Renormalization for Siegel Discs

by

Andrew D Burbanks

A Doctoral Thesis

Submitted in partial fulfilment of the requirements
for the award of Doctor of Philosophy
of Loughborough University

1997

(Revision 1.1, 1998)

Abstract

This thesis is concerned with a domain of linearizability, otherwise known as a *Siegel disc*, around an irrationally indifferent fixed point of a complex analytic map. In particular, we investigate the existence of Siegel discs and examine the properties of their boundary curves for golden mean rotation number. The key tool used is the idea of a *renormalization operator* acting on a space of functions.

Firstly, a computer-assisted proof is discussed and verified, which establishes the existence of a fixed point of the relevant renormalization operator. In particular, the proof yields a ball of functions around an approximate fixed point that is guaranteed to contain the true fixed point. The rigorous computational techniques which allow computers to be used for this purpose are then discussed.

Given the existence of the renormalization fixed point, we verify certain topological conditions, known as the *necklace hypotheses*, on the action of the maps making up the fixed point. This proves the existence of a Siegel disc having a Hölder continuous (invariant) boundary curve for all maps attracted to the fixed point. Further, it is shown that the motion on the boundary is conjugate to a pure rotation, that the boundary curve passes through a critical point of the map, and that the conjugator is not differentiable on a dense set of points.

Finally, by viewing the invariant curve as the limit set of an *iterated function system* (IFS), a further investigation is made to get rigorous bounds on the fractal dimension of the Siegel disc boundary. This involves calculating bounds on the *contractivities* and *coercivities* of the maps of the IFS and solving corresponding *partition equations*. In particular, a rigorous upper bound on the dimension of 1.08523 is obtained.

Acknowledgements

To my parents and family, who have always given their unconditional support during my studies. I would like to thank my supervisor, Andy Osbaldestin, for putting up with my frequent (and sometimes lengthy) interruptions, for being willing to offer advice and help at a moment's notice and, most importantly, for his friendship.

Special thanks must also go to Andreas Stirnemann, in particular for his extensive contributions to the theory presented in Chapter 5, for his willingness to put up with a bombardment of questions by email, and for instilling in me an enthusiasm for object-oriented programming. Thanks also for the friendship and hospitality extended to me by Andreas and his wife Lynda, which made my visits to Exeter and Edinburgh thoroughly enjoyable.

To my friends, including the original crew of office W164: Keith Beardmore, Ros Sinnamon, and Jason Henry. To Cristina, Rachel, Jon, Ben, Steve, and the other maths department postgraduates. To Natalie Duff for being a great friend and drinking partner. To Peter, Paul, and the lunch-timers. To the the youth hostellers, the Babblings crew, the East-midlands pub crew (especially Beverley, Frankie, Chris, and Jon), and my friends back home. (And apologies to anyone else I forgot!)

Andy Burbanks,
Loughborough, April 1997.

Grants: During the first year (1993) this work was funded jointly by the Department of Mathematical Sciences and by a Loughborough University Studentship. In the second and third years funding was by an EPSRC Studentship (94001339).

Technical Preparation: Thesis typesetting with LaTeX2ε, using the Euler math fonts designed by Herman Zapf. Pictures produced using a `C++/Postscript` interface written for this project by the author, and included in LaTeX2ε using the `epsfig` package. Main programs written using GNU-`C++` and the literate programming tool `Noweb` (by Norman Ramsey), within the padded cell provided by the GNU `emacs` editor. The source code for the rigorous programs is available from the author and will also be submitted to a software archive. Thanks to Keith Watling for computer support.

Contents

Chapter 1

Introduction

1.1 Overview

This thesis investigates the existence and properties of *Siegel discs* for iterated complex maps. A Siegel disc is a neighbourhood of a fixed point of the map on which iteration is conjugate to a rigid rotation of the plane. It turns out that the existence of such a region is governed by the number-theoretic properties of the *multiplier* of the map. This thesis details rigorous computer proofs of a number of conjectures concerning Siegel discs and their boundary curves (outside which the conjugacy to pure rotation breaks down).

The question of the existence of Siegel discs lies in a class of related problems that are collectively called *KAM theory*. KAM theory plays a crucial role in explaining the transition from orderly to chaotic behaviour in models of physical systems described by differential equations. For certain systems, the motion takes place on invariant surfaces (in fact, tori) in the phase space. KAM theory is concerned with the conditions under which these invariant surfaces will either persist (albeit in a distorted form) or break-up if the system is perturbed.

In addition to their practical importance, such studies have enjoyed a surge of interest recently, as certain characteristics of the transition to chaotic behaviour appear to be *universal* — that is, they are present across a range of systems that share just a few common features. Universality is important as it allows both qualitative and *quantitative* predictions to be made about complicated dynamical systems by looking at simpler systems that are in the same *universality class*. Indeed, many of the general predictions made by analysing grossly simplified models have been confirmed to a comparatively high degree of accuracy in a variety of physical experiments on real systems.

A key tool in understanding universality is the technique of *renormalization*. The main part of this introduction begins by discussing renormalization theory and by giving specific examples of its use. One of the simplest (and most striking) examples of universal behaviour occurs in the cascade of *period-doublings* observed in one-dimensional maps. The renormalization explanation offered by Feigenbaum [Fei79] is described, along with its verification. This will serve not only to describe the renormalization technique, but also to introduce the notion of a rigorous *computer-assisted proof* for the existence of a

fixed point of a renormalization operator. The latter involves high precision arithmetic and careful maintenance of error bounds to give a rigorous proof. This approach forms the basis for the main results of this thesis.

The second example is concerned with the breakdown of quasiperiodicity for diffeomorphisms of the circle. This is motivated in part by the analysis of the dynamics seen in dissipative systems. This example serves to introduce the concepts of *rotation number*, *conjugacy to a rigid rotation*, and the importance of continued fraction expansions. In particular, a renormalization operator will be introduced that is used to examine the breakdown of quasiperiodicity in the case of the *golden mean* rotation number (a variant of this operator will be used in the rest of this thesis). These ideas will prove useful when discussing the Siegel disc problem.

The third example concerns the persistence of invariant circles in area-preserving twist maps of the plane. This final example will mention another important notion, namely the *necklace construction*, which is used to deduce the existence of an invariant curve for maps attracted to a suitable renormalization fixed point (this will be used extensively in Chapter 4 and Chapter 5).

(The reader familiar with the above concepts may safely skip these preliminary sections.)

