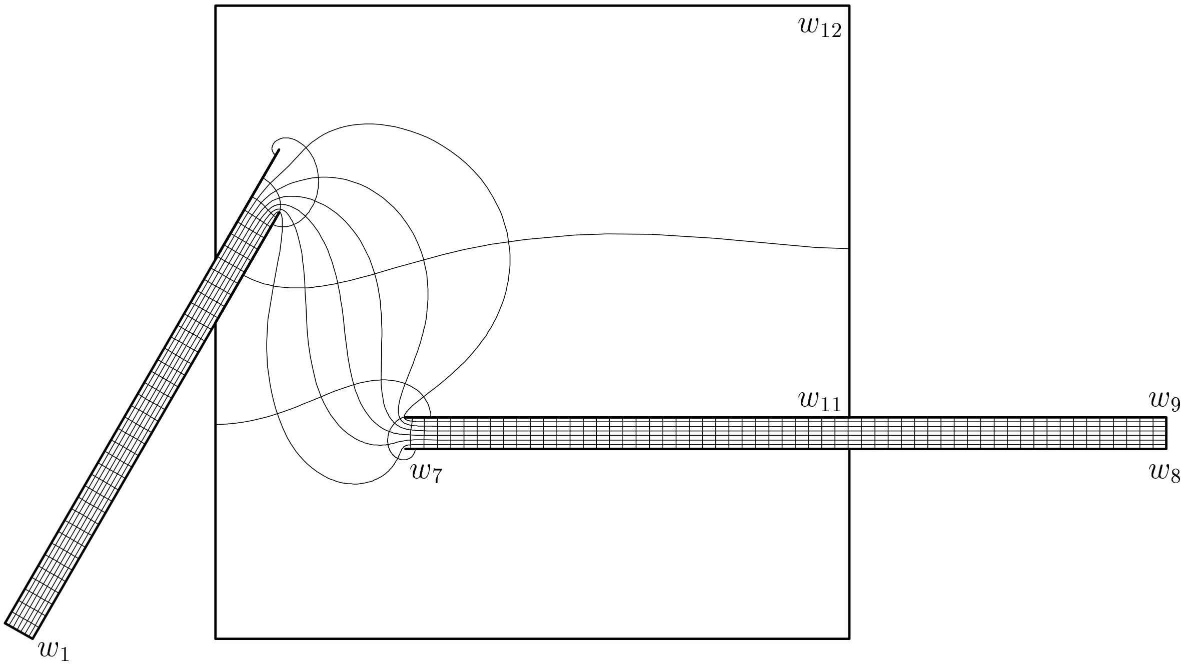

Figure 2-9: Conformal map of a rectangle onto a distorted polygon ($\mu \approx 41.812465$).

of w_k is very good—the only way to make these problems well-scaled would be to devise a radically different set of dependent variables.[6] Note that the self-intersecting nature of the first domain does not itself cause any difficulties; the problem results from the fact that changing almost any side length slightly can greatly alter the conformal modulus of the polygon. Problems with scaling do increase as the target polygon becomes more elongated, but despite this we have succeeded in mapping regions similar to these with moduli of many thousands to one. (We omit illustrations of such regions because they are too thin to show any interesting internal structure.)

As a final example, the polygon in Figure 2-11 is a bit more troublesome—it is not only poorly scaled but also involves a more severe distortion of the preimage rectangle.

[6]It is still possible to calculate accurate conformal moduli for these polygons, however, even though the parameter problems are slow to converge. So that others may reproduce these examples if they wish, in the star the width of the strip is 1/20 of the radius of the circumscribing circle, while in the spiral the width of the strip is exactly half the width of the complementary white space. The regions shown in Figures 2-9 and 2-12b have corners exactly: $\{(-.2887,0.), (.1,.6732), (0.,.4999), (0.,0.), (1.,0.), (1.,.2999), (.3,.3), (1.5,.3), (1.5,.35), (.3,.35), (1.,.3501), (1.,1.), (0.,1.), (0.,.6001), (.1,.7732), (-.3320,.025)\}$ and $\{(0.,8.), (1.8,8.), (1.8,0.), (13.,0.), (14.,0.), (14.,1.), (2.,1.), (2.,10.), (0.,10.)\}$, respectively.

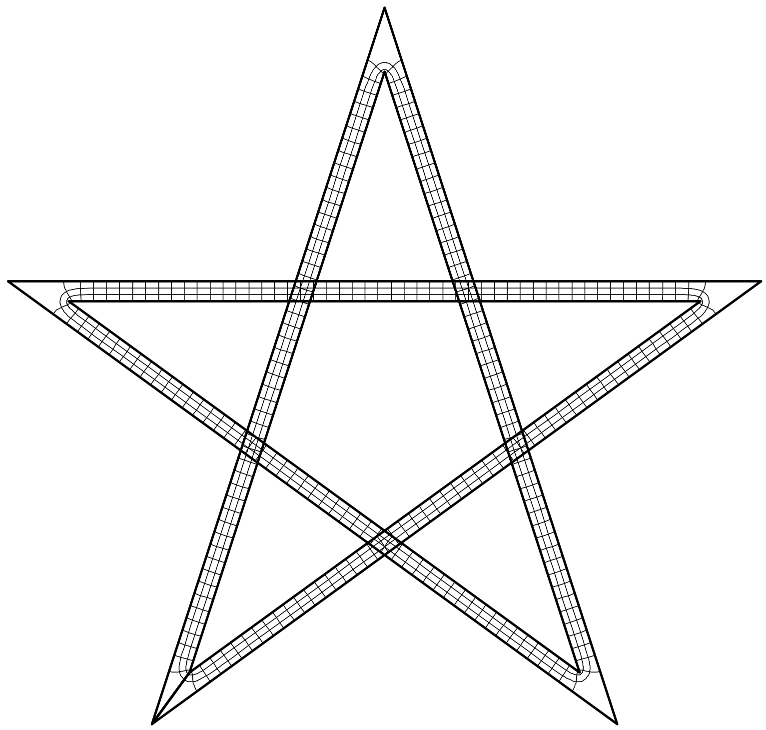

Figure 2-10a: Conformal map of a rectangle onto a star ($\mu \approx 163.28151$). Note that both ends of the polygon meet at the lower-left corner.

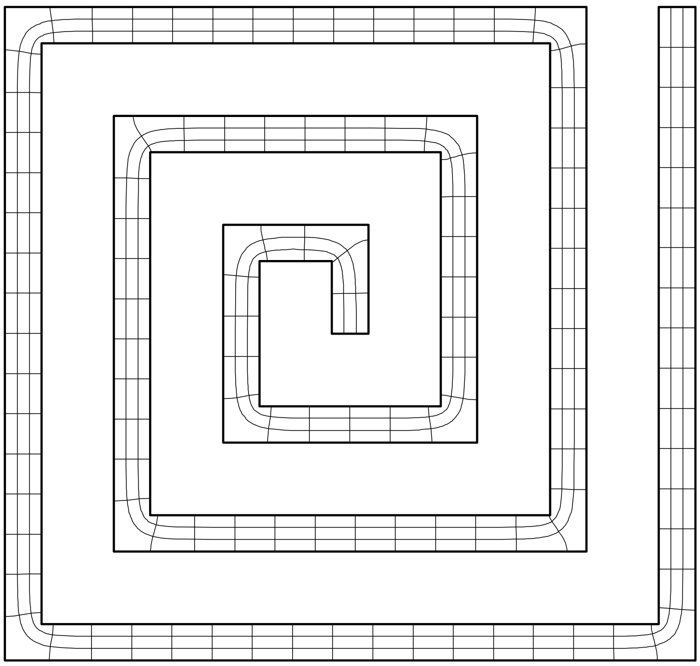

Figure 2-10b: Conformal map of a rectangle onto a finite spiral ($\mu \approx 132.70454$).

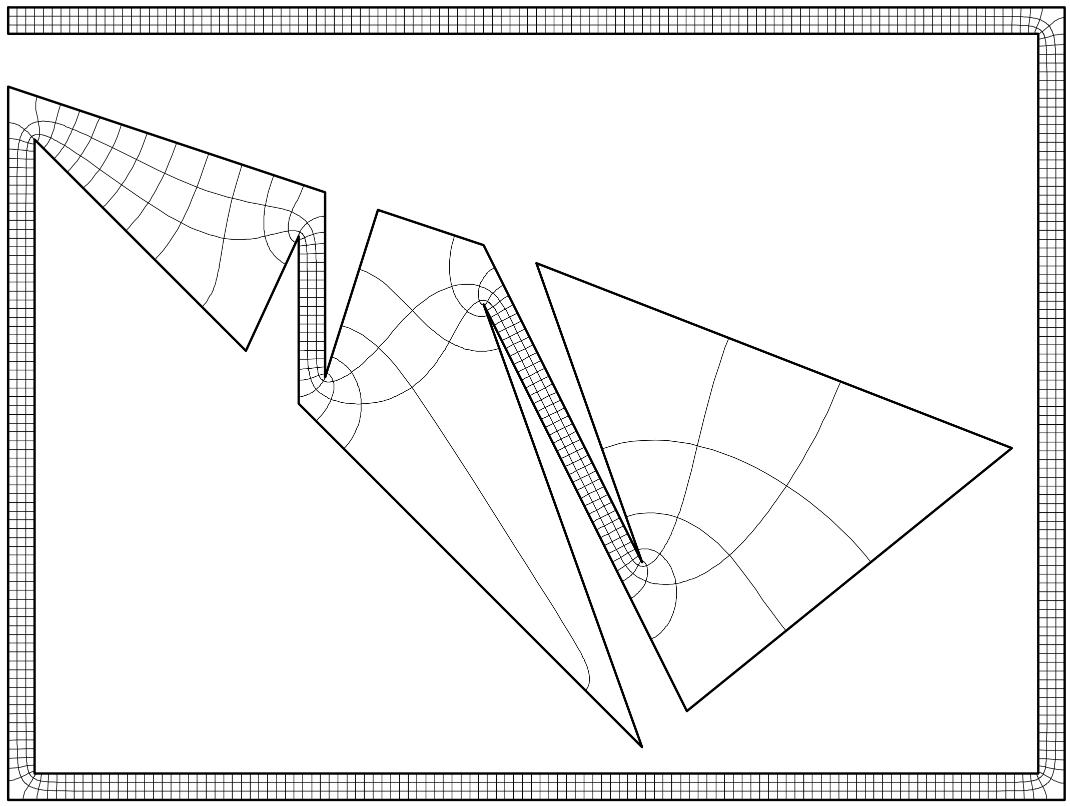

Figure 2-11: A polygon with 23 sides ($\mu \approx 156.6241139$).

Solving the parameter problem for this region to roughly 10-digit precision took about one hour on a Sun 3/50 with an MC68881 floating-point coprocessor. Since the final convergence is quite rapid the time is nearly independent of the required accuracy. The same calculation on a supercomputer would require only a few seconds.

2.5 Variations

The two problem formulations described in the preceding sections illustrate some of the choices which can be made with the strip transformation, but they by no means exhaust the possibilities. With the channel mapping, for instance, it is not necessary for either of the divergence angles θ_- and θ_+ to be positive; the channel may be bounded. Figure 2-12a shows an example where both are negative; physically this could represent an electromagnetic problem with point charges or currents, or a fluids problem with a source and a sink. (This region has significant secondary crowding effects near the point of the barb, due to the extremely acute angle there. It is impossible to plot streamlines much closer to the point than those shown using our present integration techniques. The other angles are wide enough to avoid this difficulty, though there is always some degradation of accuracy near an outward-pointing corner.)

With the rectangle mapping there is no reason why each end of the rectangle must map onto a single side of the polygon, and there are many possible ways to modify the given formulation to deal with unusual cases. In Figure 2-12c, for example, we have fixed w_{n-1} at i instead of w_n, and introduced an additional vertex w_4 with a turning angle of 0. Problems like this one could arise in integrated circuit design.

Other straightforward variations of the method described in this paper include vertices at infinity, the exterior map for a polygon, the map from a semi-infinite strip to a channel bounded at one end, and various generalized parameter problems like those described in [64] and [19].

Figure 2-13 shows a more extreme variation, an infinite logarithmic spiral. To permit our program to run to completion in a finite time, we truncated the infinite product in (2.10) by ignoring corners more than three turns away from the point of interest. This approach yields accurate results since the effect of each singularity decays exponentially along the strip. The parameter problem for this example is also quite simple, since the domain is self-similar: In addition the the normalizing constants A and B only two measurements, one expressing the magnification with each turn and the other the width of each section, are required to specify the geometry completely. A similar formulation was used by Floryan [22] to map periodic channel configurations. The ideas involved in this example might possibly be extended to permit the mapping of more complicated fractal domains, which would have applications, for example, in the study of diffusion-limited aggregation.

Davis and Sridhar [9,61], Floryan [21] and Hoekstra [37] have all worked with another variation of channel maps which supports curved boundaries, based on a formula

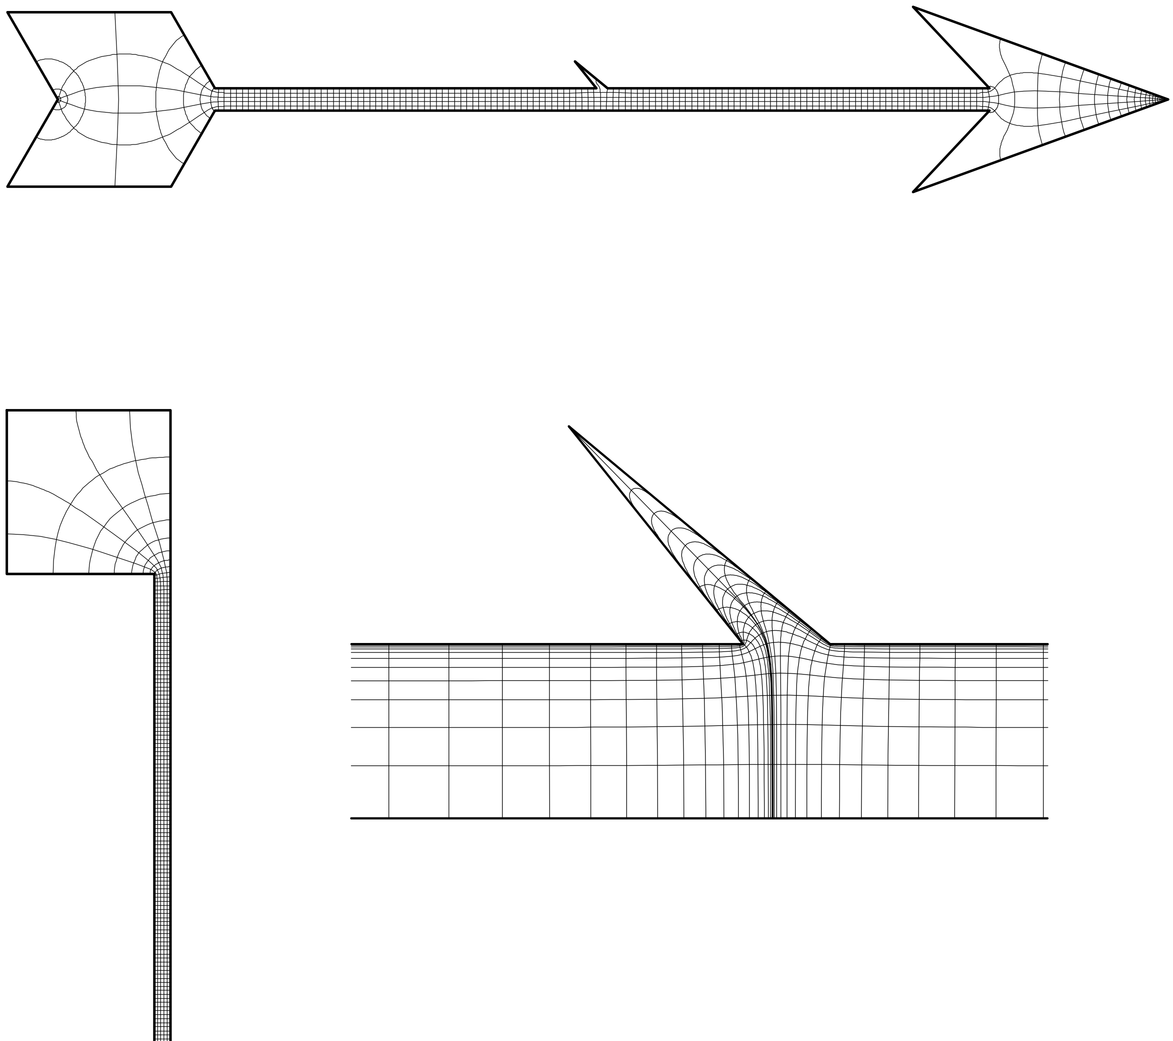

Figure 2-12a: A channel map with converging ends.

Figure 2-12b: Enlargement of the barb in Figure 2-12a.

Figure 2-12c: A rectangle map onto a polygonal conductor ($\mu \approx 49.436547$).

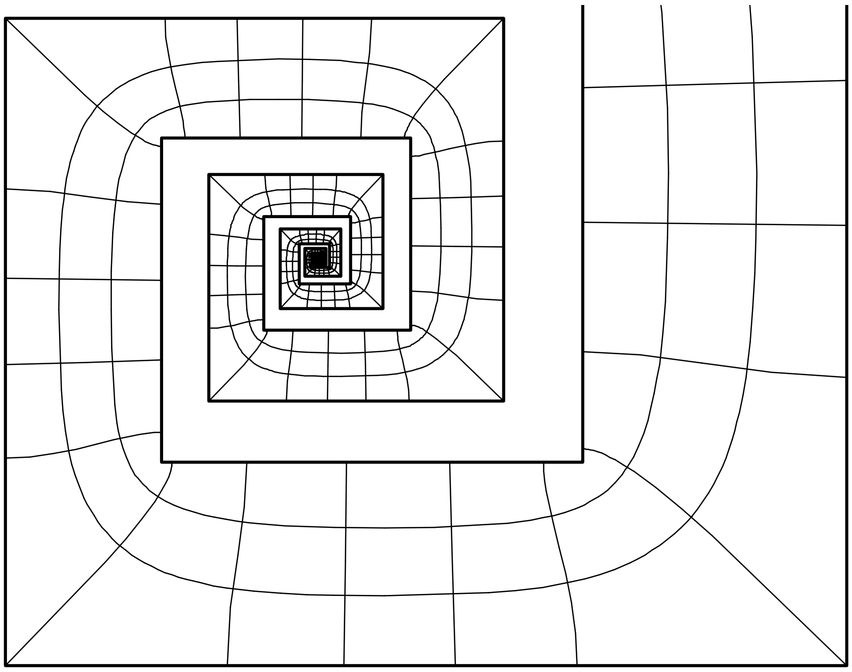

Figure 2-13: Conformal map onto an infinite spiral.

described by Woods [68]. Our implementation does not currently include this variation, but it would certainly be a valuable addition to any future software package for calculating Schwarz-Christoffel maps.

Chapter 3

Evaluating Schwarz-Christoffel Integrals

It is not possible to evaluate the contour integrals appearing in Schwarz-Christoffel transformations analytically, except in the simplest cases involving only two or three vertices. For computational purposes, therefore, some sort of numerical approximation scheme is required. A typical numerical method samples the integrand at several points in an interval and then fits a curve through those points whose integral is known. To increase the accuracy of the approximation, one needs merely to sample at more points. Schwarz-Christoffel integrals present two difficulties for this approach, both involving singularities. First, in the process of solving the parameter problem it is often necessary to integrate between two adjacent prevertices in order to determine the length of a side. Contours with endpoints at prevertices are also useful when computing values of the mapping function. The integrands in these cases are singular at one or both endpoints; this tends to cause trouble for numerical integration schemes based on polynomial fits, and generally requires the use of some method which explicitly takes the singularities into account. Second, two or more prevertices often cluster very close together due to the crowding phenomenon. The result is that even singularities which are not actually on a particular interval are often close enough to it to cause trouble for numerical integrators. The problem is usually so localized that it would be wasteful to refine the mesh over an entire interval, so the most efficient approaches use local refinement techniques based either on the geometry of the problem or the observed behavior of the integrand. Strangely enough, though general adaptive quadrature routines have long been available, their first published use for Schwarz-Christoffel integrals appears to be that of Dias [15] as late as 1986.

We have experimented with several different methods for dealing with both of these difficulties, including singularity removal, three separate adaptive integration procedures, and our own implementation of Trefethen's compound Gauss-Jacobi integration scheme [62]. In this chapter we will describe these methods in detail, summarize the results of our experiments, and compare our methods with some others found in the literature.

3.1 Paths of Integration

Since Schwarz-Christoffel integrals are complex contour integrals, the first question to ask is whether contours can be chosen which avoid the singularities entirely. In their program for solving the Schwarzian differential equation for circular arc polygons, Bjørstad and Grosse [6] integrate through the interior of the domain to points between adjacent pairs of prevertices, rather than to the prevertices themselves. Each of these points fixes a line containing a side, and corners are then identified as the points where sides intersect. This method eliminates the difficulty with endpoint singularities, but does not affect the problem of singularities close to the interval of integration. In addition, it has some disadvantages which we believe make it less practical for our computations:

- The initial guess to a solution to the parameter problem is often wildly inaccurate. We have seen examples involving the strip transformation where some sides were incorrect by factors exceeding 10^{16}. In this case some sides may be reduced to single points in floating point arithmetic, so determining the side length as the distance between two corners may be impossible. Integrating directly from one corner to the next, however, avoids this cancellation problem.

- When the turning angle is at or near an integral multiple of π the position of a corner may be poorly determined as an intersection.
- Integrating through the interior of the domain requires the use of complex arithmetic. Real arithmetic can often be used when integrating along the boundary, for a significant increase in computational speed.

While we have not experimented directly with Bjørstad and Grosse's method, we did try a somewhat similar approach in an attempt to reduce the number of separate integrations. The idea was to integrate only to every other prevertex. This would fix alternate corners in the polygon, and the other corners could then be determined as intersections. To avoid the first two difficulties listed above, we did not skip corners with turning angles near multiples of π, nor did we use the method at all when far from a solution to the parameter problem. The decrease in the number of contours, however, was offset by their increased length and complexity, so there was no great decrease in the number of integrand evaluations to compensate for the use of complex arithmetic. Since this method was also a bit less robust, we abandoned it in favor of integrating to every prevertex. Some experiments involving this approach are described below in Section 3.4.

3.2 Compound Gauss-Jacobi Quadrature

One way to efficiently compute an integral with an endpoint singularity is to use a quadrature rule which explicitly takes the singularity into account. In his Schwarz-Christoffel package SCPACK [62,63], Trefethen uses a compound Gauss-Jacobi (CGJ) method with considerable success. To avoid inaccuracies due to nearby singularities, an interval is first subdivided until the subintervals satisfy the rather arbitrary requirement that no outside singularity may lie closer to any interval than half the length of that interval. A Gauss-Jacobi quadrature formula is then used on each interval ending in a singularity, and a standard Gauss formula is used on all other intervals. The nodes and weights for the Gauss-Jacobi quadrature rules are calculated using the routine GAUSSQ by Golub and Welsch [31].

In our own implementation of the compound Gauss-Jacobi method, we have found experimentally that the number of accurate decimal places in the computed integral is approximately the same as the number of nodes used on each subinterval. This is the same behavior as was reported by Trefethen [62], and holds even for such highly distorted regions as those in Figures 2-9 and 2-11. The primary drawback of the method is that this is entirely an empirical bound. We have not been able to prove that the CGJ approach *must* always give the expected level of accuracy; we have simply observed that it tends to do so in practice.