After these preliminaries, we discuss the history and details of the Siegel disc problem itself, starting with a brief overview of the relevant topics in complex iteration theory.

1.1.1 Organization of Thesis

This thesis is divided into five main chapters, which both present new results and develop necessary background ideas. The technical details are summarised briefly below.

- After a preliminary discussion concerning renormalization and universality, the remainder of this chapter introduces Siegel discs and the universality that is observed for them. In particular, it details the conjectures of Manton and Nauenberg [MN83] concerning the boundaries of Siegel discs, and the renormalization explanation offered by Widom [Wid83]. This explanation concerns the existence of a so-called *critical fixed point* of a renormalization operator.

- Chapter 2 details the verification of an earlier proof of the existence of the critical fixed point. This enables the computer framework used to be adapted and yields a ball containing the fixed point (in a suitable space of pairs of functions) for use later in the thesis.

- Chapter 3 details the rigorous computational framework used both for the existence proof, and for the results obtained in the next two chapters. The technique involves maintaining rigorous upper and lower bounds on quantities appearing in a conceptual proof, and using only objects that the computer can deal with exactly (namely, the *representable numbers* and finitized versions of mathematical objects, e.g. functions, constructed from them).

- Chapter 4 gives a proof of the conjectures mentioned earlier, concerning the boundary of Siegel discs, in the case of golden mean rotation number. In particular, it is established that functions attracted to the fixed point have a Schröder series that converges on its circle of convergence, the image of the circle of convergence being an invariant curve that is conjugate to a rigid rotation. Further, the conjugation function is shown to be Hölder continuous, though not differentiable on a dense set of points. It is also proved that the boundary curve passes through a critical point of the map. The key tool used in the proof is the *necklace construction*, a technique which actually *constructs* a parameterization for the invariant curve.

- Finally, Chapter 5 develops some theory for obtaining upper and lower bounds on the Hausdorff dimension of the boundary curves for Siegel discs. The dimension turns out to be a renormalization invariant, which means that results obtained for the critical universal curve apply to the boundaries of maps that are attracted by the renormalization fixed point. After giving some preliminary analytical bounds, the necklace construction is modified to give a formulation of the invariant curve as the invariant set for an Iterated Function System (IFS). This enables upper and lower bounds to be obtained for the dimension by looking at the contractivity and coercivity of the maps of the IFS themselves. By verifying the *open set condition* for the lower bounds, this sets the stage for using rigorous numerics to find sharper bounds on the dimension. In particular, a rigorous upper bound of 1.08523 is found for the Hausdorff dimension.

1.2 Renormalization and Universality

In this section, we introduce the notion of *renormalization* by giving some examples of its use in recent studies of the transition to chaos in dynamical systems. The technique was motivated by the study of critical phenomena in condensed matter physics. The key insight is that certain phenomena are observed to occur repeatedly at many different scales. The technique of renormalization involves making systematic changes of scale to the system under study, in such a way that the phenomena of interest are preserved, while other particulars of the original system are lost. In essence, "complicated" systems are repeatedly replaced with "simpler" systems that share some feature of interest. The desired result is a system which has the same interesting features as the original, but now present at *all* scales, i.e. a system which exhibits an exact "self-similarity" that was present only approximately in the original.

More formally, the process involves iterating a *renormalization operator*, subject to *normalization conditions* which keep the chosen features constant. By analysing the action of the operator on a suitable space of functions (and, in particular, determining the nature of any fixed points), *quantitative* predictions may typically be made about whole sets of systems that have only some rudimentary *qualitative* features in common. Such systems are then said to lie in the same "universality class". As we shall discuss below, this kind of analysis has allowed numerical quantities (called "universal constants") to be derived from entirely abstract theoretical considerations of very simple models. These quantities

have then been observed in many different, more complicated, physical systems. (The explanation being that there happens to be a common universality class to which the complicated physical systems and the simple models both belong.)

It is perhaps easiest to demonstrate the technique by looking at one of the most well-known and straightforward applications, namely the study of the period doubling cascade in one-dimensional maps.

1.2.1 Period Doubling Universality

Some systems that undergo a transition from stable periodic to chaotic behaviour (as a parameter is varied) do so by means of a *period doubling cascade*. The prototype example of this sort of behaviour occurs in the iteration of one-parameter families of one-dimensional analytic maps of the interval having a single differentiable maximum of order $z > 1$ at a point $\hat{x}$, i.e.

$$f(\hat{x}) - f(x) \sim |x - \hat{x}|^z.$$

The simplest of these families of maps is the so-called *logistic* family (for which $z = 2$):

$$f_\lambda(x) = \lambda x(1 - x).$$

In general there is a range of parameter values for which the behaviour is attraction to a periodic orbit. As the parameter (in the above case, λ) is varied, the orbit may become unstable. For a certain class of maps, namely those having negative *Schwarzian derivative*,

$$S(f) = \frac{f'''}{f'} - \frac{3}{2}\left(\frac{f''}{f'}\right)^2,$$

there will be a *period doubling bifurcation*: as the periodic orbit becomes unstable it gives rise to a stable periodic orbit with twice the period (for example, the logistic family has an attractive fixed point for $\lambda < 3$ which bifurcates into a stable two-cycle as λ increases through $\lambda = 3$). The property of negative Schwarzian derivative is inherited by the iterates of the map, with the result that if the new stable cycle becomes unstable it will give rise to a stable cycle with four times the original period. In general, an infinite cascade of such "period doublings" will be observed. The size of the change in the parameter required to produce successive bifurcations is observed to reduce at a geometric rate, with the parameter values accumulating at some critical value, at which the dynamics of the system is no longer periodic (in fact, it is chaotic on a fractal limit set).

The surprising discovery, made by Feigenbaum [Fei78] and others, is that there are certain features of the period doubling cascade that are independent of the particular map under consideration. In particular, the period doubling sequences are asymptotically self-similar for a large class of one parameter families of functions. In other words, certain features of this transition to chaos are *universal*. More formally, if we denote the parameter of such a family by λ, and the value of this parameter at which the n-th period doubling bifurcation occurs by λ_n, then the sequence $\{\lambda_n\}$ is observed to converge asymptotically

geometrically to a value λ_∞ (at which periodicity is lost) at a rate that always tends to a *universal constant*,

$$\delta = \lim_{n\to\infty} \frac{\lambda_{n+1} - \lambda_n}{\lambda_{n+2} - \lambda_{n+1}}.$$

In the case of maps having a quadratic maximum (i.e. $z = 2$), for example, it turns out that

$$\delta = 4.669\ 201\ 609\ 102\ 990\ 671\ \ldots .$$

Universal scaling is also observed in the iteration space. This becomes apparent if we plot a graph of the attractor of a given map against the parameter value (sometimes called a "bifurcation diagram"). A diagram for the map $f_\lambda : x \mapsto \lambda x(1-x)$ is shown in figure 1.1(a) for the range $1 \le \lambda \le 4$ and figure 1.1(b) for $2 \le \lambda \le \lambda_\infty = 3.569\ 945\ldots$, in which the period doubling cascade is observed.

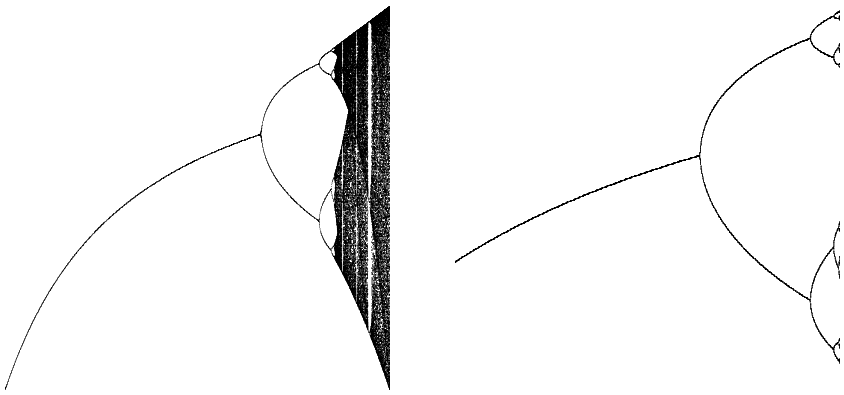

(a) Plot of attractor vs. parameter. (b) Period-doubling cascade.

Figure 1.1: Period doubling cascade for the map $f_\lambda : x \mapsto \lambda x(1-x)$.

The relative scale of successive branch splittings (that is, the vertical separation between branches at bifurcation points) is also universal: if ε_n is the separation between the two branches that resulted from the n-th bifurcation, measured at the $(n+1)$-th bifurcation, then we find that the ratio $\varepsilon_n/\varepsilon_{n+1}$ tends to another universal constant, given by

$$\alpha = \lim_{n\to\infty} \frac{\varepsilon_n}{\varepsilon_{n+1}}.$$

For $z = 2$ we have

$$\alpha = 2.502\ 907\ 875\ 095\ 892\ 822\ \ldots .$$

(Numerical values for these scaling constants, for various values of z, have been computed to very high accuracy [Bri91]. Analogous calculations have also been performed for the more general period "n-tuplings" observed in iterated complex maps [CM83] and [CM89].)

The presence of the universal scaling constants α and δ is responsible for the approximate self-similarity observed in the bifurcation diagram: asymptotically, a box around one of

the n-th bifurcation points is just a rescaled copy of a box around one of the $(n+1)$-th bifurcation points, with the horizontal and vertical scale factors mapping one box to the other converging to δ and α respectively. Another way of looking at this is to examine the parameter values s_n that give rise to a *super-stable* period 2^n orbit of f (in the case of the logistic family, an orbit that contains the critical point $x = 1/2$). The key observation is that the graph of $f_{s_n}^{2^n}$ resembles a suitably rescaled portion of the graph of $f_{s_{n+1}}^{2^{n+1}}$ around the critical point. For increasing n the portions of the graphs appear to converge, with the scale factor of the mapping from one to the other converging to α. This suggests that there may be a specific sequence of functions,

$$(-\alpha)^n f^{2^n}((-\alpha)^{-n}x),$$

having a universal limit. (The presence of $-\alpha$, rather than α, in the above indicates that the rescaling process actually "flips" the box around the critical point.) This suggests that the universality might be understood in renormalization terms. If the process of magnifying the neighbourhoods of fixed points is repeated many times, then practically all the information about the global shape of the original function tends to be lost. The result is a *universal function* for which the self-similarity of the bifurcation diagram is *exact*, i.e. the exact values of α and δ are observed in all bifurcations, rather than asymptotically in the limit (indeed, the shape of the function itself will be self-similar, in such a way that suitably rescaled portions of its iterates are exact copies of the original function).

The above observations suggest that we examine the renormalization operator $\mathcal{T}$, defined by

$$\mathcal{T}:\ f(z) \mapsto -\alpha f(f((-\alpha)^{-1}z)),$$

which represents the actions of iterating and rescaling by α. (This operator is sometimes called the *doubling transformation*.)

The Feigenbaum Conjectures

Based on such empirical findings, Feigenbaum ([Fei78] and [Fei79]) proposed some geometrical hypotheses about the action of the operator $\mathcal{T}$ on a suitable space of functions, and demonstrated how these account for the universality that is observed. In particular, he proposed that

- $\mathcal{T}$ has a fixed point, the universal function g, with $g'(0) = 0$ and $g''(0) \neq 0$. In other words, the universal function g and the scaling factor α correspond to a solution of the *functional equation*

  $$g(z) = -\alpha g(g((-\alpha)^{-1}z)).$$

  (Note that, if $x \mapsto g(x)$ is a solution to the above functional equation, then so is $x \mapsto a^{-1}g(ax)$ for any $a \neq 0$. In order to find a particular solution we need to set an absolute scale by imposing a *normalization condition*. For example, we might choose $g(0) = 1$, which means that $\alpha = -\frac{1}{g(1)}$.)

- At the fixed point, the derivative of $\mathcal{T}$ has a simple eigenvalue with modulus greater than one (in fact, it is δ); the rest of its spectrum is contained in the open unit disk. So the fixed point g has a one-dimensional unstable manifold W_U under $\mathcal{T}$, and has a codimension-one stable manifold W_S under $\mathcal{T}$.
- The unstable manifold W_U intersects transversally a codimension-one surface Σ_1 that contains all functions having a super-stable 2-cycle.

Figure 1.2 demonstrates these conjectures graphically. (Also shown in the figure is a typical one-parameter family f_λ that crosses the stable manifold W_S transversally.)