In order to obtain more reliable error bounds, we have experimented with several general-purpose adaptive quadrature routines. These use algorithms for refining the mesh based on the observed behavior of the integrand itself, and thus should be able to identify difficult areas which might be missed by the CGJ heuristic. As described in the next two sections, some of the adaptive routines give reasonably good performance when appropriate consideration is given to the treatment of singularities on the interval of integration. Unfortunately, even the best routines we have tried tend to be slower than CGJ by factors of two or more. This is because the compound Gauss-Jacobi method estimates errors and chooses an appropriate mesh based on a more complete knowledge of the problem, i.e., the positions of singularities throughout the domain, rather than by making expensive tests and comparisons in the interval of integration itself. In practice we tend to use CGJ the vast majority of the time, switching to an adaptive routine only when a strict error bound is required. If a reliable error estimate could be developed for the CGJ method, the more general adaptive schemes could perhaps be dispensed with entirely. Since most intervals do not require extensive subdivision, however, any error estimate which is much more difficult to compute than a single integrand evaluation would add significantly to the computational cost of the procedure.

3.3 Singularity Removal

There are many general purpose subroutines for adaptive quadrature available, only a few of which use approaches like the Gauss-Jacobi method that explicitly account for known singularities. Even routines which assume polynomial-like behavior can be used on singular problems, though, if the singularities can be suppressed analytically. We have implemented two distinct methods for accomplishing this: In one, singular terms are subtracted off; in the other, a different independent variable is used. Neither method completely eliminates the singularities appearing in Schwarz-Christoffel integrals, but the residual singular terms are of high order and typically have little effect on adaptive integration routines.

Let $f(x)$ be the integrand of a Schwarz-Christoffel problem with a left-endpoint singularity at the origin. There is then a function $g(x)$, analytic near the origin, such that $f(x) = x^{\gamma}g(x)$, $\gamma > -1$ [46, pp. 189–196]. (This is true both for the standard Schwarz-Christoffel transformation and for the strip transformation.) Note that f becomes infinite at 0 for $\gamma < 0$. Expanding $f(x)$ in a power series about 0, we obtain

$$f(x) = x^{\gamma}g(0) + x^{\gamma+1}g'(0) + \frac{1}{2}x^{\gamma+2}g''(0) + \mathrm{O}(x^{\gamma+3}). \tag{3.1}$$

Only the first term of the expansion is unbounded, so we can weaken the singularity by subtracting off this term and integrating it analytically. The remaining part of the integrand, though, can still have an unbounded slope, which can be enough to cause serious trouble for integrators which assume polynomial behavior. In practice, therefore, it is necessary to subtract off the second term as well; the remainder still has an unbounded second derivative, but most polynomial integrators will not require any extra refinement to obtain acceptable error bounds. The use of this two-term subtraction for Schwarz-Christoffel problems dates back at least to Kantorovich and Krylov [40].

It is instructive to examine the ways in which some specific polynomial quadrature rules react to the remaining weak singularities. Consider the function $q(x) = x^{\alpha}$, $\alpha > 0$, which we want to integrate over a subinterval $[0, h]$ of some larger interval $[0, L]$. The exact integral I is $\frac{1}{\alpha+1}h^{\alpha+1}$, and any approximation obtained from a weighted average of function values will be of the form $k(a)h^{\alpha+1}$. For example, the trapezoidal approximation T is $\frac{1}{2}h^{\alpha+1}$, and the Simpson's Rule approximation S is $(\frac{1}{6}+\frac{2}{3\cdot 2^{\alpha}})h^{\alpha+1}$. (Note that $T=I$ when $\alpha=1$, and $S=I$ when $\alpha=1$, 2, or 3.) The error will thus be $\mathrm{O}(h^{\alpha+1})$, regardless of the usual order of accuracy of the scheme. A typical adaptive procedure will require that the error on any subinterval be less than $h\epsilon/L$, where ϵ is the specified error bound for the whole interval. To satisfy this, h will have to be less than some constant times $\epsilon^{1/\alpha}$. When $\alpha<1$, as it may be when only one singular term is subtracted off, a great deal of extra refinement may be required.

The second way to remove the singularity is to find a change of variables $x=t^{\alpha}$ such that the integrand is well-behaved when expressed in terms of t. In general, we want to choose α so that for $t\approx 0$, $f[x(t)]d[x(t)]$ will behave like $t^{\beta}dt$ for some small nonnegative integer β. A little algebra gives $\alpha=\frac{\beta+1}{\gamma+1}$. The new integrand is then $\alpha t^{\alpha-1}f(t^{\alpha})$, which indeed has the expected leading-order behavior near $x=0$. Dias [15] used this method with $\beta=0$, which works quite well for $-1<\gamma\leq 0$—a similar treatment appears in [54]. However, we have found that polynomial integrators still have trouble with the transformed integrand when $\gamma>0$. To see why this happens, let us expand f again and look at what the change of variables does to the higher-order terms:

$$f(x)=x^{\gamma}(g_0+g_1x+g_2x^2+\mathrm{O}(x^3)), \tag{3.2}$$

$$f(t^{\frac{\beta+1}{\gamma+1}})=g_0t^{\frac{\gamma(\beta+1)}{\gamma+1}}+g_1t^{\frac{(\gamma+1)(\beta+1)}{\gamma+1}}+g_2t^{\frac{(\gamma+2)(\beta+1)}{\gamma+1}}+\mathrm{O}(t^{\frac{(\gamma+3)(\beta+1)}{\gamma+1}}), \tag{3.3}$$

$$t^{\frac{\beta-\gamma}{\gamma+1}}f(t^{\frac{\beta+1}{\gamma+1}})=g_0t^{\beta}+g_1t^{\beta+\frac{\beta+1}{\gamma+1}}+g_2t^{\beta+2\frac{\beta+1}{\gamma+1}}+\mathrm{O}(t^{\beta+3\frac{\beta+1}{\gamma+1}}). \tag{3.4}$$

Note that the g_1 term can have an infinite-slope singularity when $\beta=0$ and $\gamma>0$. The obvious solution to the problem is to use a larger β, but there is a tradeoff involved since the resulting large value of α makes the integrand evaluation points cluster near the singularity, so that the adaptive integrator must work harder at the other end of the interval. We have found empirically that using $\beta=0$ for $\gamma<-0.35$, $\beta=1$ otherwise, tends to give the best results, which are slightly better than those obtained using the method of Kantorovich and Krylov.

A number of other methods for dealing with singularities have been devised over the years, with varying degrees of success. In 1964, Binns [4] presented a simple method where an analytic approximation is used very near singularities and numerical methods are used farther away. It is not very robust, and was actually repudiated in a later paper by Binns, Rees, and Kahan [5], where the authors described non-adaptive implementations of the Gauss-Jacobi and Clenshaw-Curtis quadrature rules. A general survey and comparison of various methods for Schwarz-Christoffel integrals is given by Meyer in [43]. Davis [9] derived a rather clever second-order scheme based on the trapezoid rule in 1979, which has since been used by Sridhar [59,60,61] and by Floryan [21,22]. The advantage of this method is that it uses a single formula throughout the domain, so it is not necessary to break intervals into pieces near different singularities; the disadvantage of course is that it is only second order. Other low-accuracy integration formulas suitable for Schwarz-Christoffel mapping are described in a recent paper by Floryan and Zemach [24].

3.4 Adaptive Quadrature

We have tried three different adaptive quadrature package routines with the strip transformation, and have used each one with different methods of singularity removal in an attempt to determine which combinations were best for the job. The first and simplest of these routines is QUANC8 [25]. QUANC8 uses the 8-panel Newton-Cotes formula once and then twice on the requested interval, and compares the two results to obtain an error estimate. If the error is larger than the tolerance supplied by the user the same procedure is used on each half of the interval, and thereafter only those intervals which continue to give trouble are further subdivided. The Newton-Cotes formulas are not as accurate as, for instance, the Gauss formulas of corresponding size, but they have the advantage that the equally spaced abscissas reduce the work required when the mesh is refined.

Due to the simplicity of this routine we were easily able to modify it to use complex arithmetic. The only other modification we are using currently is a feature that permits the routine to halt at the first level of refinement (17 integrand evaluations); the original routine proceeded through to the second level of refinement (33 evaluations) before deciding whether or not to stop. The authors of the routine probably assumed that most calls to an adaptive quadrature subroutine would involve difficult integrals requiring a relatively fine mesh. For Schwarz-Christoffel integrals with adequate singularity removal, though, the integrand is often well behaved, making early termination quite desirable.

The two other routines we used were from QUADPACK [52]. QAGS applies a 21-point Gauss-Kronrod rule to each subinterval and uses the epsilon extrapolation algorithm to classify and account for endpoint singularities. The fact that the routine does not attempt to evaluate the function at the endpoints of the interval is an advantage since it permits operation even when the integrand is unbounded near an endpoint. QAWS uses the Clenshaw-Curtis algorithm on a subinterval containing a singularity and Gauss-Kronrod elsewhere. It takes for input the function g, the integrand with the singularity divided out, and requires evaluation of this function at the singularity. Both of these routines work in real arithmetic and are too complicated to modify easily, so they have to be called separately for the real and the imaginary parts of each complex integral. When the contour of integration is along a side of the strip the complex integral can be reduced to a real integral, so only one call is necessary.[1]

Table 3.1 summarizes the results of several experiments using the strip transformation 2.10 for a relatively simple 10-cornered polygon. The region being mapped is shown in Figure 3-1—note that though contours are shown for the map from a rectangle, the integration takes place in the usual infinite strip. For comparison purposes, the compound Gauss-Jacobi method required 112 integrand evaluations for an accuracy of 10^{-6},

[1]We have not actually implemented this feature in our test program, as the relatively poor performance of the QUADPACK routines did not seem to merit the additional effort. In the tables that follow, we give both the actual numbers of calls required by QUADPACK and our estimates of the numbers that would be required if real integrals were used whenever possible.

		QUANC8	QAGS	QAWS
Accuracy $= 10^{-6}$	Singularities Unchanged	Fails	8568 (4851)	800** (450**)
	One Term Subtracted	272**	5712 (3318)	1280** (720**)
	Two Terms Subtracted	272**	672** (378**)	2160 (1160)
	Change of Variables	272**	672** (378**)	—
Accuracy $= 10^{-10}$	Singularities Unchanged	Fails	11844 (6615)	800** (450**)
	One Term Subtracted	4272	7056 (3990)	1280** (720**)
	Two Terms Subtracted	672	4914 (2919)	7840 (4200)
	Change of Variables	576*	1008* (546*)	—

A: Contours connect adjacent corners along boundary.

		QUANC8	QAGS	QAWS
Accuracy $= 10^{-6}$	Singularities Unchanged	Fails	5376	580*
	One Term Subtracted	218	3108	800**
	Two Terms Subtracted	202*	588*	1640
	Change of Variables	202*	588*	—
Accuracy $= 10^{-10}$	Singularities Unchanged	Fails	6846	940*
	One Term Subtracted	3050	4578	1000*
	Two Terms Subtracted	634	3402	4600
	Change of Variables	538*	1260*	—

B: Contours connect alternate corners through interior.

Table 3.1: Performance of three adaptive quadrature procedures, with various singularity removal options, on a 10-corner strip transformation problem (Figure 3-1). The total numbers of integrand evaluations required are given. In the first section the contours pass around the boundary from one vertex to the next; in the second a smaller number of contours pass through the interior, visiting only alternate vertices. When any type of singularity removal is attempted, integrals from one corner to another are split into two pieces; otherwise they are calculated in one piece. Figures in parentheses give the performance expected from QUADPACK routines if real integrals were used whenever possible. A conservative implementation of the compound Gauss-Jacobi algorithm beats all of these routines significantly, requiring 112 evaluations for an accuracy of 10^{-6}, 176 for 10^{-10}.

*Some adaptive refinement was necessary, but apparently not due to trouble near endpoint singularities.

**No refinement was necessary anywhere in the domain. This is the smallest number of evaluations possible with this routine.

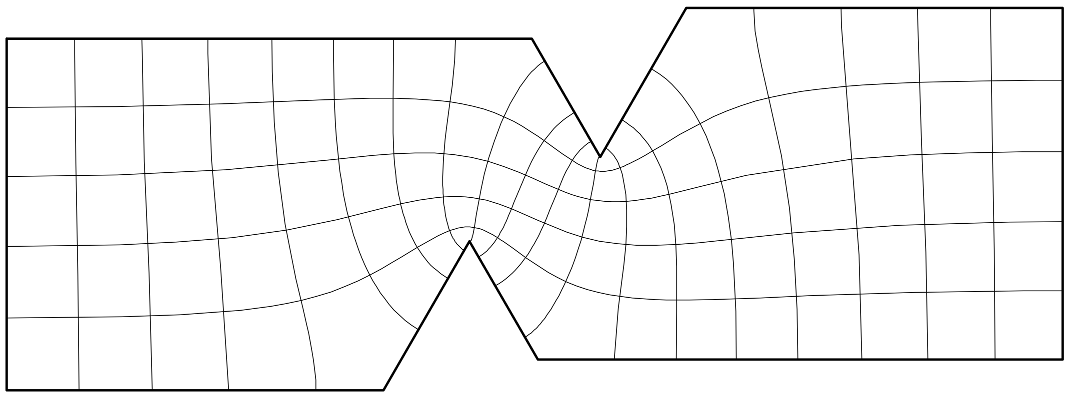

Figure 3-1: A simple polygon mapped from a rectangle using the strip transformation.

176 evaluations for 10^{-10}. (We used 7 and 11 nodes per subinterval, respectively, along contours of integration which followed the boundary of the region. This relatively conservative approach provided significantly better accuracy than required for each of the integrals computed.)

There is quite a bit of information in the table which is worthy of comment. The extrapolation feature of QAGS seems to work after a fashion, but it is quite expensive and can't compete with methods which use the known behavior of the singularity. In addition, QAGS suffers from a large (21-point) "minimum cycle", though with singularity removal and real integrals the routine would almost be competitive when high accuracy is required. The routine QAWS is also handicapped by a large minimum cycle—40 points when only one end is singular, 50 points when both are singular. At high accuracy it would certainly be competitive if real integrals were used, but at low accuracy it is too wasteful. (It should be remembered that even a "low" accuracy of 10^{-6} would be more than sufficient for most applications.)

The QAWS entries with singularity removal do not represent a serious use of this routine, but we wanted to compare its performance on weak singularities to the performance of the polynomial algorithms. Strangely, it performs well when one singular term is removed but poorly when two terms are subtracted off. Our interpretation of this

		QUANC2	QUANC4	QUANC8
Accuracy $= 10^{-6}$	Two Terms Subtracted	208^{*} 282^{*}	144^{**} 146^{*}	272^{**} 202^{*}
	Change of Variables	424^{*} 338^{*}	248^{*} 234^{*}	272^{**} 202^{*}
Accuracy $= 10^{-10}$	Two Terms Subtracted	2188 2918	1048 1090	672 634
	Change of Variables	4164^{*} 3562^{*}	1208^{*} 1082^{*}	576^{*} 538^{*}

Table 3.2: Performance of three variations on the routine QUANC8, with two different singularity removal options, on a 10-corner strip transformation problem (Figure 3-1). In each box the first entry is the number of integrand evaluations required when contours follow the boundary, while the second is the number required when contours pass through the interior to alternate vertices.
*Some adaptive refinement was necessary, but apparently not due to trouble near end-point singularities.
**No refinement was necessary anywhere in the domain. This is the smallest number of evaluations possible with this routine.

behavior is as follows: The second term involves a derivative, which we have computed via a finite difference. The precision of this difference is on the order of 10^{-8}, so a small part of the infinite-slope singularity is left in the function passed to the integrator. In the case of QAWS, this function is then divided by a factor which approaches zero at the singularity, amplifying the spurious component enough to introduce significant errors.

The clear winner out of the three codes is QUANC8, the simplest routine in the group. What the QUADPACK routines lack in efficiency, however, they make up for in accuracy and robustness. QAWS might be the preferred choice in an application where high accuracy is required. Among the singularity removal options the change of variables seems marginally preferable, but the two-term subtraction remains a viable contender. (In other experiments with more distorted regions, the change of variables technique also shows more resistance to some sources of rounding error.) As for the choice of contours, the interior paths used in the second part of the table appear to have a slight advantage. This is more than made up for, though, by the fact that integrals along the boundary can use real arithmetic. As we indicated in Section 3.1, we prefer the contours along the boundary both for this reason and for the fact that the algorithm seems to be somewhat more robust when these paths are used.

We were so impressed by the importance of a small minimum cycle that we experimented with modifications to QUANC8 using the 4-panel and 2-panel (Simpson) Newton-Cotes formulas. The results are shown in Table 3.2. It seems from these figures that QUANC8 is indeed the best compromise between accuracy and sophistication, though QUANC4 is competitive at low accuracy. In complicated problems with secondary crowding effects QUANC8's advantages become even more apparent. On the other hand, in applications where typical intervals of integration are very short, e.g. graphics, we have gotten the best performance from QUANC2.

More sophisticated adaptive integrators built on the QUANC8 structure would probably be better still. We can envision a special-purpose code for Schwarz-Christoffel problems which in the interior would use the 4- and/or 8-panel Newton-Cotes formulas, and in the subintervals at the endpoints would use formulas explicitly recognizing the singularities. The quadrature rules presented by Floryan and Zemach in [24] have equally-spaced abscissas, and thus might be particularly well-suited for such an algorithm.

3.5 Conclusions

Among several different adaptive procedures tested, the routine QUANC8 tends to give the best performance, though variations using smaller stencils may be competitive for some applications. Singularity removal via a change of variables, and paths of integration along the boundary of the region, seem to be the most robust and efficient methods to use with QUANC8 and related routines. The compound Gauss-Jacobi algorithm, however, though it does not provide strict error bounds, tends to be faster for most problems than any of the adaptive methods by a factor of two or more. We therefore recommend this algorithm for all computations where precise knowledge of the quadrature error is not required.

Chapter 4

Solving the Parameter Problem[1]

Exactly three real parameters must be specified in order to make any conformal map between simply connected regions unique. The case where the preimage domain is the unit disk is dealt with by the Riemann mapping theorem (cf. [35, pp. 324–328]), which states that there is a unique mapping function sending the center of the disk to any given point of the image and the positive real axis to a curve leaving that point in a specified direction. If the image domain is bounded by a Jordan curve, then by the Osgood-Carathéodory theorem [ibid., pp. 345–356] the mapping function can be extended continuously onto the boundary. In this case it is easy to show, using mappings of the unit disk to itself, that an alternative characterization is given by specifying the images of three distinct points on the unit circle. The only constraint on the locations of these points is that they must appear in the same order on both boundaries. The two characterizations can also be applied to maps from other standard domains, such as the upper half-plane or a strip, by composing a map from the desired domain to the unit disk with a map from the disk to the target region.

The Schwarz-Christoffel transformation, and related formulas, use the locations of all the prevertices in computing the mapping function. Since there are typically more than three corners in the target domain, not all of these points can be specified arbitrarily.

[1]Portions of this chapter appeared in a substantially different form in [38].

In general, however, there is no known procedure for determining the prevertices analytically. The method used in practice is to make an initial guess of their locations and use the mapping function to compute the resulting image polygon. All of the angles in this polygon will be correct, but some of the side lengths will be incorrect due to the incorrect placement of the prevertices (see Figures 1-1 and 1-3). An iterative procedure can then be used to adjust the prevertex locations so as to make the computed polygon match the desired polygon to any required accuracy.

Two general methods for adjusting the prevertices to their proper values have appeared in the recent literature. Davis [9] uses an iterative procedure based directly on heuristic assumptions about the geometry of the problem. In the first section below we show that these assumptions are not always justified, and that counterexamples can be found for which the method does not converge. Trefethen [62] uses a more general algorithm involving a formulation of the problem as a system of nonlinear equations, which can then be solved by existing software packages. This method has the advantage of being more adaptable, in that it can be modified to work on variations of the Schwarz-Christoffel problem for which geometric insight fails completely. In addition, it is often successful even on polygons for which Davis's algorithm diverges. In Section 4.2 we show, however, that Trefethen's system also diverges in some cases. In the final section we describe some obstacles which must be overcome to construct a method that will always converge, and discuss three approaches which may in time yield such an algorithm.

(Each of the methods mentioned above has since been used by several other authors, as cited in the sections below. We should also mention some precursors. Woods [68] works out iterative solutions for several specific examples on a case-by-case basis. If generalized, his approach would somewhat resemble Davis's algorithm. It is likely that other similar techniques were developed for hand calculation before the widespread use of computers. Reppe [56] and Prochazka [55] both use systems of nonlinear equations similar to Trefethen's, which they solve using their own implementations of Newton's method rather than by applying existing software packages.)

4.1 Davis's Iteration: Local Divergence

In [9], Davis presents the following iteration for solving the parameter problem:[2]

$$(z_{j+1} - z_j)_{\text{new}} := k(z_{j+1} - z_j) \cdot \frac{|w^*_{j+1} - w^*_j|}{|w_{j+1} - w_j|}, \tag{4.1}$$

where the map is taken from the upper half-plane. The implicit assumption is that the length of each side is roughly proportional to the separation between the corresponding prevertices on the real axis. By iterating this procedure it is hoped that the correct solution can be obtained to the desired accuracy. More specific details, e.g., the starting distribution, which points vary and which are held fixed, the precise choice of the factor k, etc., are to be settled on a case-by-case basis. In one problem, for example, Davis selects a new value of k at each step of the process in order to keep two particular prevertices a fixed distance apart. In other examples k is omitted entirely.

The iteration (4.1) often converges within at most a few tens of iterations, particularly on relatively undistorted regions like those in Figures 1-1 and 2-4. The method's rapid adoption by many researchers, e.g., Anderson et al. [3], Floryan [21,22] and Sridhar [59,60,61], is evidence both for its utility and for its ease of implementation.

We have found that it is not the best choice for a general algorithm, however, for the following reasons:

- Determining the details of the algorithm on a case-by-case basis is often impractical. Different geometries suggest different formulations, and each example appearing in the literature uses (4.1) in a slightly different way. Sridhar and Davis [61] themselves acknowledge at one point that increasing distortion of the target polygon with respect to the standard region leads to slower convergence or in some cases even divergence. It is not clear which variation, if any, would be best suited to a wide range of problems.

[2] Throughout this chapter we will use z^* and w^* to denote the (unknown) correct prevertices and the desired corners, respectively, for the polygon. Without the superscript these variables represent an approximation based on incorrect prevertex positions.

- When used with the strip transformation, the method gives no information about the position of z_n, the leftmost prevertex on the top side of the strip. Sridhar avoids this problem by restricting attention to channels where symmetry implies $z_n = i$. Floryan uses a double iteration for the asymmetric case—a 1-dimensional secant iteration determines a value for z_n at each step of the global iteration (4.1). In our own experiments with Davis's method we have used a one-dimensional finite difference to establish the dependence of the angle between vectors $w_n - w_1$ and $w_2 - w_1$ on the position of z_n. A single Newton step using this derivative is applied to z_n at the same time that the other prevertices are adjusted according to (4.1). This approach yields results similar to Floryan's and requires only two function evaluations per global step.

- Many Schwarz-Christoffel problems that arise in applications come with additional conditions to be satisfied. For example, the conformal modulus μ might be specified in advance and one of the side lengths left unspecified. In such situations one has a "generalized parameter problem" to solve [64,66]—which may not be an easy matter if one is using an iteration like (4.1) that is dependent on geometric insight.

- The convergence rate near the solution is only linear, which may be a disadvantage for high-accuracy computations.