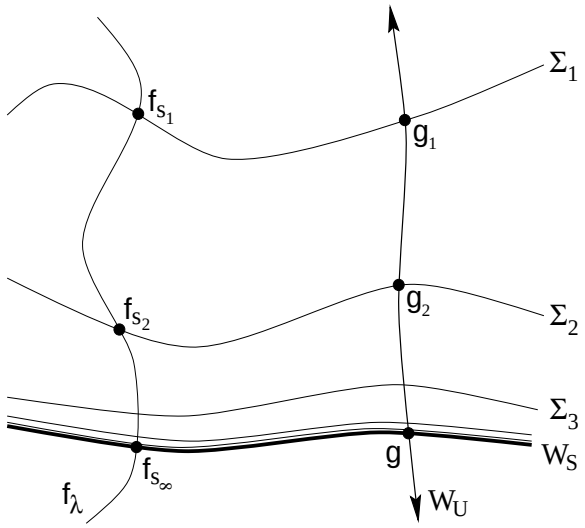

Figure 1.2: Schematic of the function space.

These conjectures account for the universality of δ. Consider the successive inverse images Σ_2, Σ_3, ... of the codimension-one surface Σ_1 of functions with super-stable 2-cycles: it may be shown that Σ_j is a codimension-one surface of functions with super-stable 2^j-cycles. The key idea is that the successive Σ_j's accumulate towards the stable manifold W_S with the separation between Σ_j and W_S decreasing asymptotically like δ^{-j}.

Suppose that we have a one-parameter family f_λ of functions that crosses the stable manifold W_S transversally at $\lambda = s_\infty$. Then, for large enough j, there should be a unique value of λ, say $\lambda = s_j$, such that f_{s_j} has a super-stable 2^j-cycle, and thus lies on Σ_j, so that

$$\lim_{j\to\infty} \delta^j (s_j - s_\infty)$$

exists and is non-zero.

In summary, there is a renormalization operator that involves a composition of maps, along with a rescaling process that enforces normalization conditions. The operator has a hyperbolic fixed point, the existence of which implies interesting behaviour for a whole universality class.

Lanford [Lan82] gave a computer-assisted proof that the universal function g exists, and that it is a hyperbolic fixed point in the space of functions. The proof used a computer

to perform rigorous calculations which showed that a variant of Newton's method for the fixed point equation (1.2.1) is a contraction on a small ball around an approximate fixed point. Further details of this method of proof will be given in Chapter 2, where it will be used for the corresponding proof in the case of Siegel discs for complex maps. (More recently, Epstein [Eps86] has given an analytic proof of the first conjecture, using the theory of Herglotz functions [EL81].)

The above theory has also been generalised to period "n-tuplings" for iterated complex maps [CM83],[CM89]. (In addition, iterated complex maps may be viewed as iterated two-dimensional mappings, which are important for modelling dynamical systems. In particular, period n-tuplings occur in Hamiltonian systems and are also characterised by universal scaling constants. However, whereas Hamiltonians are area-preserving, analytic maps are angle-preserving and it is questionable to what extent the characterisations are related.) These ideas also extend to period doubling in higher-dimensional systems [CEK81],[GMVF81]. In particular, for systems that are dissipative, volumes in phase space may contract onto an attractor. Typically, there will be one direction in which the contraction is slowest and, intuitively, one might expect that the dynamics on the attractor would be essentially one-dimensional.

1.2.2 Breakdown of Quasi-periodicity in Diffeomorphisms of the Circle

A second problem in which the ideas of universality have been applied is the transition from quasi-periodicity to chaos for diffeomorphisms of the circle. These kinds of maps arise in the study of many different physical systems. For example, for Hamiltonian systems, the KAM theorem establishes that trajectories in phase space are confined to travel on tori. A simplified system is obtained by ignoring the radial motion and concentrating on the angular motion, yielding a map of the circle.

A prototypical two-parameter family of circle maps is given by

$$f_{w,\kappa}(x) = x + w - \frac{\kappa}{2\pi}\sin 2\pi x, \tag{1.1}$$

where w may be thought of as a forcing frequency, and κ controls the magnitude of the nonlinearity present in the model. (At $\kappa = 0$ the the above equation reduces to a rigid rotation by w.) For $|\kappa| < 1$ the equation defines a diffeomorphism of the circle (in particular, the function is invertible). For $|\kappa| > 1$ the map is not invertible, the transition to chaos occuring as the map starts to lose invertibility at the *critical value* $|\kappa| = 1$. At $|\kappa| = 1$ the equation defines a homeomorphism; in particular, the inverse exists but is not differentiable at the origin, where there is a cube root singularity.

(Here, the circle T^1 is represented by the real numbers $\mathbb{R} \bmod 1$. Note that every homeomorphism f of T^1 can be represented by a periodic homeomorphism of $\mathbb{R}$, i.e. a homeomorphism f such that $f(x+1) = f(x)$ which is equivalent to $x \mapsto f(x) \bmod 1$. Equivalently, one would say that the map on the reals is the *lift* of the map on the circle.)

Definition 1 (Rotation Number)
The rotation number *(or* winding number*) of* f *is defined to be*

$$\rho(f) = \lim_{n\to\infty} \frac{1}{n}(f^n(x) - x). \tag{1.2}$$

(For homeomorphisms, this limit exists and is independent of x, [Arn83].) *Rational* rotation numbers correspond to periodic motion: $\rho(f) = p/q$ with $p \perp q$ (i.e. p, q are relatively prime integers) if and only if there exists some x_0 such that

$$f^q(x_0) = x_0 + p.$$

On the other hand, *irrational* rotation numbers correspond to quasi-periodic motion: for irrational values of ρ, f is topologically equivalent (conjugate) to a simple rotation by some value Ω, i.e. there exists a continuous invertible (conjugacy) function ϕ such that

$$f(\phi(x)) = \phi(R_\Omega(x)), \tag{1.3}$$

where $\phi(x+1) = \phi(x) + 1$, $\phi(0) = 0$, and $R_\Omega(x) = x + \Omega$ (i.e. R_Ω represents rotation by Ω). Equivalently, we have the following commutative diagram:

$$\begin{array}{ccc} x & \xrightarrow{R_\Omega} & x+\Omega \mod 1 \\ {\scriptstyle \phi^{-1}}\uparrow & & \downarrow{\scriptstyle \phi} \\ y & \xrightarrow{f} & f(y) \end{array} \tag{1.4}$$

Of particular interest for this problem is the question of what happens as the system is perturbed from rigid rotation by increasing κ from 0 to 1.

When $\kappa = 0$, irrational values of w give rise to quasiperiodicity (rotation by $\Omega = w$), whereas rational rotation numbers give rise to periodic orbits.

For $\kappa \leq 1$, the rotation number is monotonic in w, having the structure of a *devil's staircase*: it is constant on an infinite set of intervals corresponding to every rational rotation, and is irrational elsewhere. This phenomenon, in which rational periodic motion occurs for a finite range of forcing frequencies, is known as *mode-locking*. Figure 1.3(b) is a diagram of the parameter space (κ against w): as κ increases from 0 to 1, regions called *Arnold tongues* spread out from every rational number, within these regions the rotation number is rational and corresponds to the value of w where the tongue touches the axis $\kappa = 0$. More formally,

$$A_{p/q} = \{(\kappa, w) : \ f^q_{\kappa,w}(x) = x + p \text{ for some } x\}.$$

The figure shows the edges of the Arnold tongues $A_{p/q}$ corresponding to rotation numbers p/q with $q \leq 12$ (those with $q \leq 6$ are labelled at the point where they touch the axis $\kappa = 0$). The horizontal scale of the diagram is greater than the vertical, in order emphasise the structure of the tongues.

At $\kappa = 1$ the map has a cubic critical point at the origin, which means that the inverse map is not differentiable there. Figure 1.3(a) shows the Devil's staircase for this parameter value, in relation to the Arnold tongues. (It is conjectured that the structure of this

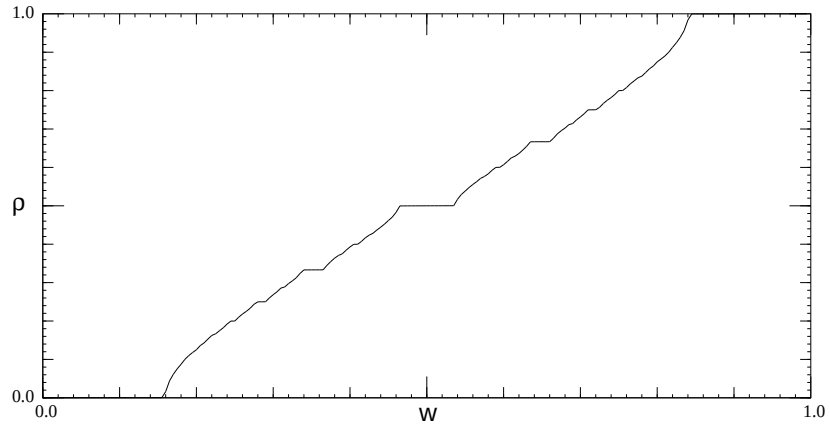

(a) Devil's staircase of mode lockings for $\kappa = 1$.

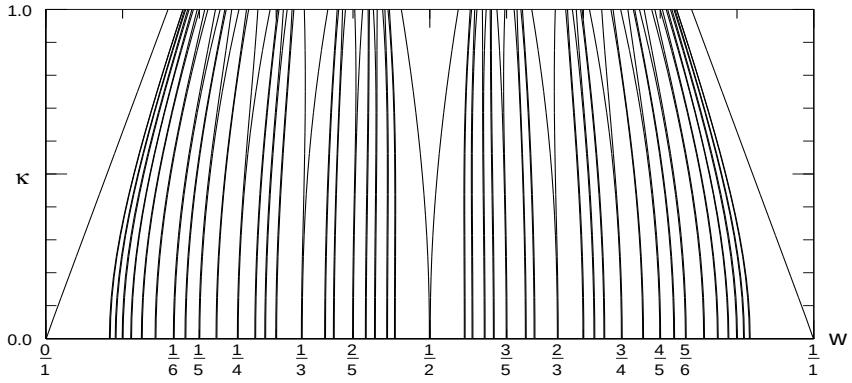

(b) Arnold tongues for $0 \leq \kappa \leq 1$.

Figure 1.3: Mode locking for the family of circle maps

staircase is universal [JBB83].) Critical circle maps are important as they correspond to the point of breakup of invariant circles in dissipative diffeomorphisms of the multi-dimensional annulus (this is discussed in [FKS82]). These occur as Poincaré sections of invariant tori for dissipative differential equations, and are therefore of direct physical significance.

For $\kappa > 1$ (not shown in the figure) the tongues overlap, the rotation number is no longer uniquely determined, and the dynamics of the map is complicated (period doubling also occurs in this regime).

Arnold's Theorem The question of the existence of a function ϕ, conjugating the system to rigid rotation, is plagued by the so-called *small divisors* (see section 1.3.2); it is heavily dependent on the number-theoretic properties of the rotation number. In particular Arnold's theorem says that ϕ exists for small enough values of the perturbation κ, provided that the rotation number satisfies a *Diophantine condition*:

Definition 2 (Diophantine)
An irrational number Ω is said to be Diophantine, *satisfying* a Diophantine condition of

order $k > 0$, *if there exists* $\varepsilon > 0$ *such that for all integers* $p, q \geq 1$*:*

$$|\Omega - \frac{p}{q}| > \frac{\varepsilon}{q^k}, \tag{1.5}$$

This property is intimately connected with the *continued fraction expansion* of the number Ω:

Definition 3 (Continued Fraction Expansion)
For a real number $0 < \Omega \leq 1$*, the* continued fraction expansion *is given by*

$$\Omega = [a_1, a_2, a_3, \ldots, a_n, \ldots] = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{\ddots}}}} \tag{1.6}$$

where the a_i *are positive integers.*

One further definition will be needed:

Definition 4 (n-th Convergent)
The n-th convergent *of a continued fraction expansion is obtained by truncating the continued fraction to give*

$$[a_1, a_2, \ldots, a_n, \infty] = \frac{p_n}{q_n}. \tag{1.7}$$

For a given real number Ω, it turns out that the n-th convergent p_n/q_n is the *closest rational approximation* to Ω having denominator not exceeding q_n, i.e.

Lemma 5 (Closest Rational Approximation)
Let p_n/q_n *be the* n*-th convergent to* Ω*. Then for any* $p/q \neq p_n/q_n$ *with* $0 < q \leq q_n$*:*

$$|\Omega - \frac{p_n}{q_n}| < |\Omega - \frac{p}{q}|. \tag{1.8}$$

Intuitively, then, the numbers which satisfy a Diophantine condition are those which are *badly approximable by rationals.* The rotation number that is most widely studied is

$$\omega = \frac{\sqrt{5} - 1}{2} = 0.618\ 034\ \ldots,$$

which is the *golden mean* (mod 1). This number is (in the above sense) the *worst approximable* by rationals. It has a particularly simple continued fraction expansion,

$$\omega = [1, 1, 1, \ldots, 1, \ldots] = \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{\ddots}}}}$$

(From now on, unless otherwise stated, the symbol ω will be used to denote $(\sqrt{5}-1)/2$.)