- Most critically, the iteration may not converge at all, even when started near the solution. The reason is that the proportionality assumption can be violated by difficult problems in at least two different ways. First, it assumes that only the prevertices at the endpoints of an interval have a major effect on the length of the corresponding side, and this condition is violated when crowding occurs. Second, if the two singularities at the endpoints are strong, i.e., the interior angles are acute, then the length of the side may actually decrease as the prevertex separation increases. Even with the standard Schwarz-Christoffel formula there are geometries for which (4.1) fails to converge, and with the strip transformation we have the

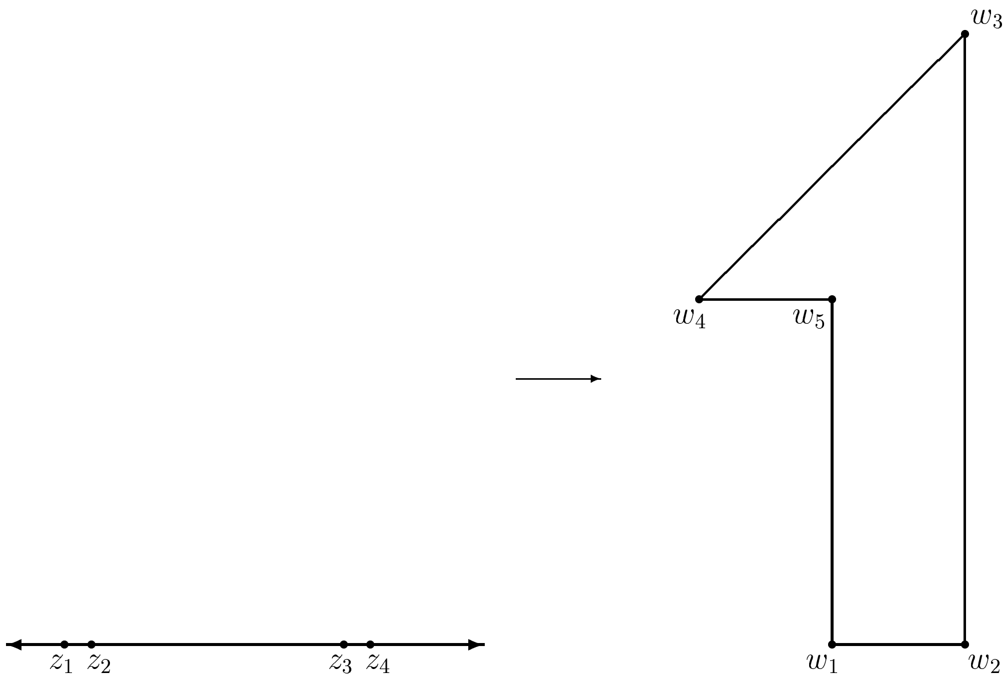

Figure 4-1: This conformal map from the upper half-plane to a polygon causes one version of Davis's iteration to diverge.

added problem that singularities on the opposite side of the strip may also strongly influence an interval.

We now discuss this last point, the divergence of the Davis iteration, in more detail. Here is our basic claim:

Theorem 4.1 *Davis's iteration (4.1) is not always locally convergent.*

Proof: To prove this theorem it is enough to exhibit examples in which the iteration diverges; for such examples, consider the regions shown in Figures 4-1 and 4-2. ∎

For the map illustrated by Figure 4-1, we hold z_1 and z_2 fixed, use a linear transformation to send w_1 and w_2 to their proper positions, and let the point at infinity map to w_5. The positions of z_3 and z_4 are then determined from (4.1) with $k \equiv 1$. Because of the slight crowding effect, these prevertices are relatively isolated from z_1 and z_2; changing their separation by moving z_4 will have little effect in the neighborhood of the left pair because the slope of $(z - z_4)^{-3/4}$ becomes flatter at increasing distance. We can therefore examine the effects on corners w_3 and w_4 individually, without significantly altering the normalization of the problem.

Consider an idealization of the Schwarz-Christoffel integral involving only two prevertices:

$$I = \int_0^a x^\alpha (a-x)^\beta dx. \tag{4.2}$$

All other factors being roughly equal, I determines the length of the side between the images of 0 and a. If we change the variable of integration to $t = x/a$, the dependence becomes

$$I = a^{\alpha+\beta+1} \int_0^1 t^\alpha (1-t)^\beta dt. \tag{4.3}$$

We see that whenever $\alpha+\beta$ is less than -1, changes in the prevertex separation a actually have an inverse effect on the length of the corresponding side. The z_3 and z_4 terms in our example mapping each have powers of $-3/4$, so the error in the distance between w_3 and w_4 will become progressively larger as the iteration (4.1) proceeds.

It would not be difficult to patch up this particular example to make it convergent, but as the complexity of the region increases it becomes more and more difficult to foresee and account for all of the trouble spots. (For an extreme example, see Figure 4-4 below.) The simple dependence given by (4.3) is just an illustration—in larger problems the joint interactions as many prevertices move around simultaneously are much less clear. In most cases there is probably no simple explanation for the observed behavior. These considerations are the basis for

Conjecture 4.2 *No iteration analogous to (4.1) can be locally convergent for all polygons.*

In Figure 4-2 we show a locally divergent example using the strip transformation, which is essentially a modification of the closed polygon of Figure 2-9 with the right angles at corners w_1 and w_n replaced by straight angles and corners w_8 and w_9 eliminated entirely. This example diverges even when started very near the solution. For all calculations on this problem we have used the high-accuracy adaptive quadrature

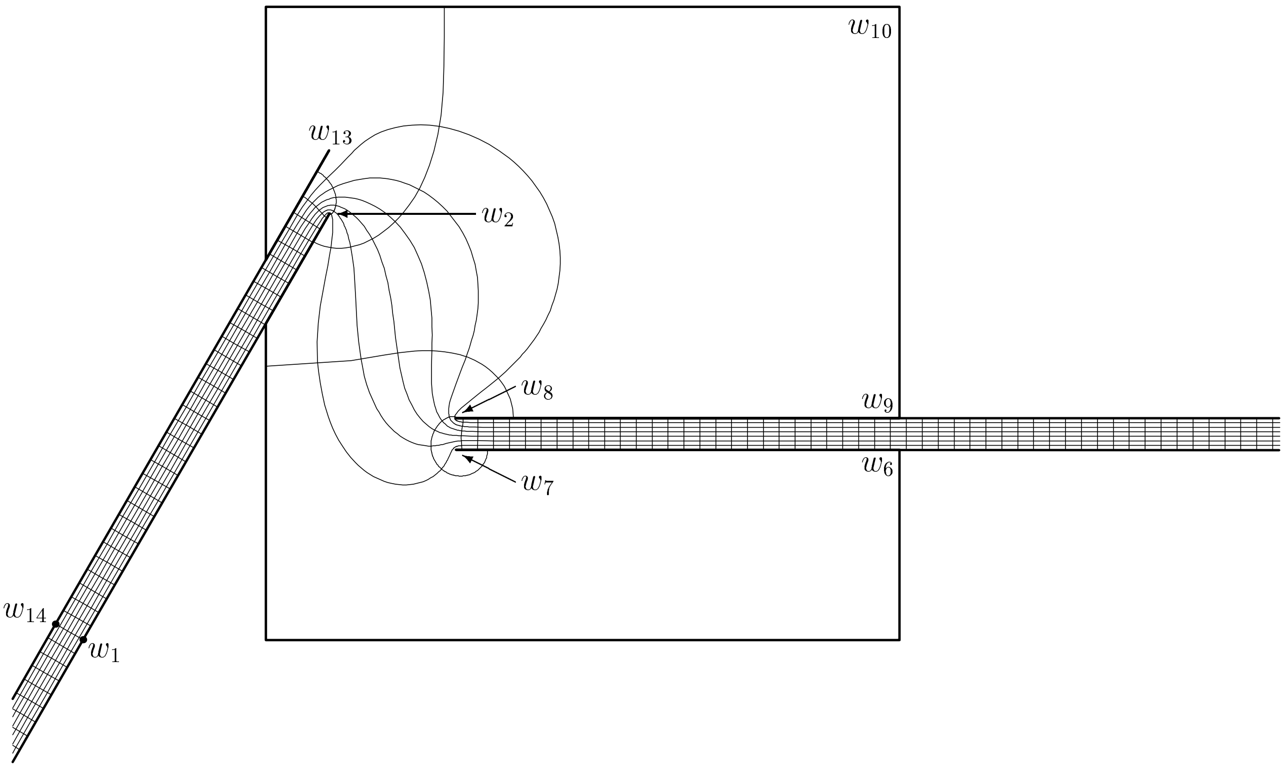

Figure 4-2: Davis's iteration is locally divergent when used with a strip transformation to this distorted channel.

methods described in Chapter 3—though somewhat inefficient, these routines give quite reliable error bounds, so the failure cannot simply be caused by an inaccurate quadrature algorithm.

To confirm that our treatment of z_n did not cause the divergence, we experimented with variations where adjustments to this corner were either underrelaxed or overrelaxed at each iteration. The divergence followed substantially the same course in all of these cases. Interestingly enough, if the straight angles w_1 and w_{14} are removed from the problem, so that prevertices z_2 and z_{13} become the farthest to the right, this new formulation for the same geometry is just barely convergent. (The error norm decreases by a few percent per iteration.) We were quite surprised by this change since in the divergent case most of the distortion appears at the other end of the figure. Apparently, though z_1 and z_{14} are too far "upstream" to have a significant effect on the geometry, the motions of z_2 and z_{13} under (4.1) are instrumental in driving an instability further downstream.

While we do not have a precise explanation for this behavior, it appears to be a result of interactions among many widely separated prevertices.

4.2 Nonlinear Equations: Local Convergence, Global Divergence

In SCPACK, Trefethen [62,63] formulates the parameter problem as a general system of nonlinear equations, $F(x) = 0$, which can be passed to any of several existing software packages for solution. The actual choice of package does not seem to be critical—we have experimented with three different routines on strip transformation problems (see Chapter 2), Dias [15] and Trefethen have tried two alternatives for the Schwarz-Christoffel transformation, and Bjørstad and Grosse [6] have used yet another package for circular arc polygons, all with similar results.

Most of the available packages for this type of continuous optimization problem use some variation on the Newton iteration

$$x_{\text{new}} = x - J^{-1}F(x), \tag{4.4}$$

where J is an approximation to the Jacobian of F at x, with various error-checking and backtracking algorithms added to prevent wild leaps into oblivion. Davis's iteration (4.1) can also be interpreted as an approximate Newton scheme using a particular approximate Jacobian and no backtracking. The examples in the previous section establish that the approximate Jacobian does not always yield a descent direction. (In experiments where a backtracking procedure was used in conjunction with Davis's method, the success rate actually declined. This is because even a single failure to choose a descent direction terminated the program. A package routine faced with this situation simply computes a new Jacobian.)

SCPACK maps the unit disk onto arbitrary polygons with or without corners at infinity. The prevertex positions are determined by an unconstrained set of $n-1$ independent variables,

$$x_j = \log \frac{\theta_j - \theta_{j-1}}{\theta_{j+1} - \theta_j}, \tag{4.5}$$

where θ_j is the angular position of prevertex z_j on the unit circle. A single prevertex, z_n, is held fixed at 1, yielding the two end conditions $\theta_0 = 0$ and $\theta_n = 2\pi$. This system includes two more degrees of freedom than the $n-3$ required to determine all the side lengths; Trefethen uses these to position the image w_c of the center of the disk at a chosen point inside the polygon. Specifically, once equation (1.2) has been integrated to determine the shape of an approximate polygon, a linear transformation is used to give w_c and the corner w_n their proper positions. The $n-1$ components of F are then defined as follows:

$$F_j = \begin{cases} \operatorname{Re}(w_1 - w_1^*) & j = 1, \\ \operatorname{Im}(w_1 - w_1^*) & j = 2, \\ |w_{j-1} - w_{j-2}| - |w_{j-1}^* - w_{j-2}^*| & 3 \le j \le n-1. \end{cases} \tag{4.6}$$

The first two elements of this vector are the complex error in the position of w_1; each of the other $n-3$ is the error in the length of a side. A somewhat more elaborate system is used when the polygon has corners at infinity, but that is a complication that need not concern us here.

Even if the position of w_c is important to a given application, it is not necessary to include it as part of the parameter problem. Given any set of prevertices that generates the correct polygon, we can invert the mapping function by the method described in [62] and map the desired location of w_c back onto the unit disk. We then find a Möbius transformation (cf. [2]) mapping the disk to itself and sending $f^{-1}(w_c)$ to the origin. The new distribution of prevertices thus obtained will still generate the same polygon, and will correctly map the origin onto w_c.

This procedure permits us to reduce the size of the parameter problem back to $n-3$ unknowns. One possible way to organize the reduced system uses the upper half-plane as the standard domain. Three prevertex positions are specified: w_1 is placed at -1, w_{n-1}

at 1, and w_n is implicitly placed at infinity. An unconstrained set of $n-3$ independent variables similar to (4.5) controls the distribution of the remaining prevertices between -1 and 1. Upon computing an approximate polygon, a linear transformation can be used to send w_n and w_1 to the correct locations, and $n-3$ side lengths then determine the residual vector F. Many other variations on this theme can be developed; for example, a similar formulation for the rectangle map is discussed in Section 2.4. Since all of the nonlinear systems considered here and in Chapter 2 are based primarily on side lengths, they have many features in common. The theoretical results discussed in the remainder of this section apply with equal force regardless of which of these systems is used.

The method of solving the parameter problem via a nonlinear equations algorithm, as implemented in SCPACK and its variations, is generally quite robust. It converges for a wide range of problems, and most of the cases where it has been observed to fail are so strongly crowded that no competing method for computing the same map could be expected to give better performance. The reason for this success is quite simple: The solution norms for these systems have no suboptimal local minima (Theorem 4.3). One feature common to all robust classical algorithms for solving systems of equations (as opposed to specialized methods like simulated annealing) is that they always march downhill—that is, they decrease some norm of the solution vector at each step (cf. [30]). There is obviously a global minimum with norm zero at any solution to the system. Any other local minimum, however, can trap an algorithm so that it never reaches the solution. The only recourse when this happens is to restart the iteration at a different starting point. Since the systems used for the parameter problem have no local minima, a robust iteration can not converge to any point but the solution itself.

Theorem 4.3 *For any of the nonlinear systems we have considered, the 2-norm of the residual vector has no stationary points other than at the desired solution.*

Proof: We will first show that there are no local minima by using a simple argument based on continuity. Ruling out other types of stationary points is more difficult, and requires the differentiability conditions we will establish in Lemma 4.4.

To see that there are no local minima with any of the systems discussed here or in Chapter 2, consider making a small change in the length of just one side. Given any legal polygon which could be mapped from some distribution of prevertices, it is always possible to change the length of any side, and thus any of the components of the residual vector F, without affecting any of the other residuals. In order to keep the polygon closed one or both of the two "extra" sides must also be changed, but this has no effect on F. The residual vectors defined in (4.6) and in (2.12) have one or two components which are not side lengths, but these too can be altered individually. Since any approximation to the desired polygon can be gradually improved, no local minimum is possible.

This argument depends on the fact that F is a homeomorphism, a continuous 1-to-1 function with a continuous inverse. We establish this as part of a stronger result which will be useful later:

Lemma 4.4 *Each of the nonlinear systems we have considered defines a diffeomorphism between the set of parameter vectors x and the set of residual vectors $F(x)$.*

Proof: The conditions for a diffeomorphism are that F be a 1-to-1 function and that both it and its inverse be continuously differentiable (see [20]). To see that F is 1 to 1 we note that in the systems we have discussed, each vector $F(x)$ uniquely determines the shape of a polygon. This polygon in turn determines a unique set of prevertices up to the three degrees of freedom permitted by the Riemann mapping theorem. The definition of x takes these degrees of freedom into account, so there is a unique x for each residual vector in the range of F. That F itself is continuously differentiable is readily apparent from the form of the Schwarz-Christoffel mapping function.[3] The same result can be obtained for F^{-1} by considering any of a number of other algorithms for directly computing the inverse map, e.g., the formulation as a Dirichlet problem given in [35, pp. 370–374]. ∎

[3]The system (2.12) which we used for the channel map includes one angle which could vary discontinuously as it crosses a branch cut. For this geometry, however, it is always possible to position the branch cut away from the range of possible angles so that no ambiguity can arise.

The impossibility of local minima follows immediately from the fact that F is a homeomorphism, since a local minimum of a norm of F with respect to x would also be a local minimum of a norm of F with respect to F itself, which is possible only at the solution $F(x) = 0$. In fact, $\|F\|_2 = (F^T F)^{1/2}$ can have no suboptimal critical points of any kind—extrema, saddle points, cubic points, etc.—because at such points the Jacobian $F'(x)$ would be singular, which is forbidden by Lemma 4.4. (If we seem to be belaboring an obvious point, it is because both approaches, that of improving one residual without affecting the others, and that of considering F as a diffeomorphism, can be useful in evaluating nonlinear systems for many variations of the Schwarz-Christoffel transformation.) [End of proof of Theorem 4.3.] ∎

A lack of suboptimal critical points is an important asset to consider in formulating a nonlinear system, but to guarantee convergence we need in addition a condition of compactness. In particular, an algorithm may fail if there is a descent direction towards infinity.[4] Though a rigorous discussion of iterative algorithms for nonlinear systems is beyond the scope of this thesis, the following theorem establishes the conditions a system must satisfy in order for us to prove convergence from a given starting point:

Theorem 4.5 *Assume that $g : D \subset \mathbf{R}^n \to \mathbf{R}$ is continuously differentiable on the open set D and that there is an $x^0 \in D$ such that the level set $L(x^0) = \{x \in D | g(x) \leq g(x^0)\}$ is compact. Assume that g has a unique critical point x^* in $L(x^0)$. Then under suitable conditions concerning the selection of descent directions and step lengths—which are satisfied by any robust iterative scheme—the sequence of iterates x^k converges to x^*.*

Proof: This is an adaptation of Theorem 14.3.2 of [47], to which the reader is referred for more information. ∎

Applying this to our nonlinear systems, we immediately obtain

Theorem 4.6 *Assuming exact arithmetic, a robust iterative algorithm converges locally when applied to any of the nonlinear systems we have considered.*

[4]Excessive crowding can also render some solutions inaccessible which in theory are perfectly well-defined.

Proof: We have established that $g = \|F\|_2^2$ can not have a suboptimal critical point for any finite value of x. The remaining condition of Theorem 4.5 is that the level set $L(x^0)$ be compact. This follows from Lemma 4.4 for points x^0 sufficiently close to x^*. ∎

When we consider arbitrary starting points, however, $L(x^0)$ may not always be compact. This occurs when some point outside the range of F, i.e., a residual vector which does not correspond to a legal polygon, has a norm less than $\|F(x^0)\|$. A descent algorithm could then asymptotically approach the boundary of the range of F, with the norm of x increasing without limit. As we have defined x, this could either mean that two prevertices merge into a single point, or one prevertex recedes towards infinity.

In terms of the image polygon, a descent direction towards infinity would involve one of three possibilities: A side could grow without limit, it could shrink to nothing, or the interior of the polygon could split into two pieces, i.e., the region could become infinitely crowded. The first of these cases is not a problem because in every system we have considered $\|F\|$ would have to increase. The second case is not so trivial; specific nonlinear systems must be treated separately. If side lengths are considered directly, as in Trefethen's system (4.6), then any side could become arbitrarily small. If logarithms of side lengths are used, however, as in (2.12), then only the two "extra" sides could become small without $\|F\|$ becoming large. It is the third case, however, which causes the worst behavior. Not only may there be a descent direction towards infinity, but from some starting points there may in fact be *no* descent paths leading to the solution.

These issues are not academic. In fact, we have found that SCPACK does not always converge:

Theorem 4.7 *The SCPACK algorithm is not always globally convergent. More generally, the nonlinear systems we have considered can cause divergence of any iterative solution algorithm which requires a decrease in the solution norm at each step.*

Proof: The situation is best explained by means of an example. Figure 4-3 shows two polygons, one the desired solution and one an approximation to the solution which might easily arise while attempting to solve the parameter problem. The two bumps can not

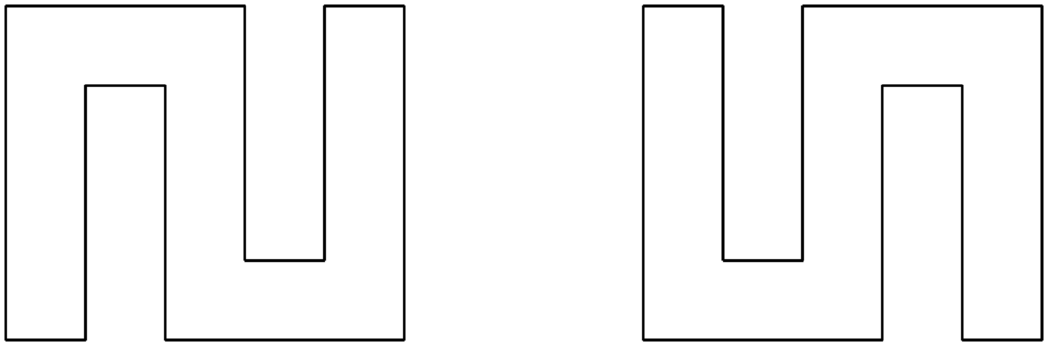

Figure 4-3: No descent algorithm based entirely on side lengths can transform one of these polygons into the other.

pass through one another—as they approach, the region becomes arbitrarily crowded and $\|x\|$ diverges to infinity. The only way one region can be transformed into the other is by temporarily shrinking one or both bumps or stretching one of the other vertical sides. Since this would cause an *increase* in the norm of F, no descent path to the solution exists. A similar example can be constructed for the strip transformation system (2.12). Note that two bumps were required for this demonstration, which considered only side lengths as contributing to the norm of F. The system (4.6) used by SCPACK also involves the image w_c of the center of the unit disk—by using this point it is possible to construct divergent example with only one internal bump. ∎

A nonlinear system like (4.6) can not guarantee convergence from arbitrary starting points, and no efficient method is known for selecting starting points which are close to the desired solution. On the other hand, this type of situation has rarely, if ever, arisen in a practical problem. None of the dozens of SCPACK users has reported any such difficulty, and it was not until the present work that the phenomenon was even recognized. Regions with internal bumps that can catch on each other tend to be at least a bit crowded, and those shown in Figure 4-3 are probably beyond the range of SCPACK. (We have confirmed the divergent behavior using a less crowded problem.) It would still be desirable to have a completely dependable algorithm, however, and regions of this nature may be more frequently encountered when mapping to elongated polygons with methods such as the strip transformation.

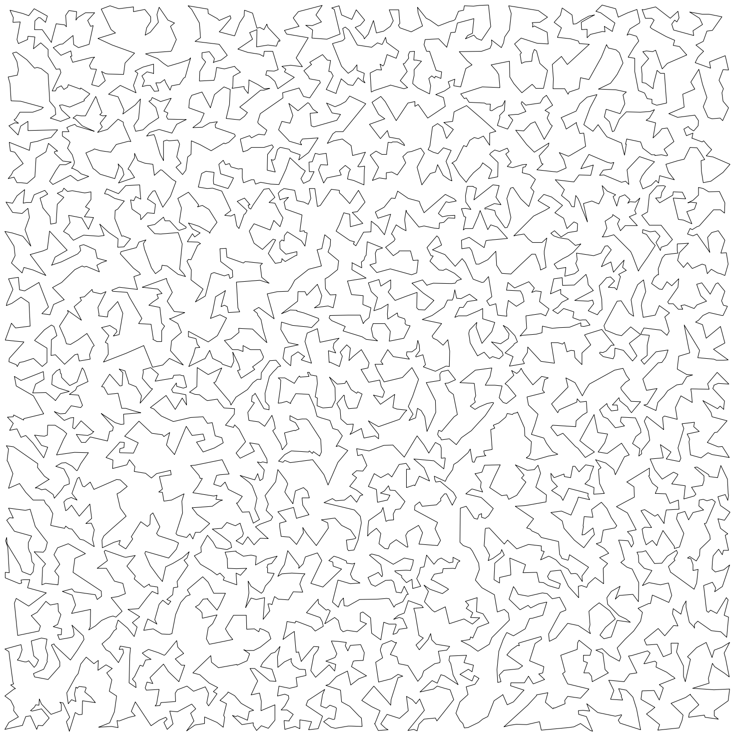

Figure 4-4: A completely general algorithm should in principle be able to map even this polygon. (As a practical matter, though, this particular region is far too crowded to be mapped from any reasonable standard domain.)

4.3 Conditions for Global Convergence

How can we design an algorithm guaranteed to converge for any conceivable polygon? Given a particular geometry like that of Figure 4-3, it is not difficult to design a specialized system of equations with which convergence is achieved from any starting point. The real difficulty is in trying to account for all of the possibilities that could arise in a huge problem, even a 5000-sided monstrosity like Figure 4-4. In this section we present sketches of five different approaches which could be used in cases where an SCPACK-type algorithm fails. None are completely satisfactory, as they range from the tedious to the impractical to the undefined. It may be possible in the future to design a globally convergent algorithm using one of these methods, but it seems likely that the extra bookkeeping involved would make such an algorithm uncompetitive with those discussed in the previous section, except of course in the rare cases where the latter actually fail.

In order to prove that an algorithm is globally convergent for a given system of equations, we have to show that the level set $L(x^0)$ is compact for any starting point x^0 (Theorem 4.5). Since we do not have a good general method for picking starting points, we must assume that x^0 could be any point whatsoever. The range of F must therefore be either an open ball about the origin or all of $\mathbf{R}^n$. It is easy to transform one of these cases into the other, so without loss of generality we can require the range to be $\mathbf{R}^n$. Finally, this must be accomplished without violating any of the other conditions of Theorem 4.5—specifically, without giving $g = \|F\|_2^2$ any suboptimal critical points.

In Theorem 4.6 we established that g has no critical points by using the fact the F is a diffeomorphism. For this approach to work, F must completely characterize the image polygon, i.e., if two legal polygons yield the same value of F then they must be identical. With the systems we have discussed so far this is a triviality: Given a value for F, simply draw one side at a time until the polygon is complete. The position of each side depends only on its predecessors, with no reference to sides that come later. The examples in Figures 4-5 and 4-6, however, illustrate the difficulty of both generating F "one side at a time" and insuring that $\|F\| \to \infty$ whenever $\|x\| \to \infty$. If, on the other hand, all of the sides are considered at once when computing F, it becomes much more difficult to show that the process can be inverted to obtain a unique polygon.

In order to avoid splitting the interior of a polygon into two pieces, as is threatened in Figure 4-3, an algorithm must be able to detect some situations where the boundary nearly intersects itself. This problem is complicated, however, by the fact that some other self-intersections are perfectly legal, as in Figure 4-5. Intuitively speaking, a polygon becomes infinitely crowded when the "inner" side of the boundary approaches itself, but no problem at all occurs when the "outer" side approaches or overlaps itself. The trouble is, it may not be possible to tell which case is which when drawing sides one at a time. The question of whether an intersection point like the one highlighted between w_4 and w_5 is legal depends in part on the behavior of all the other sides of the region.

In Figure 4-5 we can deduce that the questionable intersection is legal without know-

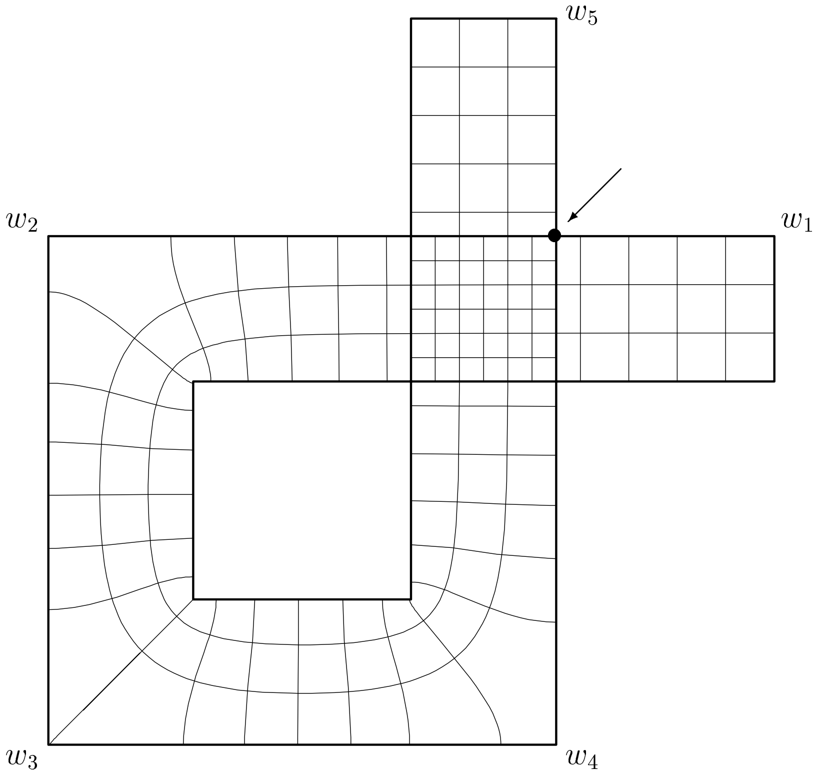

Figure 4-5: This type of self-intersection is permitted. However, if we examine sides sequentially starting at corner w_1, we encounter what *appears* to be an illegal intersection between w_4 and w_5.

ing the lengths of the later sides. The angles alone, which are always known in Schwarz-Christoffel problems, are sufficient to show that the remaining sides must unwrap around the highlighted point and make it a legal intersection. That this type of reasoning is not always possible, however, is demonstrated by Figure 4-6. The lowest drawing here shows the outline of a polygon which could not be mapped from any simply connected region. A "one side at a time" algorithm would have to identify the approach to such a condition from the behavior of the lower boundary alone, starting at w_1. The figure in the middle, though, shows a legal polygon with all of the same angles and identical side lengths up to corner w_{14}. At what point does the polygon actually become illegal? Our hypothetical algorithm would have to permit the intersection on the lower boundary, and then later identify the critical point on the upper boundary where the legality of that intersection is actually determined.

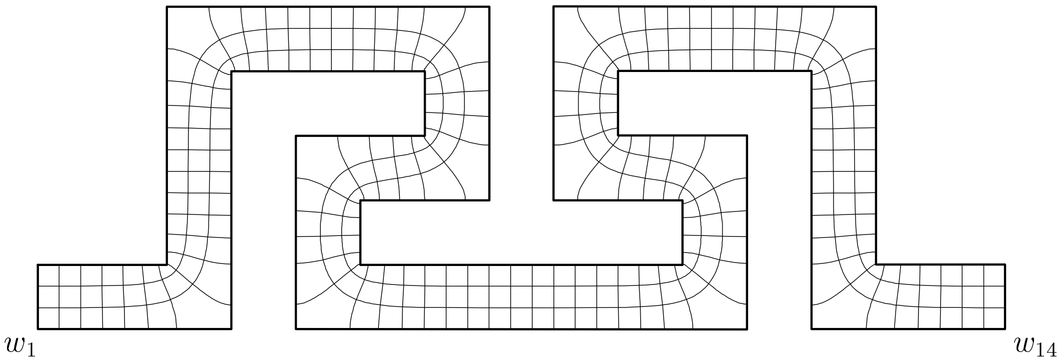

Figure 4-6a: This polygon does not intersect itself.

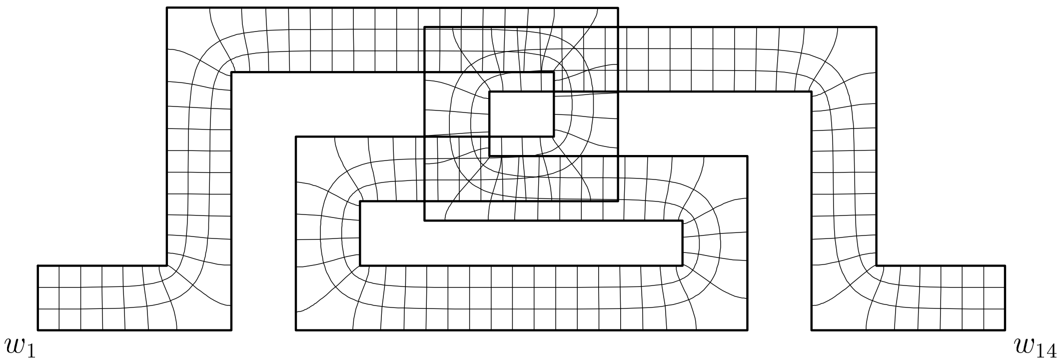

Figure 4-6b: A self-intersecting version of the above polygon, this figure could easily be produced as an intermediate step while solving the parameter problem.

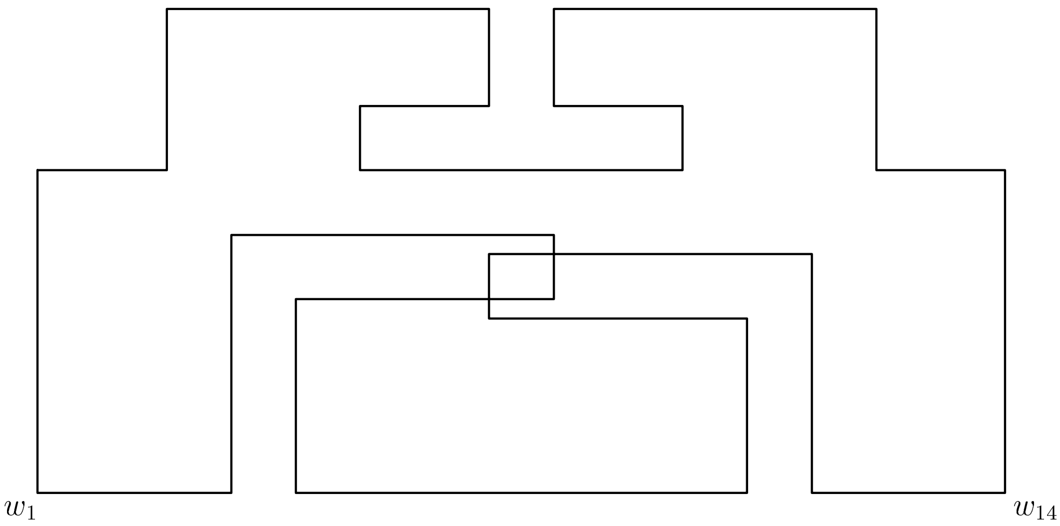

Figure 4-6c: The lower boundary of this illegal polygon is identical to that of Figure 4-6b up to corner w_{14}, including the region where the intersection occurs.

Having established that designing a globally convergent system of equations is not a trivial proposition, we will now suggest some ways in which it might be accomplished.

Algorithm 4.1 (Continuation Method) *Instead of trying to obtain convergence with a single call to a nonlinear equations solver, use a series of polygons of increasing difficulty until the desired solution is obtained.*

For example, with the polygons shown in Figure 4-3, we could first compute a solution where the bumps are short enough to pass each other freely. By using this as an initial condition, we should then be able to obtain the desired solution with no difficulty. With suitable safeguards, this method could be made globally convergent, and it has the advantage of using simple systems of equations which are easy to evaluate. The disadvantages, of course, are that it requires human intervention as well as multiple calls to the solution routine.

Algorithm 4.2 (Inverse Mapping) *Use a completely different conformal mapping algorithm to compute the boundary correspondence function for the map from the polygon to the standard domain. Evaluate this function at the corners to obtain the prevertices for the Schwarz-Christoffel formula.*

This approach does not preclude use of the Schwarz-Christoffel transformation entirely; it merely provides a solution to the parameter problem. The Schwarz-Christoffel formula might still be the most efficient way to calculate specific values of the mapping function for an application.

Neither of these first two methods is intended to be taken seriously as a practical algorithm. Our purpose is merely to show that with sufficient computational effort the prevertices can be obtained for any polygon whatsoever. Given these alternatives, it would definitely be more desirable to have a single system of equations with guaranteed convergence. That such a system is at least possible is shown by the inverse mapping algorithm: If we can calculate the prevertices for the desired polygon, the trivial system

$F(x) = x - x^*$ has all the sufficient properties for global convergence. The idea, of course, is to find another system with the same properties that is easier to calculate.

The next algorithm we consider is an attempt to satisfy the need for some "look-ahead", as demonstrated by the examples, without sacrificing the ability of F to characterize the image polygon. In order to make the function continuous we need to detect near-intersections as well as actual intersections, so that the components of F can approach infinity smoothly when necessary. This is accomplished via an auxiliary function $c_j(F_1, F_2, \ldots, F_{j-1}, s)$, which serves as a very crude "crowding detector" and depends only on those components of F which have already been calculated.

Algorithm 4.3 (Look-Ahead) *Assume that the sides corresponding to the first $j-1$ components of F have already been determined, and suppose the side corresponding to F_j has length s. For any legal polygon P which includes these sides, let $d_j(P)$ be the minimum distance, measured through the interior of the region, from any corner to any of the $n-2$ sides not adjacent to that corner. Define c_j as follows:*

$$c_j(F_1, F_2, \ldots, F_{j-1}, s) = [\max d_j(P)]^{-1}, \tag{4.7}$$

where the maximum is taken over all possible completion polygons P. (This is not the only way c_j could be defined—the key property is that $c_j \to \infty$ as the best completion polygon becomes infinitely crowded.) Then define the component F_j by

$$F_j = \int_{s_0}^{t} c_j(F_1, F_2, \ldots, F_{j-1}, s)ds - I_j, \tag{4.8}$$

where s_0 is the value of s for which c_j is a minimum, t is the actual length of the side in the current approximate polygon, and I_j is the value the integral takes in the desired polygon. Note that s_0 is not the minimum value of s; c_j becomes infinite for some $s \in [0, s_0)$. In some cases c_j achieves its minimum over a range of values; if this happens, take s_0 to be the smallest value in the range.

The look-ahead algorithm as defined here gives an idea of what might be accomplished, but is not itself completely rigorous. The distance function d_j should be defined along the image of a path interior to the standard domain in order to avoid trouble when the region intersects itself. More seriously, we have implicitly assumed that the auxiliary function c_j is convex as a function of s. That we have not been able to either prove or disprove such a simple result is a clear indication that we do not have a good mathematical framework within which to rigorously describe maps to polygons. Even if the assumption is true, there may be no efficient method by which c_j could actually be calculated.

The final two algorithms are merely shells—incomplete ideas which we only present because they may provide some insight for other researchers.

Algorithm 4.4 (Continuous Deformation) *Any general method for continuously deforming an incorrect image polygon into a desired target polygon with the same angles might serve as a basis for a globally convergent system. At the very least, such a method could provide a way to automate the continuation algorithm so that it would not require human intervention.*

Algorithm 4.5 (Springs and Rubber Bands) *The idea is to conceptually connect various corners to each other by stretching springs through the interior of the polygon. Thus in Figure 4-3, several springs would be stretched around each bump. The hope is that an appropriate distribution of springs would force the bumps down as they passed each other, avoiding the collision that makes the region infinitely crowded. While we have not been able to prove anything definite about such methods, this still seems to be one of the more promising avenues of investigation.*

This section ends on an inconclusive note. While it may be possible to design a globally convergent nonlinear system based on some of the ideas given here, such a method might nevertheless not be competitive with simpler systems with smaller bookkeeping requirements. The vast majority of uncrowded problems encountered in practice do not pose any difficulty for existing techniques, and the continuation method provides an inelegant but effective approach to those that do cause simple systems to diverge.

Chapter 5

Circular Arc Polygons

5.1 The Schwarzian Differential Equation

As a practical conformal mapping technique, the main drawback of the Schwarz-Christoffel transformation is its limitation to regions with straight sides. This restriction can be partially overcome by breaking up a curved boundary into many short line segments, as in Figure 2-7; as the number of corners becomes large, the individual turning angles approach zero. By taking the limit of an infinite number of shallow corners the products appearing in (1.2) and (2.10) can be be replaced by integrals, giving general expressions for conformal maps onto regions with both corners and curved sides. Of course, solutions to the resulting integral equations must still be approximated numerically, but with an appropriate discretization this can be both more efficient and more accurate than simply using a large number of straight sides.

Analytic treatments of the curved boundary formulas date back at least to Woods's work in the 1950's—a detailed overview appears in [68]. Though Woods discusses approximation techniques for some specific examples, the first general computational implementation appears to be that of Davis [9], which divides each curved section of the boundary into a relatively small number of curved elements. These curved segments will not in general match the exact shapes of the desired arcs in a target region, but by piecing together enough of them the given arcs can be approximated to any required accuracy. Normally this is done using curves with continuously varying tangents, but higher degrees of smoothness can be attained if necessary. The method has since been used by Davis's student Sridhar [59,60,61] and by others including Floryan [21,22] and Hoekstra [37].

There is an entirely different approach, however, which is not as general but which gives an exact mapping function from the upper half-plane to a **circular arc polygon**: a region bounded by arbitrary arcs of circles, including straight lines (Figure 5-1). In 1869 Schwarz [57] showed that any mapping $w = f(z)$ to such a region satisfies the ordinary differential equation

$$\{f, z\} = \frac{1}{2}\sum_{j=1}^{n} \frac{1-\alpha_j^2}{(z-z_j)^2} + \sum_{j=1}^{n} \frac{\beta_j}{z-z_j}, \tag{5.1}$$

where z_j is the location of the jth prevertex on the real axis, α_j is $1/\pi$ times the interior angle at the vertex $w_j = f(z_j)$, and the β_j are n additional real parameters. The notation $\{f, z\}$ is shorthand for the expression

$$\{f, z\} = \left[\frac{f''(z)}{f'(z)}\right]' - \frac{1}{2}\left[\frac{f''(z)}{f'(z)}\right]^2, \tag{5.2}$$

which is variously known as the Schwarzian differential operator, the Schwarzian derivative, the Schwarzian differential parameter, or simply the Schwarzian. Predictably enough, (5.1) is often referred to as the Schwarzian differential equation.

Though Schwarz was the first to apply it to conformal mapping, Hille [36, p. 374] reports that the operator (5.2) appeared in a paper by Kummer in 1834 and was known to Riemann as early as 1857. It plays an important role in the theory of univalent functions, as recorded by Nehari [45; 46, ch. 5] and by Hille [36, ch. 10].

The differential equation (5.1) was solved by Schwarz for $n = 3$ in terms of hypergeometric functions. (The case $n = 2$ can be solved by much simpler means.) No analytic solutions for higher values of n are known, and as a result, for over a century little attention has been paid to the use of this equation for conformal mapping.

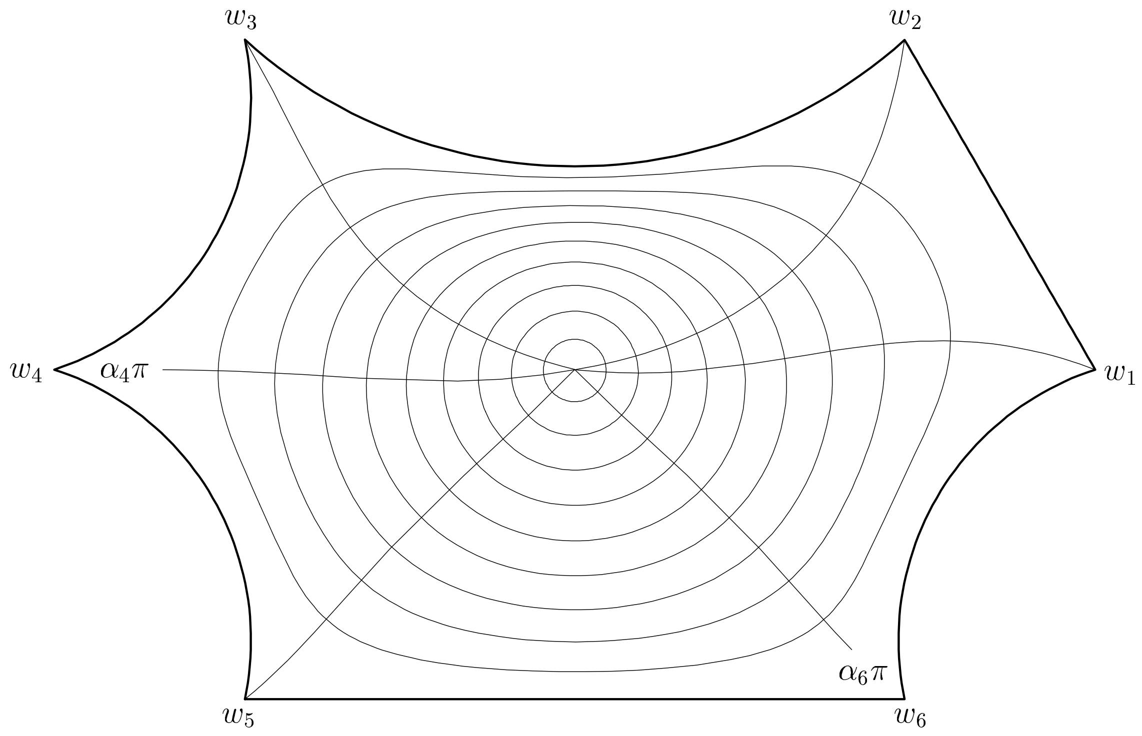

Figure 5-1: Conformal map of the unit disk onto a circular arc polygon. The lines shown are images of radii and concentric circles in the disk.

The only previous computational implementation of (5.1) that we are aware of is that of Bjørstad and Grosse [6] in 1987, which uses numerical methods both for integrating the differential equation and for determining the unknown prevertices z_j and the parameters β_j. Their code[1] features a fast technique for evaluating the mapping function $f(z)$ by means of power series expansions once the solution to the parameter problem is known. The procedure for solving the parameter problem itself, however, is not entirely satisfactory. Bjørstad and Grosse saw evidence of severe ill-conditioning in their system of nonlinear equations, which they attribute to poor understanding of the role of the accessory parameters β_j. In some cases their method failed to converge, and they were forced to resort to continuation techniques to obtain solutions.

In this chapter we will explore some of the problems involved with the use of equation

[1]Available from Netlib [17].

(5.1), and discuss methods for their resolution. While we have not entirely overcome the difficulties encountered by Bjørstad and Grosse, we have obtained improvements in a number of areas and have identified further obstacles to convergence that were not previously understood. Our goal, not yet realized, is to use the techniques of Chapter 4 to find a nonlinear system which is guaranteed to converge in all cases not dominated by crowding.

5.2 Derivation

The enigma in equation (5.1) is the behavior of the accessory parameters β_j. Unlike the constants z_j, which give the locations of the prevertices, and the α_j, which control the angles between boundary segments, the β_j do not have a straightforward geometric interpretation. In order to show why these parameters appear in the transformation at all, and to explain the dependencies among them which *are* known to exist, we will now briefly sketch a derivation of the Schwarzian differential equation. Many details omitted here can be found in [46] or [36].

Any linear transformation $L(w) = Aw + B, A \neq 0$, maps straight lines into straight lines and therefore polygons into polygons. Two polygons which are related by a linear transformation can be mapped from the upper half-plane using the same set of prevertices, so it is not surprising that the Schwarz-Christoffel formula (1.2) contains an arbitrary linear transformation in the guise of the two complex constants A and B. For circular arc polygons like the one shown in Figure 5-1, the corresponding role is played by the Möbius transformation

$$M(w) = \frac{aw+b}{cw+d}, \quad (ad - bc \neq 0). \tag{5.3}$$

(In the literature this expression is often referred to as a bilinear transformation or, even more confusingly, a linear transformation.) Thorough discussions of the properties of Möbius transformations can be found in [2, pp. 76–88] and [34, pp. 299–321]. For our purposes the most important property is that these functions always map circles into circles—straight lines being considered circles with zero curvature. We also note that the inverse of a Möbius transformation, and the composition of two Möbius transformations, are themselves Möbius transformations; and that exactly three complex or six real degrees of freedom are required to characterize any Möbius transformation.

Let $w = f(z)$ be a conformal map from the upper half-plane to the region shown in Figure 5-1. We can find an analytic continuation of this function into the lower half-plane via the symmetry principle (cf. [46, pp. 183–188]) by reflecting through any of the sides of the polygon. The procedure is similar to that illustrated in Figure 1-5 for the standard Schwarz-Christoffel transformation. If z is a point in the lower half-plane, then the reflected function can be expressed as $M^{-1}(\overline{M(f(\overline{z}))})$, where the bars denote complex conjugation and M is a Möbius transformation mapping the chosen side onto a segment of the real axis. The image of the lower half-plane is therefore also bounded by circular arcs. Reflecting again through another side of this new polygon, we find a continuation of the function back up into the upper half-plane. This is a new branch of the function f, reached by following a path around one or more prevertices in the z plane. The key point is that this new branch is simply a Möbius transformation of the original mapping.