The golden mean has some nice algebraic properties which make it a convenient rotation number to analyse. In particular, it is intimately connected with the *Fibonacci numbers* $\{Q_n\}$, defined by

$$Q_{-1} = 0 \tag{1.9a}$$
$$Q_0 = 1 \tag{1.9b}$$
$$Q_n = Q_{n-1} + Q_{n-2} \quad \text{for } n \geq 1. \tag{1.9c}$$

This definition may be conveniently expressed in matrix form:

$$\begin{pmatrix} Q_n \\ Q_{n-1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} Q_{n-1} \\ Q_{n-2} \end{pmatrix}. \tag{1.10}$$

Denoting the matrix on the right-hand side by F, we have

$$F^n = \begin{pmatrix} Q_n & Q_{n-1} \\ Q_{n-1} & Q_{n-2} \end{pmatrix}.$$

It turns out that $(-\omega)$ is an eigenvalue of the matrix F corresponding to the eigenvector $(-\omega, 1)$, so that $(-\omega)^n$ is an eigenvalue of F^n corresponding to eigenvector $(-\omega, 1)$. Therefore

$$Q_{n-1} - Q_n\omega = (-\omega)^{n+1}. \tag{1.11}$$

For the rotation R_ω, we have

$$(R_\omega)^{Q_n}(x) - x = Q_n\omega \bmod 1,$$

which on the circle yields (using equation 1.11 and the fact that $|\omega| < 1$)

$$|(R_\omega)^{Q_n}(x) - x| = \omega^{n+1}.$$

Since $|\omega| < 1$, it follows that the Fibonacci-numbered iterates $(R_\omega)^{Q_n}$ of the rotation R_ω must converge geometrically to the identity function. Notice also that the quotients of successive Fibonacci numbers are exactly the convergents for the continued fraction expansion of the golden mean: $p_n/q_n = Q_{n-1}/Q_n$. (These convergents oscillate about ω, converging geometrically.) It follows, from the fact that the convergents are the closest rational approximations to ω, that the Fibonacci iterates of a given point x_0 are, in fact, the "nearest hits" to x_0.

These facts suggest that interesting scaling behaviour might be analysed for the original function f by looking at its Fibonacci iterates. By analogy with the period-doubling scenario outlined in the previous section, it is natural to construct a renormalization operator involving a composition of functions, along with suitable rescaling. Since the Fibonacci numbers satisfy $Q_n = Q_{n-1} + Q_{n-2}$, the Fibonacci iterates of the map are given by repeated compositions like the following:

$$(R_\omega)^{Q_n} = (R_\omega)^{Q_{n-1}} \circ (R_\omega)^{Q_{n-2}} = (R_\omega)^{Q_{n-2}} \circ (R_\omega)^{Q_{n-1}}.$$

(Recall that the desired result of renormalization is a universal function expressing some *exact* self-similarity. The bounded length of the circle imposes a largest length scale which would prevent this. It would therefore be necessary to consider the *lift* of the map of the circle to a map of the real line.)

By analogy with the definition of the Fibonacci numbers in terms of a matrix (equation 1.10), Feigenbaum et. al. [FKS82] and Ostlund et. al. [ORSS82] pointed out that it is easier to work with an operator on a space of *pairs* of functions, having a general form like either of the following

$$\begin{bmatrix} U \\ T \end{bmatrix} \mapsto \begin{bmatrix} UT \\ U \end{bmatrix} \tag{1.12a}$$

$$\begin{bmatrix} U \\ T \end{bmatrix} \mapsto \begin{bmatrix} TU \\ U \end{bmatrix}, \tag{1.12b}$$

combined with suitable rescaling operations. In future, we will refer to the composition order UT (equation 1.12a) for such an operator as the *accretive* order, and the composition order TU (equation 1.12b) as the *inverse* order. The ideas behind this choice of names will become apparent in Chapter 4.

It is worth reiterating that KAM-type results are perturbative in nature (in particular, Arnold's theorem talks of conjugacy existing provided that the perturbation from rigid rotation is not too large). More recently, a very powerful non-perturbative result has been proved by Herman [Her79], namely that for almost all rotation numbers *every* sufficiently smooth (C^3 or smoother) diffeomorphism of the circle is smoothly equivalent to rigid rotation.

Universality

Using the concepts introduced above, we will now briefly describe the universality observed and its explanation in terms of renormalization. An important feature is that there is a "simple" scaling behaviour present for the sub-critical regime which corresponds to a so-called *simple fixed point* of the relevant renormalization operator. A more complicated scaling behaviour observed for the critical regime corresponds to a *critical fixed point* of the operator.

First, consider the subcritical case. For the family of maps introduced above, this means that $|\kappa| < 1$. Now, fix the value of κ and consider the resulting diffeomorphism f_w $(= f_{w,\kappa})$. Suppose that w_* is the value of the forcing frequency w that gives rise to golden mean rotation, i.e. $\rho(f_{w_*}) = \omega$, and let w_n be the value nearest to w_* that gives rise to $\rho(f_{w_n}) = p_n/q_n$ (so that the values w_n give rotation numbers that are the successive rational approximations to the golden mean).

From Herman's theorem (above) we know that there is a smooth conjugacy to rigid rotation. From the above discussion of the golden mean the following scaling scenario may be deduced:

- $(f_w)^{q_n}(0) - p_n$ decreases like a^n where $a = -\omega$.

- For a defined as in the previous statement, the sequence of functions,

  $$\{\ a^{-n}(f_w^{q_n}(a^n x) - p_n)\ \},$$

  converges (up to a scale change) to the rigid rotation $x \mapsto x + \omega$.

- The limit $\lim_{n\to\infty} \frac{w_n - w_{n-1}}{w_{n+1} - w_n}$ exists and is equal to $\delta = -\omega^{-2}$.

The important question now is what happens at the critical parameter value when the conjugacy starts to break down. Recall that the critical parameter value for the family introduced above is $\kappa = 1$.

Numerical experiments by Shenker [She82], and others, suggest the following scaling scenario for these critical maps:

- $(f_{w_*})^{q_n}(0) - p_n$ decreases like a^n where a is no longer the golden mean, but takes the value $a = -0.766\ldots$. This is a universal constant in the same way as Feigenbaum's constant α in the period doubling scenario of the previous section.

- For a defined as in the previous statement, the sequence of functions,

  $$\{\ a^{-n}(f_w^{q_n}(a^n x) - p_n)\ \},$$

  converges to an analytic function of x^3. (Analogous to the universal function g of the previous section.)