Now consider the Schwarzian $\{f, z\}$, defined by (5.2). This expression is invariant under Möbius transformations, i.e., $\{M \circ f, z\} = \{f, z\}$, where M is any function of the form (5.3). Therefore $\{f, z\}$ is the same on every branch of f—it is a single-valued function defined on the entire complex plane with the exception of n isolated singularities at the prevertices z_j.

To determine the nature of these singularities, we can examine the asymptotic behavior of the mapping function near the corners of the circular arc polygon. Nehari [46] and Hille [36] both do this by performing a Möbius transformation sending a corner $w_j = f(z_j)$ to the origin, one of the sides meeting at this corner to the positive real axis, and the other to the ray $\arg[M(w)] = \alpha_j \pi$. Using the symmetry principle, they express the mapping function into this corner in the form $(z - z_j)^{\alpha_j} g_j(z)$, where g_j is a nonzero analytic function with a real Taylor series about z_j. The Schwarzian of this expression then yields

$$\{f, z\} = \{M \circ f, z\} = \frac{1}{2}\frac{1-\alpha_j^2}{(z-z_j)^2} + \frac{\beta_j}{z-z_j} + R_j(z), \tag{5.4}$$

where R_j is analytic at $z = z_j$ and

$$\beta_j = \frac{1-\alpha_j^2}{\alpha_j}\frac{g_j'(z_j)}{g_j(z_j)}. \tag{5.5}$$

There is a small gap in this treatment, which we point out because it will figure in the discussion of integration methods in the next section. When α_j is an integer—0, 1 and 2 are all possible—the two sides meeting at w_j are on tangent circles, so there is no Möbius transformation sending both into straight lines with a finite intersection. The best that can be done is to let $M(w_j) = \infty$; this sends both sides to straight lines which are parallel when α_j is an integer. Taking a derivative now sends these to lines of constant argument, which can be the positive real axis and the line $\arg[M(w)] = \alpha_j\pi$ if an appropriate rotation is included in M. The behavior of this mapping near z_j is of the form

$$\frac{d}{dz}(M \circ f)(z) = (z-z_j)^{-(\alpha_j+1)}G_j(z), \tag{5.6}$$

where again G_j is a nonzero analytic function with a real Taylor series about z_j. (The reason for taking the derivative is now apparent: When α_j is an integer, the expansion of $(M \circ f)(z)$ near z_j includes a logarithmic term.) The Schwarzian of $M \circ f$ is again of the form (5.4), where β_j is now

$$\beta_j = (\alpha_j+1)\frac{G_j'(z_j)}{G_j(z_j)}. \tag{5.7}$$

If we subtract off the singular terms at each prevertex, we obtain a function

$$\{f, z\} - \frac{1}{2}\sum_{j=1}^{n}\frac{1-\alpha_j^2}{(z-z_j)^2} - \sum_{j=1}^{n}\frac{\beta_j}{z-z_j} \tag{5.8}$$

which is analytic in the entire complex plane. Moreover, since the polygon is bounded, this function is also bounded and thus reduces to a constant. The Laurent expansion of $\{f, z\}$ near ∞ begins with the term of order z^{-4}, so the constant is zero and we are left with (5.1), the Schwarzian differential equation.

We can also obtain three relations among the parameters z_j, α_j, and β_j by expanding the right hand side of (5.1) near ∞ and setting the z^{-1}, z^{-2}, and z^{-3} terms to zero:

$$\sum_{j=1}^{n} \beta_j = 0, \tag{5.9}$$

$$\sum_{j=1}^{n} [2\beta_j z_j + (1 - \alpha_j^2)] = 0, \tag{5.10}$$

and

$$\sum_{j=1}^{n} [\beta_j z_j^2 + (1 - \alpha_j^2) z_j] = 0. \tag{5.11}$$

By considering the number of independent parameters required to describe a general circular arc polygon of n sides, Nehari [46] shows that no other such relations are possible.

Some minor changes to these formulas are necessary if one of the prevertices, say z_n, is at infinity. The constraint (5.11) implies $\beta_n = 0$, with $\beta_n z_n \sim -(1 - \alpha_n^2)$ as $z_n \to \infty$. The other two constraints thus become

$$\sum_{j=1}^{n-1} \beta_j = 0 \tag{5.12}$$

and

$$\sum_{j=1}^{n-1} [2\beta_j z_j + (1 - \alpha_j^2)] = (1 - \alpha_n^2); \tag{5.13}$$

these relations can also be derived directly by considering the behavior of $\{f, z\}$ near the prevertex at infinity. In the mapping function (5.1) itself the terms corresponding to z_n become 0, so both sums simply run from 0 to $n - 1$. While we do not use this version of the mapping function in any of our calculations, we mention it here because it will be useful in deriving the map from an infinite strip in Section 5.6.

5.3 Numerical Integration

In their algorithm, Bjørstad and Grosse use the unit disk instead of the upper half-plane as their standard domain. This requires a slight modification of the Schwarzian differential equation, as follows:

$$\{f, z\} = P[h(z)]h'(z)^2, \tag{5.14}$$

where $P(z)$ is the RHS of (5.1) and the Möbius transformation

$$h(z) = i\frac{1-z}{1+z} \tag{5.15}$$

maps the unit disk to the upper half-plane. In the course of solving the parameter problem, their program must integrate equation (5.14) many times to determine the shape of every approximate polygon. Each of the contours of integration runs from the center of the disk to a point midway between two prevertices; corner positions are then found by computing the points where sides intersect. Since there are no singularities on the paths of integration, a standard subroutine (Shampine and Gordon's ODE [58]) can be used to solve the equivalent system of six first-order real differential equations.

In order to compute a solution to (5.14) as an initial value problem, $f(0)$, $f'(0)$ and $f''(0)$ must be used. It is not possible to determine the correct values a priori, but fortunately the invariance of the Schwarzian under arbitrary Möbius transformations permits the initial values to be adjusted after the fact. Bjørstad and Grosse's program starts out with arbitrary initial conditions at the origin. Integrating along the radial contours reveals the shape of the image polygon corresponding to those initial values. Next, a Möbius transformation is used to send $f(0)$ to w_0, $f(z_n)$ to w_n, and the side containing w_n and w_1 to the proper circle. This last condition fixes both the curvature of this side and its tangent at w_n, so a total of six real degrees of freedom are involved—the number required to make the Möbius transformation, and the initial conditions, unique. It is now known what the initial conditions should have been to start with. On each succeeding iteration, with slightly different parameters z_j and β_j, the integration package uses the computed initial values from the *previous* iteration, thus yielding a polygon which typically requires only a small renormalization. Bjørstad and Grosse report that this procedure decreases the cost of the later integrations by a factor of up to four.

There are three drawbacks to this integration scheme that we have attempted to improve upon in our own algorithm. First, the method of determining corners as intersections breaks down when the two circles are tangent to one another, as occurs when α_j is an integer. Second, the contours of integration all reach from the center of the disk to its edge. Often a number of prevertices will be clustered together in a small group, making the integrals along each of several contours particularly difficult due to nearby singularities. Third, in integrating through the interior of the domain Bjørstad and Grosse are forced to use complex arithmetic, with a resulting decrease in efficiency.

All of these difficulties can be overcome by integrating along the boundary of the standard region from one prevertex to the next. In our own program we have returned to the upper half-plane as the standard domain, since this makes it trivial to use real arithmetic on the boundary. The behavior of the function at each singularity is known, and while we know of no alternative formulation of the problem which would remove every singularity from the real axis at once, it is not difficult to obtain analytic behavior at each prevertex in turn using n different changes of variables.

Separate solutions around each of the prevertices can be pieced together using Möbius transformations. Starting with arbitrary initial conditions at each prevertex, we integrate halfway to each of the adjacent prevertices. This gives us two sets of f, f' and f'' values at the midpoint of each boundary segment. Let $f_1(z)$ be the computed solution starting from z_1, the leftmost prevertex, and let $f_2(z)$ be the solution computed from z_2. We want to find the constants a, b, c, d for a Möbius transformation sending solution 2 into solution 1:

$$f_1 = \frac{af_2 + b}{cf_2 + d}. \tag{5.16}$$

Differentiating this equation twice yields

$$f_1' = \frac{ad - bc}{(cf_2 + d)^2} f_2' \tag{5.17}$$

and

$$f_1'' = \frac{ad - bc}{(cf_2 + d)^2} f_2'' - 2c\frac{ad - bc}{(cf_2 + d)^3}(f_2')^2, \tag{5.18}$$

from which we can solve for the desired constants:[2]

$$c = f_1' f_2'' - f_2' f_1'', \tag{5.19}$$

$$d = 2 f_1' (f_2')^2 - c f_2, \tag{5.20}$$

$$a = c f_1 + (c f_2 + d) \frac{f_1'}{f_2'}, \tag{5.21}$$

$$b = (c f_2 + d) f_1 - a f_2. \tag{5.22}$$

The transformation thus obtained is then applied to the values of f_2 and its derivatives computed for the point midway between z_2 and z_3, a new transformation is found to make f_3 match those values, and so on around the rest of the polygon. This gives a solution for the entire region, based on the initial conditions chosen at z_1, which can then be modified further to give any desired normalization.

As for the singularity removal scheme itself, we have two different methods currently in use. An advantage to working with a differential equation, rather than an integral, is that there is a much broader range of manipulations available than was discussed in Chapter 3. In particular, it is possible to obtain completely regular behavior near a prevertex, so there is no trouble with residual singular terms. Consider the method used in the previous section to classify the singularities. For nonintegral values of α_j there is always a solution near z_j which can be expressed in the form

$$(z - z_j)^{\alpha_j} g_j(z), \tag{5.23}$$

where g_j is a nonzero analytic function with a real Taylor series about z_j. From this solution, which maps the real axis near z_j onto two straight lines meeting at the origin, any other desired solution can be obtained via a Möbius transformation. By differentiating (5.23) three times and substituting for f and its derivatives in (5.1), we generate a nonsingular differential equation which can be solved for g_j. This equation is somewhat

[2]This is only one possible solution; if each of the values a, b, c, d given here is multiplied by the same nonzero constant, an equivalent transformation results.

complicated, and is most easily expressed in the first order system form used for the actual computation:

$$
\begin{aligned}
y_1' &= y_2 & (5.24)\\
y_2' &= y_2 y_3 + \frac{\alpha_j}{z - z_j} y_1 y_3 - \frac{\alpha_j + 1}{z - z_j} y_2 & (5.25)\\
y_3' &= \frac{1}{2} y_3^2 + \frac{\alpha_j - 1}{z - z_j} y_3 + \frac{\beta_j}{z - z_j} + R_j(z), & (5.26)
\end{aligned}
$$

where $y_1 = g_j$, $y_2 = g_j'$, $y_3 = \frac{f''}{f'} - \frac{\alpha - 1}{z - z_j}$, and $R_j = \sum_{k \neq j} [\frac{1}{2} \frac{1 - \alpha_k^2}{(z - z_k)^2} + \frac{\beta_k}{z - z_k}]$. Note that though the equations appear to be singular at z_j, the definition of β_j provides two initial conditions, $y_2(z_j) = \frac{\alpha_j \beta_j}{1 - \alpha_j^2} y_1(z_j)$ and $y_3(z_j) = \frac{\beta_j}{1 - \alpha_j}$, which cause the singularities to cancel each other out. The remaining initial value, $y_1(z_j)$, is arbitrary and we simply set it to 1.

The first-order system (5.24–5.26) is nonsingular and can be solved by a standard software package. When integrals over long contours are required, we follow Bjørstad and Grosse and use ODE [58] for this purpose. For shorter contours, such as those required for graphics, the simpler routine RKF45 [25] tends to be more efficient. (Ultimately Bjørstad and Grosse's series expansion method should replace RKF45 in this application, but we have not yet implemented this feature.)

Note: Some care must be taken in the implementation to compute accurate values for the derivatives at z_j. Though the initial conditions cause the $(z - z_j)^{-1}$ terms to cancel to leading order, thus eliminating the singularity, these terms do provide a finite contribution to y_2' and y_3'. Expanding g_j in a Taylor series about z_j provides the correct derivatives,

$$y_3'(z_j) = \left[\frac{1}{2} y_3^2(z_j) + R_j(z_j)\right] / (2 - \alpha_j) \tag{5.27}$$

and

$$y_2'(z_j) = [(1 + \alpha_j) y_2(z_j) y_3(z_j) + \alpha_j y_1(z_j) y_3'(z_j)] / (2 + \alpha_j). \tag{5.28}$$

In practice, this integration scheme works well when we integrate away from the real axis into the upper half-plane, but it occasionally fails when used along the real axis. This is because the transformation sending the two circular arcs adjacent to z_j into two straight lines may actually send a point on one of those arcs to infinity. The result is a movable singularity, one which is not associated with a singular term in the governing differential equation. In order to overcome this problem, we have made use of an alternative formulation of equation (5.1) described by Nehari [46] and Hille [36]. If u_1 and u_2 are two linearly independent solutions to the equation

$$u'' + Q(z)u = 0, \tag{5.29}$$

where

$$Q(z) = \sum_{j=1}^{n} \left\{ \frac{1}{4} \frac{1-\alpha_j^2}{(z-z_j)^2} + \frac{1}{2} \frac{\beta_j}{z-z_j} \right\}, \tag{5.30}$$

then the quotient $f(z) = u_1/u_2$ is a solution to (5.1). The effect is to replace the solution of a third-order nonlinear equation with two solutions of a second-order linear equation. For analytic work this is an immense improvement, and most existing knowledge about the analytic properties of the Schwarzian comes indirectly from the study of equation (5.29). For computational purposes the advantage is less clear, but the linearity of (5.29) does provide one key benefit: There are no spontaneous singularities. The points where f goes to infinity are simply the zeroes of u_2, which cause no difficulty for an integration package.

The coefficient function Q, and hence the solutions u_1 and u_2, do have singularities at the prevertices z_j, and these must be removed by a procedure similar to that which we used for f. If we assume that these singularities take the form $(z-z_j)^\gamma v_j(z)$ for some analytic function v_j and substitute this into equation (5.29), we obtain

$$v_j'' + \left\{ \gamma(\gamma-1) + \frac{1}{4}(1-\alpha_j^2) \right\} v_j (z-z_j)^{-2} + \left[2\gamma v_j' + \frac{1}{2}\beta_j v_j \right] (z-z_j)^{-1} + Q_j(z) v_j = 0. \tag{5.31}$$

(Here $Q_j = (R_j/2)$ denotes the terms of Q which are not singular at z_j.) In order to eliminate the $(z-z_j)^{-2}$ term the quantity in braces $\{\}$ must be zero, so γ has the two possible values $\gamma_1 = (1+\alpha_j)/2$ and $\gamma_2 = (1-\alpha_j)/2$ corresponding to the two linearly independent solutions. The quantity in square brackets [] must also be zero at z_j, so the initial condition $v_j'(z_j) = -(\beta_j/4\gamma)v_j(z_j)$ is determined automatically. Written as a first order system, equation (5.31) thus becomes

$$\begin{aligned}
y_1' &= y_2 && (5.32)\\
y_2' &= -\left[\frac{1+\alpha_j}{z-z_j}y_2 + \left(\frac{1}{2}\frac{\beta_j}{z-z_j} + Q_j\right)y_1\right] && (5.33)\\
y_3' &= y_4 && (5.34)\\
y_4' &= -\left[\frac{1-\alpha_j}{z-z_j}y_4 + \left(\frac{1}{2}\frac{\beta_j}{z-z_j} + Q_j\right)y_3\right], && (5.35)
\end{aligned}$$

with initial conditions $y_1(z_j) = 2(1+\alpha_j)$, $y_2(z_j) = -\beta_j$, $y_3(z_j) = 2(1-\alpha_j)$, and $y_4(z_j) = -\beta_j$. We use a single system of four equations rather than two systems of two equations so as not to duplicate the effort of computing Q_j. As before, since y_2 and y_4 are not constants, the $(z-z_j)^{-1}$ terms do contribute to y_2' and y_4' at z_j. A Taylor series expansion of v_j about z_j reveals that the proper initial behaviors are $y_2'(z_j) = -[\frac{1}{2}\beta_j y_2(z_j) + Q_j(z_j)y_1(z_j)]/(2+\alpha_j)$ and $y_4'(z_j) = -[\frac{1}{2}\beta_j y_4(z_j) + Q_j(z_j)y_3(z_j)]/(2-\alpha_j)$.

Once a solution to the system (5.32–5.35) has been computed—we again use ODE for this purpose—the mapping function f and its derivatives can be obtained via the relations

$$f(z) = (z-z_j)^{\alpha_j}\frac{y_1}{y_3}, \tag{5.36}$$

$$f'(z) = f(z)\left(\frac{y_2}{y_1} - \frac{y_4}{y_3} + \frac{\alpha_j}{z-z_j}\right), \tag{5.37}$$

and

$$f''(z) = f'(z)\left(\frac{\alpha_j - 1}{z-z_j} - 2\frac{y_4}{y_3}\right). \tag{5.38}$$

These values will differ from those computed from the nonlinear system (5.24–5.26) by the constant factor $(1+\alpha_j)/(1-\alpha_j)$, due to the differing initial conditions. This factor need only be taken into account when both methods are used from the same prevertex, as the matching procedure described earlier will otherwise make the proper correction automatically.

There are three cases remaining in which both methods described here will break down, those being when α_j takes on one of its possible integer values. As we showed in the previous section, $f(z)$ then has a logarithmic singularity, so the analytic function $g_j(z)$ does not exist. As for the linear equation (5.29), changing the dependent variable to v_j either provides only one solution or else fails to remove the singularity. When α_j is 0, the two values of γ are the same; the solution not accounted for has a logarithmic singularity. When α_j is 1 there are two distinct γ's, 1 and 0. Expanding both solutions v_j in Taylor series about z_j, however, shows that $z^0 y_3 = z^1(\beta_j/4)y_1$, so the computed solutions are identical up to a constant and another, logarithmic, solution exists. The case $\alpha_j = 2$ fails in a different way. The two distinct γ's yield two linearly independent solutions y_1 and y_3. Only one of these is analytic: y_3 has a logarithmic singularity starting with its second derivative. Note that the initial behavior for y_4', given above, becomes infinite as α_j approaches 2.

None of these problems is insurmountable. In each case the exact nature of the logarithmic singularity can be determined and an appropriate nonsingular formulation for the equation (5.29) derived. We have not yet implemented these special cases in our code, but plan to do so when the more pressing problems described in the following sections are resolved. The instance $\alpha_j = 1$, in particular, seems to be an important one for many applications.

5.4 Ill-Conditioning and the Accessory Parameters

In their computational experiments, Bjørstad and Grosse observed that the problem of determining the accessory parameters β_j for a given polygon is often ill-conditioned. In other words, the β's are nearly linearly dependent—small perturbations in different sets of parameters can produce almost identical changes in the image. The system of nonlinear equations for the parameter problem is thus poorly conditioned, and may be difficult or impossible to solve. Bjørstad and Grosse attributed this problem to a poor understanding of the geometric significance of the β's, but did not suggest any remedies.

In this section we present the idea that the β's are ill-conditioned because they rep-

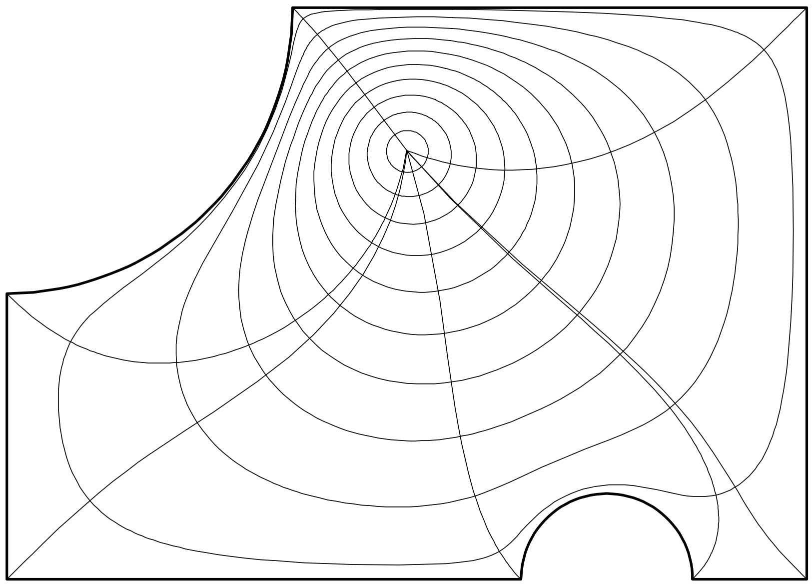

Figure 5-2: This bounded region is more naturally mapped from the unit disk than from the upper half-plane.

resent derivatives of the mapping function from the upper half-plane. Since these derivatives vary greatly between prevertices near the origin and those far out on the real axis, they are not the best choice for characterizing the image polygon. A better set of parameters can be obtained by mapping from the unit disk. (Though Bjørstad and Grosse computed integrals in the unit disk, they continued to use the parameters β_j defined in the half-plane.) We first derive these new parameters using a transformation from the upper half-plane to itself, then show how they can be used in formulating an equation for the disk map with no "preferred orientation". Numerical experiments illustrate how the conditioning of the nonlinear system is improved. Finally, we show that as an added benefit the new variables often generate less distorted initial guesses for use in solving the parameter problem.

Consider the region shown in Figure 5-2. In computing this map most of the actual computations were performed in the upper half-plane, for reasons discussed in the previous section. Some point on the boundary is therefore the image of the point at infinity.

The choice is essentially arbitrary—except for the difficulties involved in integrating the differential equation when one of the prevertices is at or near ∞, there is no pressing reason for picking any one point over another. Whatever point is chosen, however, the derivatives of the mapping function nearby will be very small. The β parameters at the adjacent prevertices will therefore be very sensitive to small perturbations, and the parameter problem will be poorly conditioned.

It is much more natural to consider mapping such a region from the unit disk. In fact, in our own computations we choose prevertices equally spaced around the circle as the initial guess, and then adjust those points via (4.5) while solving the parameter problem. (This does not prevent us from integrating equation (5.1) along the real axis, since it is a trivial matter to transform one domain into the other.) If we so much as rotate the prevertices around the disk, however, the β's all change! It seems clear from the outset that a set of parameters which were invariant under such rotations would give better results than the present system defined in the half-plane.

If we proceed directly to the mapping function from the disk at this point, we obtain a well-conditioned set of parameters which unfortunately is complex. In order to find a *real* set of variables which can be used in the parameter problem, it is better to first consider a transformation of the upper half-plane to itself. Let us define the correspondence between the half-plane and the disk by the two maps

$$z = h(\zeta) = i\left(\frac{1-\zeta}{1+\zeta}\right) \tag{5.39}$$

and

$$\zeta = h^{-1}(z) = \frac{i-z}{i+z}, \tag{5.40}$$

where for clarity we use ζ to denote points in the disk. Taking any prevertex z_j on the real axis, we first map it to the disk, then rotate the disk so as to bring that point to 1, then invert the map to send it back to a transformed version of the half-plane. Composing these maps gives a transformation

$$s = \frac{z - z_j}{z_j z + 1} \tag{5.41}$$

of the upper half-plane to itself which corresponds to a rotation of the disk and sends the prevertex z_j to the origin.

The motive behind all this is to define a new parameter β'_j as the value that β_j assumes when z_j is rotated to the origin. Using equations (5.6) and (5.7) along with the relations

$$z - z_j = \frac{(z_j^2+1)s}{1 - z_j s} \tag{5.42}$$

and

$$\frac{ds}{dz} = \frac{(1 - z_j s)^2}{z_j^2 + 1}, \tag{5.43}$$

we proceed as follows:

$$(z - z_j)^{-(\alpha_j+1)} G(z) = s^{-(\alpha_j+1)} G_s(s) \frac{ds}{dz}, \tag{5.44}$$

$$s^{-(\alpha_j+1)} \left(\frac{z_j^2+1}{1 - z_j s} \right)^{-(\alpha_j+1)} G(z) = s^{-(\alpha_j+1)} G_s(s) \frac{(1 - z_j s)^2}{z_j^2+1}, \tag{5.45}$$

$$G(z) = (z_j^2+1) \left(\frac{1 - z_j s}{z_j^2+1} \right)^{-(\alpha_j-1)} G_s(s), \tag{5.46}$$

$$\frac{G'(z)}{G(z)} = \frac{(1 - z_j s)^2}{z_j^2+1} \left[\frac{G'_s(s)}{G_s(s)} + (\alpha_j - 1) \frac{z_j}{1 - z_j s} \right], \tag{5.47}$$

$$\frac{G'(z_j)}{G(z_j)} = \frac{1}{z_j^2+1} \left[\frac{G'_s(0)}{G_s(0)} + (\alpha_j - 1) z_j \right], \tag{5.48}$$

$$\beta_j = \frac{1}{z_j^2+1} \left[\beta'_j - (1 - \alpha_j^2) z_j \right]. \tag{5.49}$$

Theorem 5.1 *The new parameters β'_j defined by (5.49) are invariant under arbitrary rotations of the unit disk.*

Proof: Each prevertex ζ_j is rotated to 1 before being mapped back to the upper half-plane, so the effects of any prior rotations of the disk are eliminated. ∎

We have not yet shown that these new parameters play any particular role in the map from the disk itself. For this purpose, and because this form of the mapping function may be useful for some applications, we will now proceed to derive a differential equation for the map from the unit disk to a circular arc polygon. We have found no references to this or any similar equation in the literature, though given the age of (5.1) it seems unlikely that the disk map should never have been derived before.

To begin with, it is easy to show that

$$\{f, \zeta\} = [h'(\zeta)]^2 \{f \circ h^{-1}, z\}, \tag{5.50}$$

where $f(\zeta)$ is the mapping function from the disk and $(f \circ h^{-1})(z)$ is therefore the mapping function from the upper half-plane. (This property of the Schwarzian holds for any Möbius transformation of the independent variable, not just h.) We can now write the Schwarzian differential equation in terms of ζ:

$$\{f, \zeta\} = \left[\frac{-2i}{(1+\zeta)^2}\right]^2 \left\{\frac{1}{2}\sum_{j=1}^{n} \frac{1-\alpha_j^2}{[h(\zeta)-h(\zeta_j)]^2} + \sum_{j=1}^{n} \frac{\beta_j}{h(\zeta)-h(\zeta_j)}\right\}. \tag{5.51}$$

Expanding the individual terms and simplifying then gives, after a bit of algebra,

$$\{f, \zeta\} = \frac{1}{2}\sum_{j=1}^{n} \frac{1-\alpha_j^2}{(\zeta-\zeta_j)^2}\frac{(1+\zeta_j)^2}{(1+\zeta)^2} + \sum_{j=1}^{n} \frac{-2i\beta_j}{(\zeta-\zeta_j)}\frac{(1+\zeta_j)}{(1+\zeta)^3}. \tag{5.52}$$

The most striking features of equation (5.52) are the factors $(1+\zeta)$ in the denominators of both sums. The mapping function from the disk should be completely symmetrical; it should not distinguish any special points on the circle; and it certainly should not have singularities which seem to be completely independent of the prevertex positions! These apparent singularities are not really present; they are artifacts of the change of variables which sends the point at infinity in the upper half-plane to the point -1 on the circle. To eliminate them we must separate them from the legitimate singular factors $(\zeta - \zeta_j)$, then apply the relations (5.9–5.11) which control the behavior of (5.1) at ∞.

A partial fraction expansion of each term yields, after much manipulation,

$$\begin{aligned} \{f, \zeta\} \;=\; & \sum_{j=1}^{n} \left\{ \frac{1}{2}\frac{1-\alpha_j^2}{(\zeta-\zeta_j)^2} - \frac{1-\alpha_j^2}{(\zeta_j+1)(\zeta-\zeta_j)} - \frac{2i\beta_j}{(\zeta_j+1)^2(\zeta-\zeta_j)} + \right. \\ & \frac{2i\beta_j}{(\zeta+1)^3} + \\ & \frac{1}{2}\frac{1-\alpha_j^2}{(\zeta+1)^2} + \frac{2i\beta_j}{(\zeta_j+1)(\zeta+1)^2} + \\ & \left. \frac{1-\alpha_j^2}{(\zeta_j+1)(\zeta+1)} + \frac{2i\beta_j}{(\zeta_j+1)^2(\zeta+1)} \right\}, \end{aligned} \tag{5.53}$$

where each row contains the terms featuring a different power of $(\zeta+1)$. If we express the constraints (5.9–5.11) in terms of ζ, they become

$$\sum_{j=1}^{n} \beta_j = 0, \tag{5.54}$$

$$\sum_{j=1}^{n} \left[2i\beta_j \frac{1-\zeta_j}{1+\zeta_j} + (1-\alpha_j^2) \right] = 0, \tag{5.55}$$

and

$$\sum_{j=1}^{n} \left[-\beta_j \left(\frac{1-\zeta_j}{1+\zeta_j} \right)^2 + i \left(\frac{1-\zeta_j}{1+\zeta_j} \right) (1-\alpha_j^2) \right] = 0; \tag{5.56}$$

with these relations all the terms of (5.53) except the first row drop out, and we are left with

$$\{f,\zeta\} = \sum_{j=1}^{n} \left\{ \frac{1}{2} \frac{1-\alpha_j^2}{(\zeta-\zeta_j)^2} - \frac{1-\alpha_j^2}{(\zeta_j+1)(\zeta-\zeta_j)} - \frac{2i\beta_j}{(\zeta_j+1)^2(\zeta-\zeta_j)} \right\}. \tag{5.57}$$

The differential equation (5.57) is free of singularities caused by factors of the form $(\zeta+1)$. It is not yet symmetrical, though, because the factors (ζ_j+1) still distinguish -1 as a "special" point. All of these factors should now be absorbed into a new definition for the accessory parameters. We could, for example, simply declare a new set of variables β_j^c to be defined as

$$\beta_j^c = -\frac{1-\alpha_j^2}{\zeta_j+1} - \frac{2i\beta_j}{(\zeta_j+1)^2}, \tag{5.58}$$

putting the equation for the disk map into exactly the same form as (5.1). These variables could not be used in the parameter problem, however, because they are complex; this is the situation we alluded to above in the discussion of the parameters β_j'. If we use the definition (5.49) to substitute for β_j in (5.57), we can finally simplify the equation to the following symmetrical form:

Theorem 5.2 *Any conformal map from the unit disk to a circular arc polygon satisfies the differential equation*

$$\{f,\zeta\} = \sum_{j=1}^{n} \frac{1}{2} \left\{ \frac{1-\alpha_j^2}{(\zeta-\zeta_j)^2} - \frac{1-\alpha_j^2}{\zeta_j(\zeta-\zeta_j)} - \frac{i\beta_j'}{\zeta_j(\zeta-\zeta_j)} \right\}, \tag{5.59}$$

where the parameters β'_j *are real and subject to the three equality constraints*

$$\sum_{j=1}^{n} \frac{1}{\zeta_j}[i\beta'_j + (1-\alpha_j^2)] = 0, \tag{5.60}$$

$$\sum_{j=1}^{n} \beta'_j = 0, \tag{5.61}$$

and

$$\sum_{j=1}^{n} \zeta_j[i\beta'_j - (1-\alpha_j^2)] = 0. \tag{5.62}$$

Proof: We showed in Section 5.2 that any map from the upper half plane to a circular arc polygon satisfies the Schwarzian differential equation (5.1). Equation (5.59) is derived from (5.1) by a change of variables. The constraints can be obtained from the equivalent relations (5.9–5.11) by substituting for z_j and β_j and rearranging terms. ▮

It is not difficult to confirm that the three constraints on the parameters β'_j correspond to the first three terms of the Laurent expansion for $\{f, \zeta\}$ at infinity. These terms must all be 0 since the reflection principle requires all branches of f to be analytic at ∞, just as in Section 5.2.

It is now something of an anticlimax to observe that, experimentally, the parameters β'_j are typically better-conditioned than the β's, but still not well-conditioned. The improvement is usually about an order of magnitude, and varies with the problem and the specific nonlinear system being used. For the region shown in Figure 5-2, for instance, using pure coordinates for the nonlinear system—see the next section—the condition number of the Jacobian at the solution decreases from 6564 to 3145. Using curvatures and arc lengths for the system, the change is more noticeable, from 961059 to 46176. A distorted version of this polygon (Figure 5-4b) has a decrease from 2641 to 490, also using curvatures and arc lengths. The polygon shown in Figure 5-3 had decreases of 79610 to 22989 and 73559 to 10700, respectively, measured at the solution in the two systems mentioned. There are definite improvements; the condition numbers almost always decrease; but the nonlinear systems remain somewhat difficult to solve.

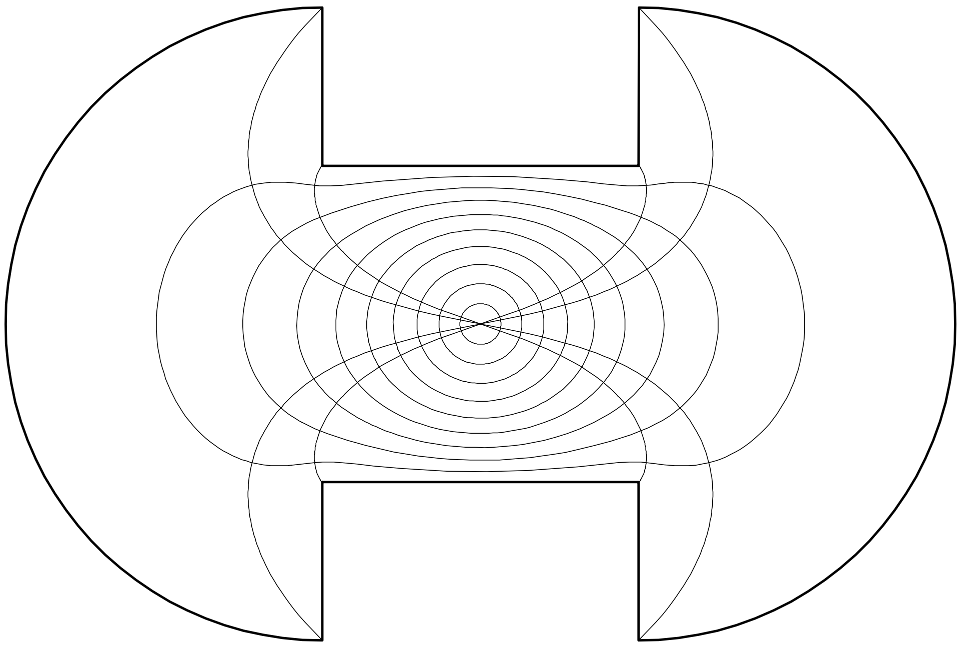

Figure 5-3: An 8-sided circular arc polygon. No advantage has been taken of symmetry in computing this map.

These regions do not seem sufficiently distorted to deserve condition numbers in the thousands, so something may still be missing from our understanding of the behavior of the mapping function. Both polygons have regions that are very slightly crowded, for example near the lower right corner of Figure 5-2. This degree of crowding is almost unavoidable in any problem. Acting on the assumption that crowded prevertices would tend to have larger β's than prevertices in uncrowded regions, we have experimented with adding additional scaling factors to the parameter definitions based on the distances to neighboring prevertices. So far these experiments have not produced any conclusive results one way or the other.

Fortunately, there is one area in which the new set of parameters has provided a definite improvement. Lacking any scheme for making a good initial guess to the parameter problem, we have followed Bjørstad and Grosse in letting the prevertices be equally spaced on the disk and setting the first $n-3$ accessory parameters to 0. (The remaining 3 parameters must then be chosen to satisfy the constraints (5.9–5.11).) Since the β''s

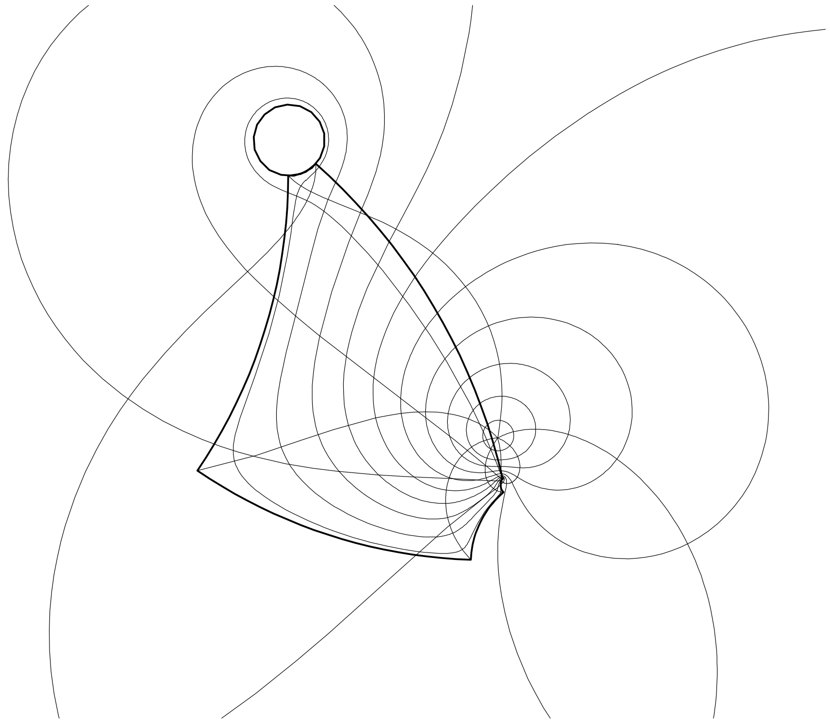

Figure 5-4a: Initial guess for the 7-sided polygon of Figure 5-2 using the standard parameters β_j.

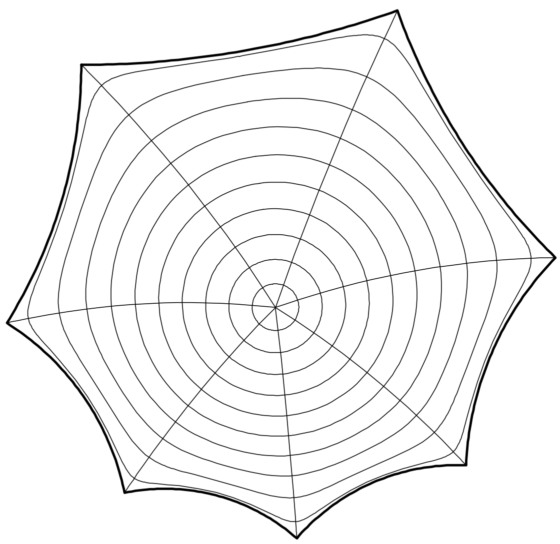

Figure 5-4b: Initial guess for the same polygon using the modified parameters β_j'.

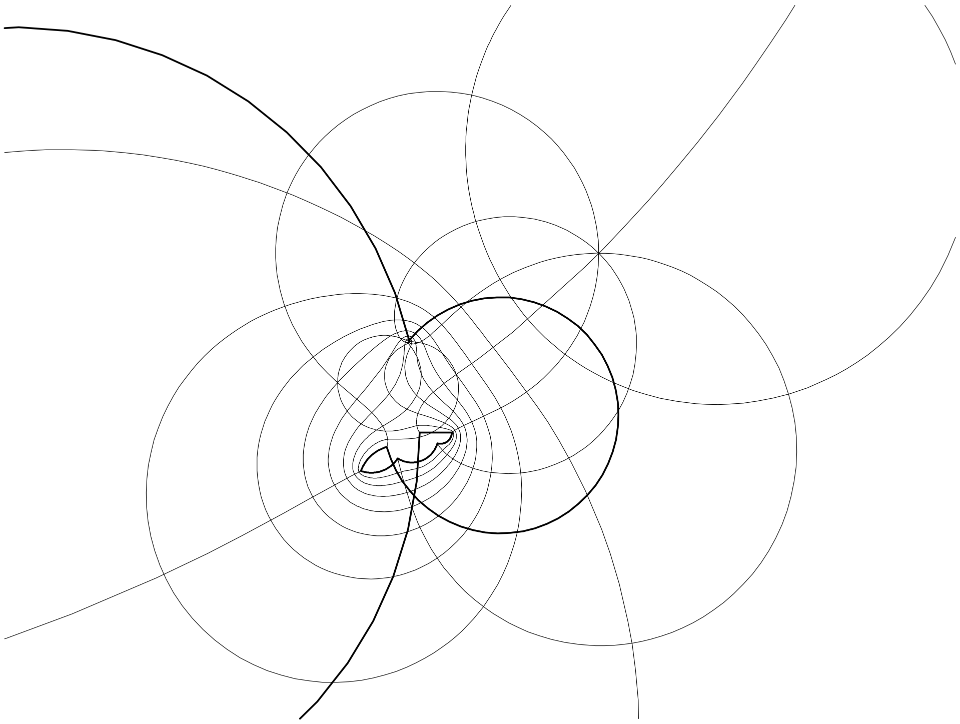

Figure 5-5a: Initial guess for the 8-sided polygon of Figure 5-3 using the standard parameters β_j.

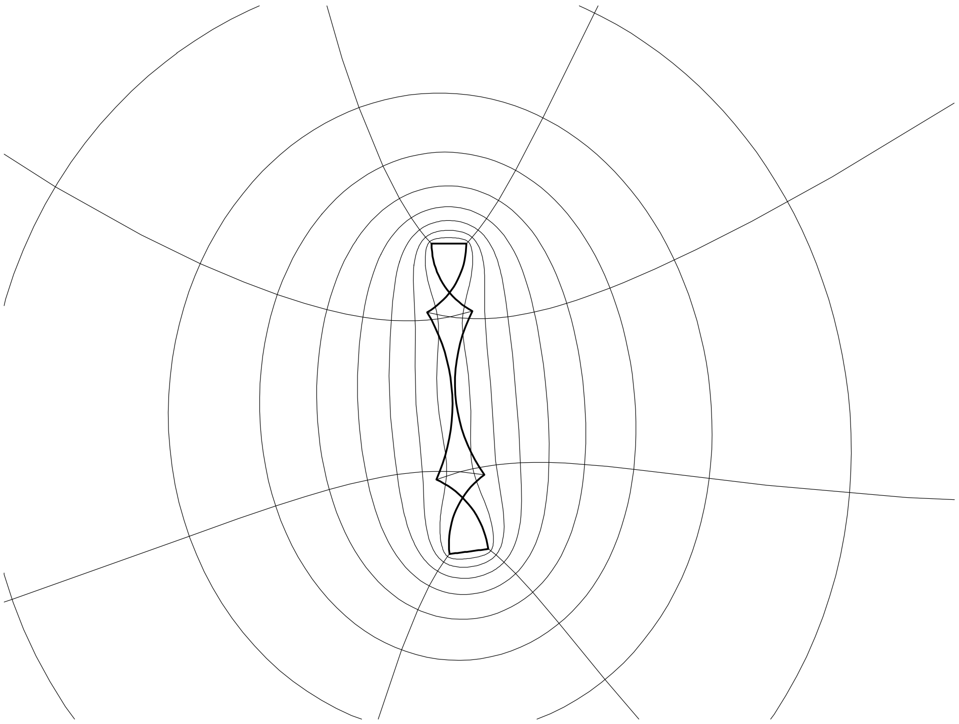

Figure 5-5b: Initial guess for the same polygon using the modified parameters β'_j.

are in a sense more balanced than the β's, it is not surprising that they tend to yield less distorted initial guesses. Figure 5-4 shows the results of both initial guesses for the polygon in Figure 5-2; Figure 5-5 shows the same thing for the polygon in Figure 5-3. In both sets of plots the β version is above and the β' version is below. We should point out that in both of these problems the β' version takes advantage of the symmetries present in the distribution of angles—we have observed significant improvements, however, even on polygons without such symmetries.

5.5 Formulation of the Nonlinear System

Despite the difficulties discussed at the end of Chapter 4, any straight-sided polygon can at least be uniquely characterized by a linear transformation and a nonlinear system of equations. All points of a bounded polygon must remain finite no matter how the region is distorted. Normalization is therefore quite trivial: Just fix the positions of any two adjacent corners. As for the nonlinear system, not only are the angles between sides constant, but each side always maintains the same orientation relative to the others. Consequently, any two sides will have the same inclination to each other in every possible version of the polygon. So long as these sides are not parallel, specifying the lengths of the $n-2$ others will uniquely characterize the polygon, not only at the desired solution, but in all possible distorted solutions as well. The resulting diffeomorphism between the space of parameters and the space of residuals makes the nonlinear system much easier to solve, even though it does not by itself provide sufficient conditions for global convergence of a nonlinear equations algorithm. This picture becomes slightly more complicated when some of the vertices are at infinity, but no insurmountable difficulties arise.

Circular arc polygons, on the other hand, are much more difficult to characterize uniquely. In Figures 5-4 and 5-5, note the ways in which such a region can be distorted: It can turn itself inside out by passing a boundary through the point at infinity, sometimes more than once. It can intersect itself in extremely complicated ways. Individual sides can sometimes wrap around through more than 360 degrees. Even problems as well-behaved as the ones shown here can tax the ability of the human eye, much less a single nonlinear system, to identify the way to the solution. And these examples all have prevertices evenly spaced around the unit circle! In crowded problems the situation can become far worse.

Bjørstad and Grosse recognize some of these difficulties in [6], and attempt to remedy them by using an overdetermined system of equations to solve the parameter problem. In this section we will show that their approach has several deficiencies, most of which can be attributed to a failure to adequately consider the behavior of the model far from the desired solution. Specifically, their normalization is not always well-defined, their characterization does not always uniquely determine the target polygon, and their use of an overdetermined system is inefficient and may actually make the algorithm less robust than one based on a square system of equations. Though we do suggest a few improvements, however, we must make it perfectly clear that we have been unable to find satisfactory solutions to these problems ourselves. We regret that the results of this section are mostly negative, but hope that by identifying many of the trouble spots encountered by the current methods we may suggest some new insight which may eventually produce a successful solution algorithm.

To begin with, it is not difficult to see that the corner positions alone do not completely determine a circular arc polygon. Bjørstad and Grosse give a straightforward example to demonstrate this: Take a square, and squeeze two opposite sides inward. The other two sides bulge outward, keeping all four angles constant as well as leaving the corners in their original positions. The only change has been in the orientations of the corners, i.e., each side now approaches each of its endpoints from a different direction. Figure 5-6 shows a slightly more complicated example of the same phenomenon, a distortion of Figure 5-3. The symmetry of these figures is not important; the same type of deformation can be applied to any circular arc polygon with an even number of sides.