- The limit $\lim_{n\to\infty} \frac{w_n - w_{n-1}}{w_{n+1} - w_n}$ exists. It is no longer equal to $-\omega^{-2}$, instead it takes the universal value $\delta = -2.834\ldots$. (The analogy is with Feigenbaum's δ.)

The important point to note about this situation is that there is a smooth conjugacy to irrational rotation which breaks down by becoming non-smooth in a scale-invariant manner at a critical parameter value. Essentially the same scenario occurs when examining Siegel discs (see section 1.4).

A renormalization explanation for the universality described above was constructed by Ostlund, Rand, Sethna, and Siggia [ORSS82] and Feigenbaum, Kadanoff, and Shenker [FKS82]. As indicated earlier, in the general discussion of golden mean renormalization, it is useful to work with an operator acting on a space of *pairs* of maps. (For critical circle maps this strategy also avoids several technical difficulties.) The idea here is to use pairs (ξ, η) of commuting analytic maps of the real line which are glued together to form a map of the circle.

On a suitable space of pairs, the action of the operator is given by

$$\begin{pmatrix} \xi(x) \\ \eta(x) \end{pmatrix} \mapsto \begin{pmatrix} a^{-1}\eta(ax) \\ a^{-1}\eta(\xi(ax)) \end{pmatrix}, \tag{1.13}$$

where the rescaling $a = a(\xi, \eta)$ is chosen to impose a suitable normalization condition (in fact, $a = -(\xi(\eta(0)) - \eta(0))$).

In general, this operator sends rotation numbers $\rho = p_n/q_n$ to $\rho = p_{n-1}/q_{n-1}$, with the result that it preserves the golden mean. (Additionally, it preserves criticality of the

map and a suitable "approximate commutativity" condition, exact commutativity being awkward to implement. A similar difficulty arises in the Siegel disc case, and is discussed in chapter 2.)

The renormalization explanation then goes as follows. The operator has a hyperbolic *simple fixed point*, given by the pair of maps

$$\begin{aligned} \xi(x) &= x + \omega \\ \eta(x) &= x + \omega - 1. \end{aligned}$$

This fixed point has a one-dimensional unstable manifold with corresponding eigenvalue $-\omega^{-2}$. The presence of this fixed point explains the "simple" scaling observed for sub-critical maps with golden mean rotation number.

The central conjecture offered by Feigenbaum et. al. [FKS82] and Ostlund et. al. [ORSS82] is that the renormalization operator also has a hyperbolic *critical fixed point* that accounts for the critical scaling.

A computer-assisted proof was given for both existence and hyperbolicity of the critical fixed point by Mestel [Mes85]. The proof is analogous to that given by Lanford [Lan82] for the period doubling scenario; it uses a variant of Newton's method in a suitable space of pairs of analytic functions of x^3. (The general method is described in Lanford [Lan86] and in chapters 2 and 3 of this thesis.)

1.2.3 KAM Theorem and Area preserving twist maps

This section will describe briefly the study of area-preserving twist maps (in particular, a special case of the KAM theorem will be considered) and how they relate to our results. The main references in this area are Arnold [Arn78], Moser [Mos73], and Siegel and Moser [SM71].

The study of *Hamiltonian* dynamical systems has been important in the modelling of physical processes by differential equations. A Hamiltonian system with N degrees of freedom is defined in terms of a *Hamiltonian* $H(q, p)$ that satisfies the equations:

$$\dot{q} = \frac{\partial H}{\partial p}, \quad \dot{p} = -\frac{\partial H}{\partial q},$$

where p, q are N-dimensional vectors (q representing some generalised coordinates, and p their corresponding "conjugate momenta"). The Hamiltonian is constant on solution curves; the system preserving a quantity called the *symplectic area*, which may be thought of as being somewhat analogous to the energy in a physical system.

A Hamiltonian system with N degrees of freedom is said to be (completely) *integrable* if there exist N invariants (integrals) for the motion. These invariants restrict the motion in the ($2N$-dimensional) phase space to the surfaces of N-dimensional tori. The flow induced on these tori may be characterised by a system of canonical coordinates (I, θ), the *action-angle* coordinates, in which the motion reduces to

$$\dot{I} = 0, \quad \dot{\theta} = \Omega(I).$$

(The elements of the vector Ω may be thought of as specifying frequencies of rotation.)

The existence of these invariant tori is extremely important in the analysis of these systems; they act as barriers to the motion in the phase space (a conjecture being that the onset of turbulence in many physical systems corresponds to the point where they break up). In its most general form, the so-called *Kolmogorov-Arnold-Moser* (KAM) Theorem says that under small enough perturbations from integrable most of these invariant tori will persist (albeit, perhaps, in a distorted form).

Several versions of this theorem were proved by Kolmogorov [Kol54], Moser [Mos62], and Arnold ([Arn63a] and [Arn63b]), and it was later extended to N degrees of freedom systems having Poincaré (i.e. *first return*) maps that are $(2N-2)$-dimensional (see, for example, Arnold [Arn78]).

Rather than considering the KAM theorem in its full generality, it is helpful to look at the special case of two degree of freedom Hamiltonians. For a given constant value of the Hamiltonian H, the flow in the four-dimensional phase space is restricted to a three-dimensional surface. Taking a local cross-section then yields a two-dimensional Poincaré map. This motivates the study of area-preserving maps of the plane, in particular area-preserving *twist maps*. The version of the KAM theorem relevant in this context is known as the *Moser twist theorem*.

Twist maps are usually taken to be defined on annuli (equivalently cylinders); the radial (or longitudinal) and angular variables being used as coordinates. In the context of the cylinder, a twist map is essentially one that sends verticals into spirals. (On the annulus the orbits of the map lie on circles, where the rotation number for each circle depends on the radial coordinate.) More formally, consider a map $f : (x, y) \mapsto (f_1(x, y), f_2(x, y))$ where $\partial f_1/\partial y$ has a fixed sign.

When working in the context of the cylinder, it is useful to regard twist maps as periodic maps of the plane, requiring that they commute with the backward rotation $R : (x, y) \mapsto (x - 1, y)$.

A well known example is the *standard family* defined by

$$\begin{aligned} x' &= x + y' \\ y' &= y - \frac{\kappa}{2\pi}\sin(2\pi x), \end{aligned}$$

the members of which are known as *standard maps*. As mentioned above, twist maps that preserve area are important because they are section maps of Hamiltonian systems that are models of physical phenomena (for example, the restricted three-body problem in celestial mechanics).