Taking this additional degree of freedom into account, Bjørstad and Grosse present

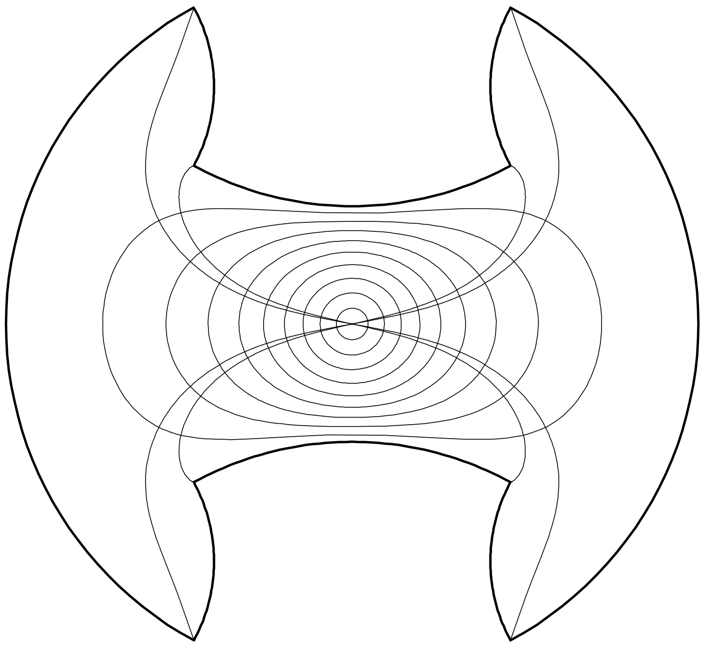

Figure 5-6: This region has exactly the same corners and angles as Figure 5-3. Only the orientations of the corners have been changed.

the following algorithm for solving the parameter problem:

Algorithm 5.1 (Bjørstad and Grosse) *Normalize the polygon using a Möbius transformation to send* w_0*—the image of the center of the disk—and* w_n *to their correct positions, and the side containing* w_n *and* w_1 *to the correct circle. This fixes both the direction of the tangent at* w_n *and the curvature of the side, so six real degrees of freedom are involved and the Möbius transformation is completely determined. The polygon is then characterized by the positions of the other* $n-1$ *vertices, so compute the unknown parameters by solving the* $2n-2$ *real equations*

$$w_j - w_j^* = 0, \quad 1 \le j \le n-1. \tag{5.63}$$

There are only $2n-4$ real parameters to be determined, $n-1$ prevertices and $n-3$ independent β's, so the nonlinear system (5.63) is overdetermined. Bjørstad and Grosse therefore use the nonlinear least squares routine N2F, formerly known as NL2SOL [12,13], to solve the parameter problem. We believe that this overdetermined approach has two

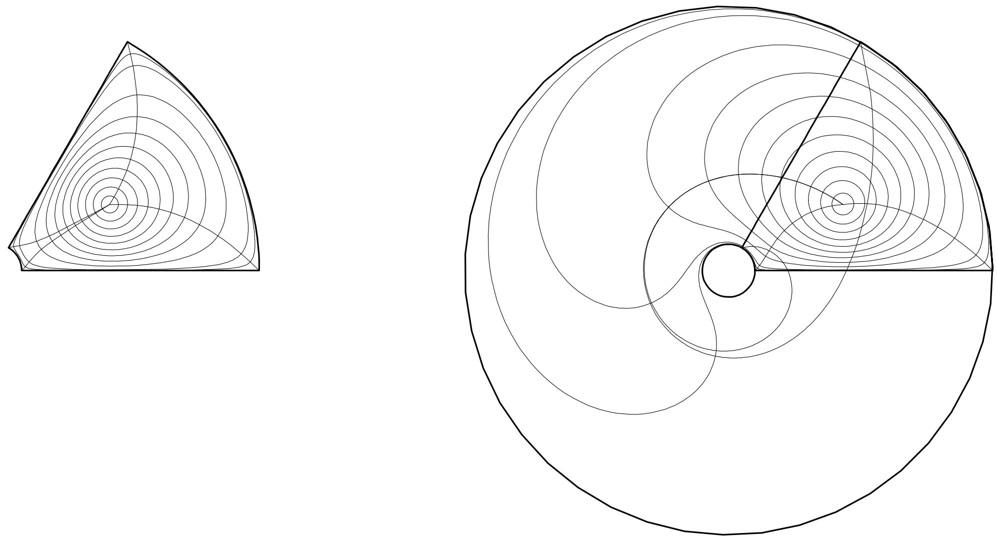

Figure 5-7: These two regions demonstrate that even specifying the corners *and* their orientations is not always sufficient to characterize a polygon.

serious shortcomings which are independent of the actual system being used:

1. It is less efficient. A nonlinear least squares solver like N2F must reevaluate the Jacobian of the transformation at every iteration, whereas methods for square systems, e.g. Minpack [44], can use secant updates. As the number of corners increases, N2F therefore takes $O(n)$ times longer per iteration, a liability which is only partially offset by the fact that it often requires fewer steps to reach the solution.

2. Instead of mapping the parameter space $\mathbf{R}^{2n-4}$ to an open subset of itself, the residual function maps $\mathbf{R}^{2n-4}$ to a $(2n-4)$-dimensional hypersurface in a higher-dimensional space. The results in Chapter 4 obtained from the diffeomorphism F therefore no longer apply, and the system may have any number of local minima which could cause a solution method to fail.

Aside from these difficulties with the solution process, however, the nonlinear system (5.63) may have more than one solution.

Theorem 5.3 *The n corners w_j and the circle containing one of the sides are not sufficient to uniquely define a circular arc polygon.*

Proof: The two regions shown in Figure 5-7 are completely identical except for the lengths of two of the sides. The error in Bjørstad and Grosse's proof ([6], Characterization 1, p. 26) is that they do not consider the possibility of a side extending through more than 360 degrees. ∎

There is also a problem with the way the polygon is normalized. Bjørstad and Grosse report that their algorithm often causes much greater distortion to the image polygon than seems necessary, but that they were not able to suggest a better alternative. They did not mention, however, that the normalization is not always even well-defined.

Theorem 5.4 *It is not always possible to find a Möbius transformation sending w_0 and w_n to the correct positions and the side containing w_n and w_1 to the correct circle. In addition, when the transformation does exist it may reverse the orientation of the side along that circle.*

Proof: Consider the region shown in Figure 5-8a. The center of the disk is mapped to a point inside the circle containing the side from w_n to w_1. Suppose, however, that the desired solution sends the center of the disk to the point shown by the black dot instead, which is outside the circle. Changing the normalization to send the center to this new point yields the region shown in Figure 5-8b. The side from w_n to w_1 is still on the same circle, but it has reversed its direction. This new figure is not an acceptable starting point for solving the parameter problem, because it can not be continuously deformed into the desired solution. Moreover, if the center of the disk were to be mapped to a point exactly on the circle, there would be no normalization sending it to the desired point. ∎

Given the shortcomings of Bjørstad and Grosse's approach, what changes would be necessary to make the solution algorithm better-behaved? Ideally, of course, we would like to find a nonlinear system with which convergence could be guaranteed from any starting point. Failing that, however, a square system of equations defining a diffeomorphism between the space of parameters and the space of residuals would at least provide the positive results discussed in Section 4.2. Ignoring differentiability for the moment, a

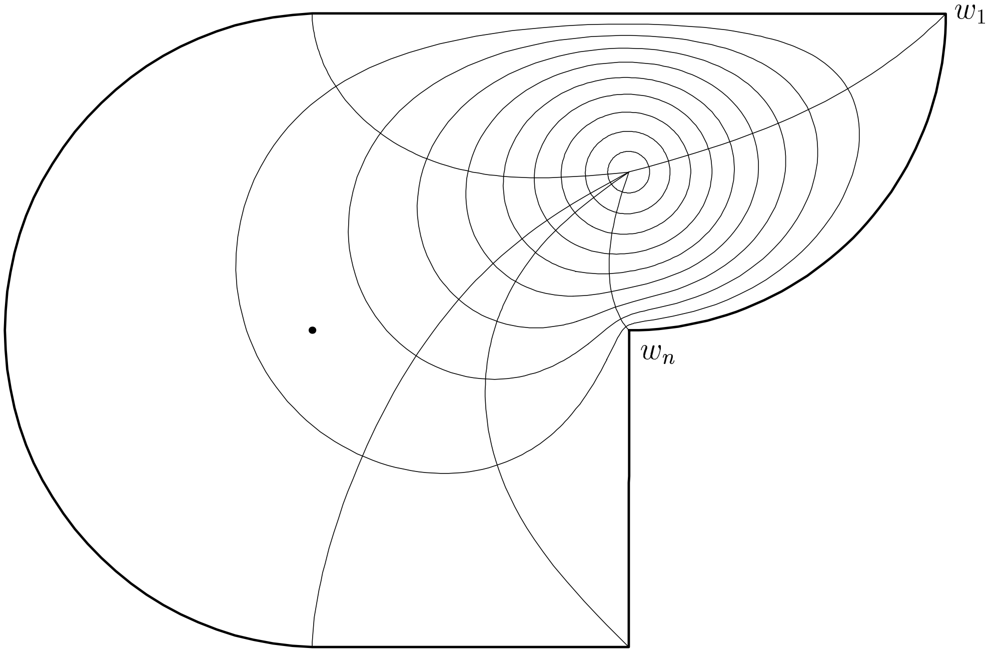

Figure 5-8a: This conformal map to a circular arc polygon does not send the center of the disk to the desired point, identified here by a black dot.

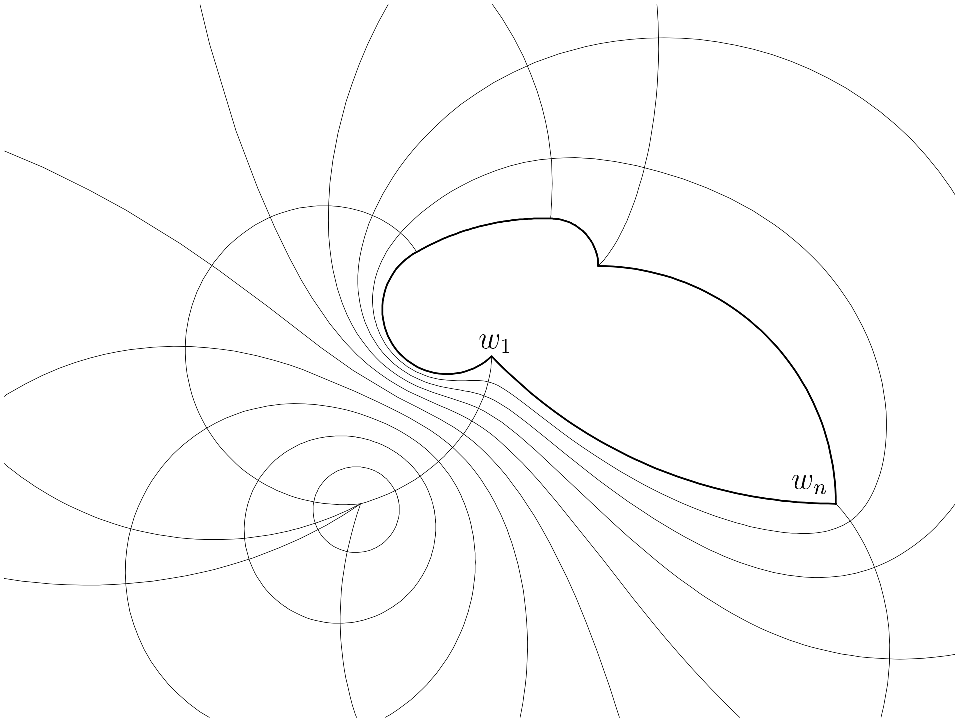

Figure 5-8b: Changing the normalization to send the center to the correct point, however, inverts the polygon and reverses the orientation of the side from w_n to w_1.

homeomorphism between the two sets requires the following conditions:

1. The normalization must be well-defined for every possible distortion of the image polygon.

2. The residuals F_j must be uniquely defined for every legal polygon, and must be finite for all finite parameter vectors x.

3. No two distinct polygons may yield the same residual vector.

Given these conditions, it should be a simple matter to insure differentiability, thus making the system define a diffeomorphism.

Any normalization which is to completely define a Möbius transformation must involve six real degrees of freedom. The trouble is, when the polygon intersects itself any points used to determine the normalization may turn out to be the same point. It is quite possible to twist a circular arc polygon so that $w_0 = w_n$, $w_1 = w_n$, or whatever else may be required to make the transformation undefined. If we restrict our attention to the behavior of the map at a single point, say w_n, there are only five real quantities that can always be specified: the real and imaginary coordinates of the point, the orientation of the angle at that corner, and the curvatures of the two sides meeting there. It is not possible to specify the length of either of those sides since they may wrap through more than 360^o. The only solution to this problem may be to determine one of the factors in the Möbius transformation by adding an additional equation to the nonlinear system. We have not yet performed any experiments with such an algorithm, however.

Since the corners alone do not uniquely define a circular arc polygon, our characterization must instead take into account some measures of the individual sides. The curvature is one obvious quantity that can be used—its inverse, the radius of curvature, is unsuitable because it is infinite when a side is a straight line. For the length of the side there are two straightforward choices: the length of the arc and the angle over which it extends. Neither is completely satisfactory. Arc length is infinite whenever the side in question passes through infinity; that this possibility can not be ruled out is amply

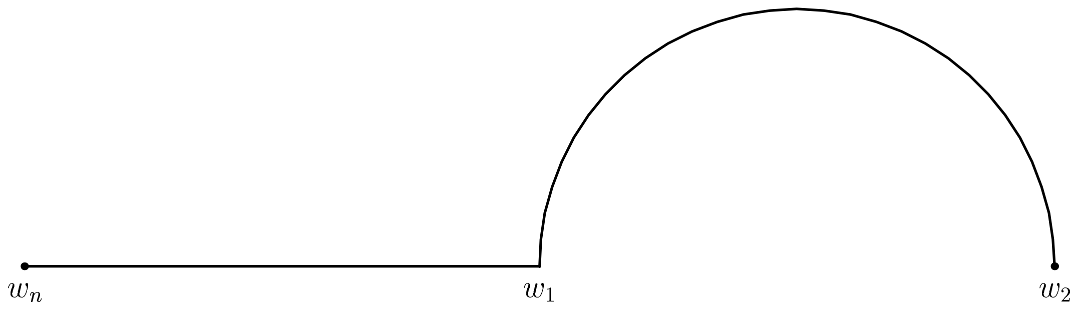

Figure 5-9: The two solid dots, and the orientations of the lines at those points, are assumed to be determined by the remainder of the polygon (not shown). Holding the curvature of the left side fixed is not sufficient to make the length of that side, and thus the position of w_1, unique.

demonstrated by the examples we have already shown. Curvature and angle, however, both become zero simultaneously when a side is a straight line, and thus do not determine the length of such a side.

One promising way to proceed is to map the problem onto the Riemann sphere (cf. [34]). Curvatures and arc lengths can be measured on the sphere without encountering either of the problems just mentioned. It is not at all obvious, however, just which projection to the sphere should be used. This difficulty is intimately connected with the normalization problem: What Möbius transformation should be applied to the polygon before mapping it onto the Riemann sphere? Ideally, we would use measurements which were in some way independent of the normalization. Our investigations have failed to discover any such system, however.

One final difficulty should be mentioned before we bring this discussion to a close. With straight-sided polygons there are always two "extra" sides which do not enter into the parameter problem, but which are determined by the other sides of the polygon. With

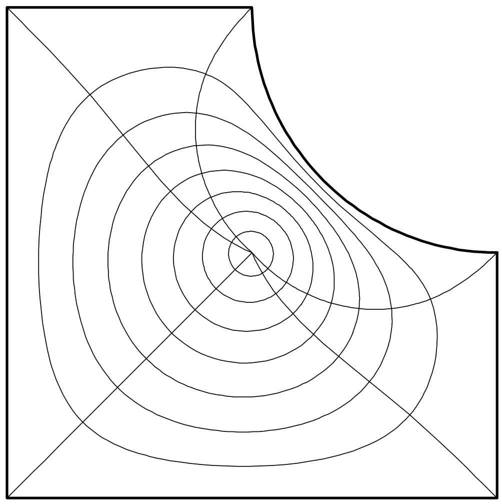

Figure 5-10: A simple example based on one computed by Bjørstad and Grosse.

circular arc polygons, whatever approach is taken to measuring sides, one always arrives at a situation where the last two sides must be determined by the last component of the residual vector. In most cases controlling the curvature of one of the sides appears to give a unique solution; the other curvature and the lengths of the two sides are determined by the known angle at the remaining corner. In a case like that shown in Figure 5-9, however, fixing the curvature of the side on the left does not prevent the semicircle from expanding or shrinking. The key factor here is that w_2 is on the circle containing the side from w_n to w_1. Controlling the length of one of the sides may be a better approach, but this still permits multiple solutions when one or both sides wraps through more than 360^o. We do not know whether or not such wrapped solutions can be eliminated by considering the geometry of the remainder of the polygon.

Though we have not succeeded in designing a single system with all of the elegant properties desired, we have implemented bits and pieces of almost all of the methods described here. (We could spend several pages describing implementational details, but without a consensus as to which approach is best there seems to be little point. Most of the formulas simply boil down to plane or spherical trigonometry in any case, and are

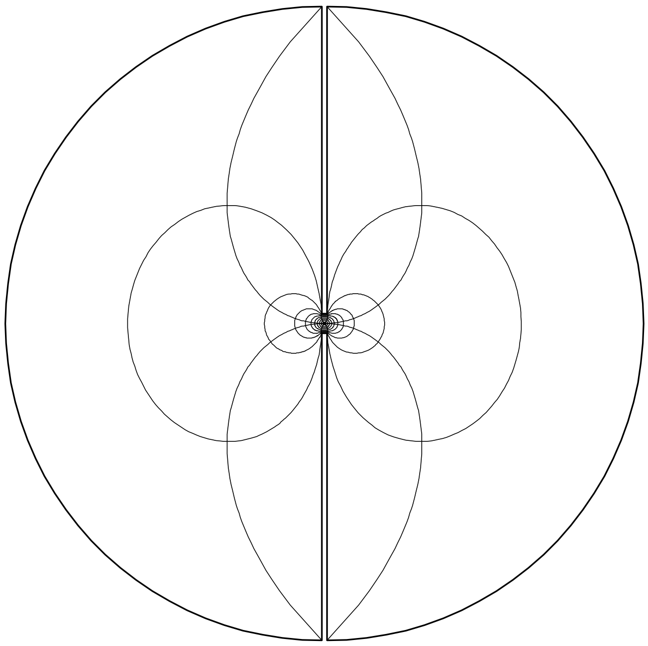

Figure 5-11: Bjørstad and Grosse's algorithm fails to converge to this region without the continuation method. Most of the Riemann sphere algorithms succeed, however.

more interesting as brain-teasers than for any real mathematical content.) Our program is therefore equipped with a truly bewildering array of "knobs and switches", with which we feel quite confident of eventually finding the solution for any given polygon within the scope of these techniques. Two final examples are shown in Figures 5-10 and 5-11: The first requires 118 function evaluations for Bjørstad and Grosse's program, while most of our current methods finish in under 50. The second was reported as requiring continuation techniques, but most of our Riemann sphere variants converge even from fairly bad starting points. These results are not much to brag about, but do show at least some improvement. There is a big difference, however, between a research computer program that works reasonably well for the researcher, and a distributable program that can be operated by anyone with a general understanding of the subject. Without a more robust formulation of the parameter problem, this latter goal must remain at least temporarily out of reach.

5.6 Mapping from a Strip

In Chapter 2 we discussed the crowding phenomenon and its specific consequences for the computation of Schwarz-Christoffel transformations. The same problem applies with equal force to the use of the Schwarzian differential equation (5.1): As the region being mapped becomes increasingly elongated, groups of prevertices cluster exponentially close together and quickly become indistinguishable in finite-precision arithmetic. The solution to the problem is also the same: Use an elongated region such as an infinite strip as the standard domain.

Lacking a geometric explanation for the effects of the individual terms in (5.1), we can not use the term-by-term approach described in Chapter 2, but must instead fall back on a direct change of variables for deriving an equation for the map from a strip. (As we discussed before, this approach involves more algebra and less understanding, but in the circular arc case there appears to be no alternative.) Specifically, we switch from the upper half-plane to a strip of unit width by substituting $e^{\pi z}$ for z in all the appropriate formulas. (We do not introduce a new variable, like the ζ of Section 5.2, for instance, in order to maintain consistency with the usage in Chapter 2.) The left and right ends of the strip therefore correspond to the points 0 and ∞ on the real axis, respectively. As with the mapping function for straight-sided channels, we distinguish the right end from the n finite prevertices by calling it prevertex number 0; the only explicit reference to the left end is by way of the angle to which it maps, which we will denote by $\alpha_\infty \pi$.

Since one of the prevertices in the upper half-plane is at ∞, the appropriate form of the transformation is the one discussed at the very end of Section 5.2. Substituting $e^{\pi z}$ for z gives the differential equation

$$\{f, e^{\pi z}\} = \frac{1}{2}(1-\alpha_0^2)e^{-2\pi z} + \beta_0 e^{-\pi z} + \frac{1}{2}\sum_{j=1}^{n} \frac{1-\alpha_j^2}{(e^{\pi z} - e^{\pi z_j})^2} + \sum_{j=1}^{n} \frac{\beta_j}{e^{\pi z} - e^{\pi z_j}} \tag{5.64}$$

and the two relations between parameters

$$\beta_0 + \sum_{j=1}^{n} \beta_j = 0 \tag{5.65}$$

and

$$(1-\alpha_0^2)+\sum_{j=1}^{n}[2\beta_j e^{\pi z_j}+(1-\alpha_j^2)]=(1-\alpha_\infty^2), \tag{5.66}$$

where z and z_j refer to the strip domain and $w=f(z)$ is the mapping function from the strip to a circular arc polygon.

The first step in simplifying these formulas is to express $\{f,e^{\pi z}\}$ in terms of f and its z derivatives. The chain rule and a bit of manipulation yields the surprisingly simple relation

$$\{f,e^{\pi z}\}=e^{-2\pi z}\left[\frac{1}{\pi^2}\{f,z\}+\frac{1}{2}\right]. \tag{5.67}$$

Substituting this into (5.64) and multiplying by $e^{2\pi z}$ then gives the intermediate form

$$\frac{1}{\pi^2}\{f,z\}+\frac{1}{2}=\frac{1}{2}(1-\alpha_0^2)+\beta_0 e^{\pi z}+\frac{1}{2}\sum_{j=1}^{n}\frac{1-\alpha_j^2}{(1-e^{-\pi(z-z_j)})^2}+\sum_{j=1}^{n}\frac{\beta_j e^{\pi z}}{1-e^{-\pi(z-z_j)}}. \tag{5.68}$$

We can further simplify this by substituting $-\sum_{j=1}^{n}\beta_j$ for β_0 and then combining the terms containing β_j. The parameter β_0 and the constraint (5.65) are thus eliminated from the problem entirely, and the differential equation becomes

$$\frac{1}{\pi^2}\{f,z\}=-\frac{\alpha_0^2}{2}+\frac{1}{2}\sum_{j=1}^{n}\frac{1-\alpha_j^2}{(1-e^{-\pi(z-z_j)})^2}+\sum_{j=1}^{n}\frac{\beta_j e^{\pi z_j}}{1-e^{-\pi(z-z_j)}}. \tag{5.69}$$

Now consider the interpretation of the parameters β_j in terms of the geometry of the mapping, as discussed in Sections 5.2 and 5.4. The derivative in the definition of β_j is taken in the upper half-plane—see equations (5.5) and (5.7). When mapping to a highly elongated region these derivatives will be of widely varying orders of magnitude, making the β_j a poor set of parameters to use in the mapping function. It is much better to take the derivative in the strip domain, thus yielding a set of parameters $\beta_j''=e^{\pi z_j}\beta_j$ which should all be roughly of order 1.[3] This final change gives us

[3]If we carry out a detailed calculation similar to that leading up to (5.49), we obtain

$$\beta_j=e^{-\pi z_j}\left[\frac{1}{\pi}\beta_j^s-\frac{1}{2}(1-\alpha_j^2)\right],$$

where z_j is the variable in the strip and β_j^s is an accessory parameter properly scaled for the strip. This same expression results whether the singularities in the strip are modeled by using $(z-z_j)^{-(\alpha_j+1)}$ or by using a sinh function as in Chapter 2, since these are asymptotically the same near a singularity. This "rigorous" parameter β_j^s differs from β_j'' only by a small constant, so the latter is indeed of order 1.

Theorem 5.5 *The differential equation for the map from an infinite strip to a circular arc polygon, and the one remaining constraint, can be written as*

$$\frac{1}{\pi^2}\{f,z\} = -\frac{\alpha_0^2}{2} + \frac{1}{2}\sum_{j=1}^{n}\frac{1-\alpha_j^2}{(1-e^{-\pi(z-z_j)})^2} + \sum_{j=1}^{n}\frac{\beta_j''}{1-e^{-\pi(z-z_j)}}, \tag{5.70}$$

and

$$(1-\alpha_0^2) + \sum_{j=1}^{n}[2\beta_j'' + (1-\alpha_j^2)] = (1-\alpha_\infty^2). \tag{5.71}$$

The differential equation (5.70) has a number of properties which should make it well-suited for numerical computation. Not only are the constant parameters z_j, α_j, and β_j'' all of roughly the same order of magnitude, but except for the inevitable poles the individual terms in both sums all remain of manageable size throughout the length of the strip. (This is in contrast to the strip transformation (2.10), where the individual sinh functions vary exponentially along the strip and special care must be taken to avoid overflows when numerically computing the integrand.) Both the ODE (5.70) and the constraint (5.71) are also invariant under translations, i.e., the prevertices can be moved as a group any distance along the strip without changing any terms in either equation.

Far to the left of the group of prevertices, all of the terms in (5.70) except the α_0 term drop out, leaving $\{f,z\} \sim -\pi^2\alpha_0^2/2$. In the other direction the denominators all approach 1, so

$$\frac{1}{\pi^2}\{f,z\} \sim -\frac{1}{2}\alpha_0^2 + \frac{1}{2}\sum_{j=1}^{n}(1-\alpha_j^2) + \sum_{j=1}^{n}\beta_j''. \tag{5.72}$$

Using the relation (5.71), this simplifies to $\{f,z\} \sim -\pi^2\alpha_\infty^2/2$. The expected asymptotic behaviors for the mapping function at the ends of the channel are $f'(z) \sim e^{\pi\alpha_0 z}$ and $f'(z) \sim e^{-\pi\alpha_\infty z}$; by taking the Schwarzians of these expressions the limits from the differential equation are easily shown to be correct.

Conformal maps to infinite channels generated by (5.70) have one important complication not encountered with the strip transformation (2.10). The ends of a channel should be at infinity, corresponding to negative values for α_0 and α_∞. In equations (5.70) and (5.71), however, these parameters only appear as α_0^2 and α_∞^2. This means that positive and negative values for these, and indeed for any of the α parameters, are indistinguishable to the governing equations. (The same is also true for all of the other circular arc formulas we have discussed.) The reason is simple: In contrast to the formulas for straight-sided polygons, which are invariant only under linear transformations, the circular arc formulas are invariant under arbitrary Möbius transformations. The point at infinity can therefore not have any special properties to distinguish it from other points of the complex plane. An angle at infinity, with $\alpha \leq 0$, appears on the Riemann sphere as an ordinary angle measuring $|\alpha|\pi$. A Möbius transformation can move this angle to any other point of the plane. The placement of the divergent points at infinity must be explicitly required by the normalization and the solution of the parameter problem; it no longer follows automatically from the form of the mapping function.

In normalizing a circular arc channel, we can use a Möbius transformation to fix the positions of any three distinct points. For example, we could send the image of the left end of the strip to infinity and w_1 and w_n to their correct positions.[4] There is then one parameter required to specify the orientation of a side leaving w_1, $2n-4$ parameters to determine the curvature and length of each of the $n-2$ finite sides, and one additional parameter to specify the image of the right end of the strip. (In an infinite channel, this image should be the point at infinity. If the solution to the parameter problem is slightly in error, then the sides of the channel will appear to diverge at the proper angle for a great distance but ultimately will curve back to meet at some large but finite point.) In the mapping function (5.70) itself, as with the strip transformation, one prevertex can be fixed in place—say w_1 at the origin—but the rest must be allowed to vary. There are also the n accessory parameters β''_j, which must satisfy the single equality constraint (5.71). Thus there are $2n-2$ independent parameters appearing in the mapping function corresponding to the $2n-2$ degrees of freedom of a circular arc channel.

Considering the difficulties involved in solving the parameter problem for even the

[4]As we discussed in the previous section, this type of normalization might not always be appropriate for highly distorted versions of a polygon. For our purposes here we will ignore those complications, however.

original Schwarzian differential equation (5.1), we have not attempted to implement the strip map (5.70). If the complications described in the previous section could be resolved, however, there is no apparent reason why this formula would not be useful for practical computations.

5.7 Summary

To summarize, we have made the following contributions to the conformal mapping of circular arc polygons:

- We have developed formulas for numerically integrating the Schwarzian differential equation near the prevertex singularities. This permits the corner positions to be determined directly, and also makes it possible to integrate along the boundary using real arithmetic.

- By considering the behavior of the map from the unit disk, we have derived an alternative set of real accessory parameters with the following advantages:

  - They typically improve the conditioning of the parameter problem by an order of magnitude.
  - They yield less distorted initial guesses for the iterative solution algorithm.
  - They play a role in a new symmetrical form of the Schwarzian differential equation based directly on the map from the disk.

- In examining the problem of formulating a useful nonlinear system for solving the parameter problem, we have identified the following shortcomings in Bjørstad and Grosse's earlier implementation:

  - Their use of an overdetermined system is inefficient, and may also introduce suboptimal local minima which could cause an iterative solution algorithm to fail.

  - Their characterization does not uniquely define the shape of the image polygon.
  - Their normalization is not always well-defined, and may yield polygons which can not be continuously deformed to reach the desired solution.

  We have discussed various approaches to overcoming these difficulties, but with only partial success.

- Finally, motivated by problems with a high degree of crowding, we have derived a differential equation for the map from an infinite strip to a circular arc polygon. Another new set of accessory parameters helps insure that the terms in the equation will all be of roughly the same order of magnitude, an advantage for numerical calculations.

Appendix A

Elliptic Functions

The use of the rectangle mapping described in Chapter 2 requires the efficient evaluation of the function $s(z) = \frac{1}{\pi} \log \mathrm{sn}(z|m)$ over a wide range of values of z and m. All of the formulas we use for these computations are well known and available in standard references, but since there are several different asymptotic regimes involved it seems appropriate to give a brief summary of our methods here. For more information about any of the following material see [1].

The function $\mathrm{sn}(z|m)$ maps the rectangle with corners K, $-K$, $K+iK'$, and $-K+iK'$ to the upper half plane with the images of the corners being ± 1 and $\pm m^{-1/2}$. The composite function $s(z)$ thus sends the corners to the points 0, i, L, and $L+i$ on an infinite horizontal strip with height i. L is determined by the modified Schwarz-Christoffel transformation, so m is immediately known as $e^{-2\pi L}$.

The parameter m decreases exponentially with $K'/2K$, the conformal modulus of the rectangle. For highly elongated regions, therefore, m may be less than the underflow limit of many computers. We will first deal with the case where m is reasonably large; in our own program this means $m > 10^{-30}$.

K and K' are complete elliptic integrals with parameter m, and can be readily found by means of the arithmetic-geometric mean (AGM) method. To calculate K we first take

$$a_0 = 1,\ b_0 = \sqrt{m_1},\ c_0 = \sqrt{m}; \tag{A.1}$$

$m_1 = 1 - m$ is called the complementary parameter. The AGM iteration, defined by the formulas

$$a_i := \frac{1}{2}(a_{i-1} + b_{i-1}), \quad b_i := (a_{i-1}b_{i-1})^{1/2}, \quad c_i := \frac{1}{2}(a_{i-1} - b_{i-1}), \tag{A.2}$$

is then carried out until at the Nth step c_N is negligible to the required accuracy. K is then equal to $\pi/2a_N$; to find K' the same procedure is followed with m and m_1 interchanged.

We need to evaluate $\mathrm{sn}(z|m)$ at points inside the rectangle with corners $-K$, K, $K + iK'$, $-K + iK'$. This function has a pole at iK', but we avoid any difficulties there by using the identity

$$\mathrm{sn}(z|m) = \frac{-1}{m^{1/2}\mathrm{sn}(iK' - z|m)} \tag{A.3}$$

whenever $\mathrm{Im}(z) > K'/2$. For $\mathrm{Im}(z) \leq K'/2$ we can approximate $\mathrm{sn}(z|m)$ by

$$\mathrm{sn}(z|m) \sim \sin(z) - \frac{1}{4}m[z - \sin(z)\cos(z)]\cos(z) \tag{A.4}$$

when m is small enough; the relative accuracy of this formula is $\mathrm{O}(m)$ for $\mathrm{Im}(z) \leq K'/2$. If m is not sufficiently small we can reduce it by applying the descending Landen transformation as many times as necessary:

$$\mu = \left(\frac{1 - m_1^{1/2}}{1 + m_1^{1/2}}\right)^2, \tag{A.5}$$

$$v = \frac{z}{1 + \mu^{1/2}}, \tag{A.6}$$

$$\mathrm{sn}(z|m) = \frac{(1 + \mu^{1/2})\mathrm{sn}(v|\mu)}{1 + \mu^{1/2}\mathrm{sn}^2(v|\mu)}. \tag{A.7}$$

The effect is to replace m and z by μ (not to be confused with the conformal modulus) and v, where $\mu \sim m^2/16$ and $v \sim z$. To avoid cancellation errors, (A.5) should be evaluated via a power series when m is small:

$$\mu^{1/2} = \left(\frac{m}{4}\right) + 3\left(\frac{m}{4}\right)^2 + 17\left(\frac{m}{4}\right)^3 + 155\left(\frac{m}{4}\right)^4 + 2025\left(\frac{m}{4}\right)^5 + 34827\left(\frac{m}{4}\right)^6 + \ldots \tag{A.8}$$

In our own program we use (A.5) when $m > 10^{-3}$ and (A.8) otherwise.

The descending Landen transformation is by no means the only approximation we could use for $\mathrm{sn}(z|m)$. Abramowitz and Stegun, for instance, recommend using the AGM scale defined above in (A.1) and (A.2), along with the relations

$$\phi_N = 2^N a_N z, \tag{A.9}$$

$$\sin(2\phi_{i-1} - \phi_i) = \frac{c_i}{a_i}\sin(\phi_i), \tag{A.10}$$

$$\mathrm{sn}(z|m) = \sin(\phi_0). \tag{A.11}$$

This works well when z is near the real axis. Since the sine of a real number never has absolute magnitude greater than 1, the right hand side of (A.10) is zero to the required accuracy when $i = N$. When z is far from the real axis, however, the right hand side may not be negligible until i is considerably larger than N, so (A.10) may not be trustworthy for highly elongated rectangles.

In order to restrict the use of the AGM to real arguments, we considered using the identity

$$\mathrm{sn}(x+iy|m) = \frac{\mathrm{sn}(x|m)\mathrm{dn}(y|m_1) + i\mathrm{cn}(x|m)\mathrm{dn}(x|m)\mathrm{sn}(y|m_1)\mathrm{cn}(y|m_1)}{\mathrm{cn}^2(y|m_1) + m\,\mathrm{sn}^2(x|m)\mathrm{sn}^2(y|m_1)}, \tag{A.12}$$

where the elliptic functions cn and dn can be found by using the additional formulas

$$\mathrm{cn}(z|m) = \cos(\phi_0) \ \text{ and } \ \mathrm{dn}(z|m) = \frac{\cos(\phi_0)}{\cos(\phi_1 - \phi_0)}. \tag{A.13}$$

Some of the six functions appearing in (A.12) have their own conditioning problems when m is small, though, so the AGM method does not appear to be a promising alternative to the Landen transformation. None of the other methods we have considered, such as the representation of $\mathrm{sn}(z|m)$ in terms of theta functions, seem to offer any significant advantages either.

When $L > 11$, i.e. $m < 10^{-30}$, we use a different set of asymptotic formulas to avoid possible underflow of m or overflow of $\mathrm{sn}(z|m)$. K and K' are calculated via the approximations

$$K \sim \frac{\pi}{2} \ \text{ and } \ K' \sim \pi L + \log 4, \tag{A.14}$$

which are accurate to $\mathrm{O}(m)$. For $\mathrm{Im}(z) \leq K'/2$ the approximation $\mathrm{sn}(z|m) \sim \sin(z)$ has a relative accuracy of $\mathrm{O}(m^{1/2})$, and $\log(\sin z)$ can be expanded to give

$$s(z) \sim \frac{1}{\pi}\left\{-iz + \log\left[\frac{1}{2i}(e^{2iz} - 1)\right]\right\}. \tag{A.15}$$

When $\mathrm{Im}(z) > K'/2$ the identity (A.3) leads to the similar formula

$$s(z) \sim L + i + \frac{1}{\pi}\left\{iu - \log\left[\frac{1}{2i}(e^{2iu} - 1)\right]\right\}, \tag{A.16}$$

where $u = iK' - z$.

By the methods described here we can evaluate $s(z) = \frac{1}{\pi}\log\mathrm{sn}(z|m)$ to close to full precision (around 15 decimal places in our calculations), for all z in the fundamental rectangle, for a range of parameters roughly $e^{-2\pi/\epsilon} \ll m < 1/2$, that is, $\mathrm{O}(1) \leq L \ll \epsilon^{-1}$, where ϵ is the machine precision.

Bibliography

[1] M. Abramowitz and I. Stegun, eds., *Handbook of Mathematical Functions,* National Bureau of Standards, 1964.

[2] L. V. Ahlfors, *Complex Analysis,* McGraw-Hill, New York, 1979.

[3] O. L. Anderson, R. T. Davis, G. B. Hankins, and D. E. Edwards, *Solution of viscous internal flows on curvilinear grids generated by the Schwarz-Christoffel transformation,* in Numerical Grid Generation, J. F. Thompson, ed., North-Holland, Amsterdam, 1982.

[4] K. J. Binns, *Calculation of some basic flux quantities in induction and other double-slotted electrical machines,* Proc. I.E.E., 111 (1964), pp. 1847–1858.

[5] K. J. Binns, G. R. Rees and P. Kahan, *The evaluation of improper integrals encountered in the use of conformal transformation,* Int. J. Num. Meth. Eng. 14 (1979), pp. 567–580.

[6] P. Bjørstad and E. Grosse, *Conformal mapping of circular arc polygons,* SIAM J. Sci. Stat. Comput. 8 (1987), pp. 19–32.

[7] H. D. Däppen, *Wind-tunnel wall corrections on a two-dimensional plate by conformal mapping,* AIAA J. 25 (1987), pp. 1527–1530.

[8] H. D. Däppen, *Die Schwarz-Christoffel-Abbildung für zweifach zusammenhängende Gebiete mit Anwendungen,* Ph.D. diss., ETH, Zurich, 1988.

[9] R. T. Davis, *Numerical methods for coordinate generation based on Schwarz-Christoffel transformations,* Proc. 4th AIAA Comput. Fluid Dynamics Conf., Williamsburg, VA, 1979.

[10] T. K. DeLillo, *On some relations among numerical conformal mapping methods,* J. Comput. Appl. Math. 19 (1987), pp. 363–377.

[11] T. K. DeLillo, *A note on Rengel's inequality and the crowding phenomenon in conformal mapping,* Appl. Math. Lett., to appear.

[12] J. E. Dennis, Jr., D. M. Gay, and R. E. Welsch, *An adaptive nonlinear least-squares algorithm,* ACM Trans. Math. Softw. 7 (1981), pp. 348–368.

[13] J. E. Dennis, Jr., D. M. Gay, and R. E. Welsch, *Algorithm 573: NL2SOL—An adaptive nonlinear least-squares algorithm,* ACM Trans. Math. Softw. 7 (1981), pp. 369–383.

[14] J. E. Dennis, Jr. and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations,* Prentice-Hall, Englewood Cliffs, NJ, 1983.

[15] F. Dias, *On the Use of the Schwarz-Christoffel Transformation for the Numerical Solution of Potential Flow Problems,* Ph.D. diss., U. Wisconsin, Madison, 1986.

[16] F. Dias, A. R. Elcrat and L. N. Trefethen, *Ideal jet flow in two dimensions,* J. Fluid Mech. 185 (1987), pp. 275–288.

[17] J. J. Dongarra and E. Grosse, *Distribution of mathematical software via electronic mail,* Commun. ACM 30 (1987), pp. 403–407.

[18] M. Dubiner, *Theoretical and Numerical Analysis of Conformal Mapping,* Ph.D. diss., Dept. of Mathematics, MIT, 1981.

[19] A. R. Elcrat and L. N. Trefethen, *Classical free-streamline flow over a polygonal obstacle,* J. Comput. Appl. Math. 14 (1986), pp. 251–265.

[20] W. Fleming, *Functions of Several Variables,* Springer-Verlag, New York, 1977.

[21] J. M. Floryan, *Conformal-mapping-based coordinate generation method for channel flows,* J. Comput. Phys. 58 (1985), pp. 229–245.

[22] J. M. Floryan, *Conformal-mapping-based coordinate generation method for flows in periodic configurations,* J. Comput. Phys. 62 (1986), pp. 221–247.

[23] J. M. Floryan and C. Zemach, *Schwarz-Christoffel mappings: a general approach,* J. Comput. Phys. 72 (1987), pp. 347–371.

[24] J. M. Floryan and C. Zemach, *Quadrature Rules for Singular Integrals with Application to Schwarz-Christoffel Mappings,* J. Comput. Phys. 75 (1988), pp. 15–30.

[25] G. E. Forsythe, M. A. Malcolm and C. B. Moler, *Computer Methods for Mathematical Computations,* Prentice-Hall, Englewood Cliffs, NJ, 1977.

[26] D. Gaier, *Konstruktive Methoden der konformen Abbildung,* Springer-Verlag, Berlin, 1964.

[27] D. Gaier, *Ermittlung des konformen Moduls von Vierecken mit Differenzenmethoden,* Numer. Math. 19 (1972), pp. 179–194.

[28] D. Gaier, *Capacitance and the conformal module of quadrilaterals,* J. Math. Anal. Appl. 70 (1979), pp. 236–239.

[29] A. F. Ghoniem and Y. Gagnon, *Vortex simulation of laminar recirculating flow,* J. Comput. Phys. 68 (1987), pp. 346–377.

[30] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization,* Academic Press, Orlando, FL, 1981.

[31] G. H. Golub and J. H. Welsch, *Calculation of Gaussian quadrature rules,* Math. Comp. 23 (1969), pp. 221–230.

[32] L. Greengard, *Potential flow in channels,* subm. to SIAM J. Sci. Stat. Comput.

[33] M. H. Gutknecht, *Numerical conformal mapping methods based on function conjugation,* in [65], pp. 31–78.

[34] P. Henrici, *Applied and Computational Complex Analysis, Vol. 1,* Wiley, New York, 1974.

[35] P. Henrici, *Applied and Computational Complex Analysis, Vol. 3,* Wiley, New York, 1986.

[36] E. Hille, *Ordinary Differential Equations in the Complex Domain,* Wiley, New York, 1976.

[37] M. Hoekstra, *Coordinate generation in symmetrical interior, exterior or annular 2-D domains, using a generalized Schwarz-Christoffel transformation,* in Numerical Grid Generation in Computational Fluid Mechanics, J. Hauser and C. Taylor, eds., Pineridge Press, Swansea, Wales, 1986.

[38] L. H. Howell and L. N. Trefethen, *A modified Schwarz-Christoffel transformation for elongated regions,* SIAM J. Sci. Stat. Comput., to appear.

[39] D. C. Ives, *Conformal grid generation,* in Numerical Grid Generation, J. F. Thompson, ed., North-Holland, Amsterdam, 1982.

[40] L. V. Kantorovich and V. I. Krylov, *Approximate Methods of Higher Analysis,* P. Noordhoff, Groningen, Netherlands, 1958.

[41] Ç. K. Koç, *Schwarz-Christoffel Transformation for the Simulation of Two-Dimensional Capacitance,* subm. to IEEE Trans. CAD/ICAS.

[42] R. Menikoff and C. Zemach, *Methods for numerical conformal mapping,* J. Comput. Phys. 36 (1980), pp. 366–410.

[43] E.-S. Meyer, *Praktische Verfahren zur konformen Abbildung von Geradenpolygonen,* Ph.D. diss., U. Hannover, 1979.

[44] J. J. Moré, B. S. Garbow and K. E. Hillstrom, *User Guide for MINPACK-1,* Argonne Nat. Lab. Rep. ANL-80-74, 1980.

[45] Z. Nehari, *The Schwarzian derivative and schlicht functions,* Bull. Amer. Math. Soc. 55 (1949), pp. 545–551.

[46] Z. Nehari, *Conformal Mapping,* Dover, New York, 1952.

[47] J. M. Ortega and W. C. Rheinboldt, *Iterative Solution of Nonlinear Equations in Several Variables,* Academic Press, Orlando, FL, 1970.

[48] N. Papamichael, C. A. Kokkinos and M. K. Warby, *Numerical techniques for conformal mapping onto a rectangle,* J. Comput. Appl. Math. 20 (1987), pp. 349–358.

[49] N. Papamichael, *Numerical conformal mapping onto a rectangle with applications to the solution of Laplacian problems,* J. Comput. Appl. Math. 28 (1989), pp. 63–83.

[50] N. Papamichael and N. S. Stylianopoulos, *A domain decomposition method for conformal mapping onto a rectangle,* Tech. Rep. TR/04/89, Dept. of Mathematics and Statistics, Brunel U., 1989.

[51] N. Papamichael and N. S. Stylianopoulos, *On the numerical performance of a domain decomposition method for conformal mapping,* Tech. Rep. TR/05/89, Dept. of Mathematics and Statistics, Brunel U., 1989.

[52] R. Piessens, E. de Doncker-Kapenga, C. W. Überhuber and D. K. Kahaner, *QUADPACK: A Subroutine Package for Automatic Integration,* Springer-Verlag, Berlin, 1983.

[53] M. J. D. Powell, *A Fortran subroutine for solving systems of non-linear algebraic equations,* Tech. Rep. AERE-R. 5947, Harwell, England, 1968.

[54] W. Press, B. Flannery, S. Teukolsky and W. Vetterling, *Numerical Recipes,* Cambridge U. Press, Cambridge, England, 1986.

[55] W. Prochazka, *Conformal mapping of the unit circle or of the upper half plane onto a polygon,* Computing 31 (1983), pp. 155–172.

[56] K. Reppe, *Berechnung von Magnetfeldern mit Hilfe der konformen Abbildung durch numerische Integration der Abbildungsfunktion von Schwarz-Christoffel,* Siemens Forsch. u. Entwickl. Ber. 8 (1979), pp. 190–195.

[57] H. A. Schwarz, *Über einige Abbildungsaufgaben,* J. reine u. angew. Math. 70 (1869), pp. 105–120.

[58] L. F. Shampine and M. K. Gordon, *Computer Solution of Ordinary Differential Equations: the Initial Value Problem,* W. H. Freeman, San Francisco, 1975.

[59] K. P. Sridhar, *Coordinate System Generation for Internal Flow Problems via Schwarz-Christoffel Transformations,* Ph.D. diss., West Virginia U., 1981.

[60] K. P. Sridhar and R. T. Davis, *A Schwarz-Christoffel method for generating internal flow grids,* Computers in Flow Predictions and Fluid Dynamics Experiments, Proc. Winter Annual Meeting of ASME, Washington, D.C., 1981.

[61] K. P. Sridhar and R. T. Davis, *A Schwarz-Christoffel method for generating two-dimensional flow grids,* J. Fluids Eng. 107 (1985), pp. 330–337.

[62] L. N. Trefethen, *Numerical computation of the Schwarz-Christoffel transformation,* SIAM J. Sci. Stat. Comput. 1 (1980), pp. 82–102.

[63] L. N. Trefethen, *SCPACK User's Guide,* Numerical Analysis Rep. 89-2, Dept. of Mathematics, MIT, 1989.

[64] L. N. Trefethen, *Analysis and design of polygonal resistors by conformal mapping,* J. Appl. Math. Phys. (ZAMP) 35 (1984), pp. 692–704.

[65] L. N. Trefethen, ed., *Numerical Conformal Mapping,* North-Holland, Amsterdam, 1986, also appeared as J. Comput. Appl. Math. 14 (1986), pp. 1–270.

[66] L. N. Trefethen and R. J. Williams, *Conformal mapping solution of Laplace's equation on a polygon with oblique derivative boundary conditions,* in [65], pp. 227–249.

[67] L. N. Trefethen, *Schwarz-Christoffel Mapping in the 1980's,* Numerical Analysis Rep. 89-1, Dept. of Mathematics, MIT, 1989.

[68] L. C. Woods, *The Theory of Subsonic Plane Flow,* Cambridge U. Press, Cambridge, England, 1961.

[69] C. Zemach, *A conformal map formula for difficult cases,* in [65], pp. 207–216.