The analogue of KAM tori for this situation are invariant curves (on which the motion is conjugate to rotation) which go *around* the cylinder, acting as a barrier to the dynamics of the map. The dynamics induced on the curve is then equivalent to that of a map of the circle, as described in the previous section. The Moser twist theorem says that these curves exist provided that the rotation number is sufficiently irrational (once again, the condition is that the number be Diophantine) and that the perturbation κ is sufficiently small: by analogy with the circle maps introduced in the previous section, $\kappa = 0$ corresponds to a

completely integrable linear shear with the invariant curves being horizontal lines, whereas for $\kappa > 0$ the invariant curves with rational rotation numbers are destroyed and only those that are sufficiently irrational persist, provided that κ is not too large. (For this particular system, the last curve to break up appears to correspond to the golden mean rotation number, as often — though certainly not always — happens to be the case.)

Proving the existence of invariant curves on which the motion is conjugate to irrational rotation involves establishing that a formal Fourier series converges, which leads to a "small divisor problem" (see section 1.3.2). The situation is analogous to the proof of existence of invariant circles for iterated complex maps, which is one of the main topics of this thesis (the first solution to this type of problem being given by Siegel).

The relevant question is once again one of the scaling behaviour present and, in particular, of what happens when the final invariant curve breaks up, i.e. what is the structure of the invariant curve at the *critical* parameter value, beyond which it ceases to exist?

Once again, a "critical" universality is observed which relates to the existence of a *critical fixed point* of the relevant renormalization operator. In this case it is generally known as *MacKay's operator*, acting on a space of area-preserving symmetric periodic twist maps. As before, it turns out to be advantageous to consider pairs $(\mathcal{U}, \mathcal{T})$ of maps (this is useful here as it allows a convenient scheme to be implemented that preserves periodicity), the operator being given by

$$\begin{pmatrix} \mathcal{U} \\ \mathcal{T} \end{pmatrix} \mapsto \begin{pmatrix} \mathcal{B}\mathcal{T}\mathcal{B}^{-1} \\ \mathcal{B}\mathcal{T}\mathcal{U}\mathcal{B}^{-1} \end{pmatrix},$$

where $\mathcal{B} = \mathcal{B}(\mathcal{U}, \mathcal{T})$ is a suitable (linear diagonal) rescaling map. (Notice, once again, that this operator has the same general form as those given equation 1.12, in the general discussion of golden mean renormalization.) A detailed study of this operator was undertaken by MacKay [Mac82] revealing the following explanation for the universality.

The sub-critical scaling is again governed by the existence of a simple commuting fixed point for the operator, corresponding to a completely integrable map. This fixed point is given by

$$\mathcal{U}(x, y) = \begin{pmatrix} x + \omega y - 1 \\ y \end{pmatrix}, \quad \mathcal{T}(x, y) = \begin{pmatrix} x + \omega + y \\ y \end{pmatrix}.$$

Maps attracted by the simple fixed point have a smooth golden circle (this follows by the necklace construction of Stirnemann [Sti93c], see also below).

The conjecture is that there is also a hyperbolic critical fixed point, attracting maps with a non-smooth golden circle (where non-smooth refers to the conjugation function, rather than to the curve itself). The idea is then that maps lying "below" the stable manifold of the critical fixed point (sub-critical) are attracted to the simple fixed point and carry a smooth golden circle, whereas those lying on the stable manifold itself have a critical golden curve that is the last invariant curve to break up. (Maps lying "above" the stable manifold would not have an invariant circle at all.) A schematic representation of the function space is given in figure 1.4.

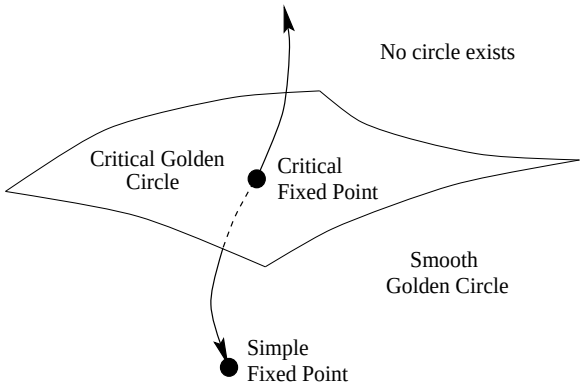

Figure 1.4: The Simple and Critical Fixed Points.

No proof of the existence of the critical fixed point has yet been published although it is, in principle, possible to construct a proof of the kind given by Lanford [Lan82] for the period doubling scenario and Mestel [Mes85] for critical circle maps (unfortunately, the numerics of this problem are much less well-behaved than those encountered previously).

Stirnemann [Sti92] has studied a "three-step" operator that has the advantage that it preserves symmetry, irrespective of whether the two maps commute or not. Attention can then be focussed on a subspace of symmetric twist maps. An attempt to prove existence of a critical fixed point by Lanford's method (as mentioned above) is described in [Sti92], but did not succeed. (More recently, however, hope has returned that a proof using this method *is* possible [Sti].)

By assuming that the critical fixed point *does* exist, and using a slightly modified operator (namely one with the composition order $\mathcal{UT}$ rather than $\mathcal{TU}$, which results in some convenient combinatorial properties analogous to those of the Fibonacci numbers, see section 4.2.3), Stirnemann [Sti93c] was able to deduce the existence of a transitive invariant golden circle for all twist maps attracted by the critical fixed point. The proof rests on the assumption that the fixed point satisfies some topological conditions known as the *necklace hypotheses*. The method of proof (developed by Stirnemann for this problem) is called the *necklace construction*, and will be adapted and used in this thesis (chapters 4 and 5) to analyse golden circles in a related problem, namely Siegel discs for iterated complex maps. The next section serves to introduce the necessary background to this problem, after which it will be described in detail.

In terms of complexity, the problem that we will concentrate on (being one-dimensional, complex) may be thought of as being midway between the Feigenbaum period doubling / circle maps problems (one-dimensional, real) and the area-preserving twist maps problem (two-dimensional, real).

1.3 Iterated Complex Maps and Siegel Discs

In this subsection, we briefly describe the subject of iterated maps of the complex plane. These are studied as dynamical systems in their own right, in addition to being regarded in some sense as models for more complicated systems. It is here that the concept of a *Siegel disc* is encountered.

By the late nineteenth century, the study of iterated complex analytic mappings was quite well advanced concerning their behaviour in the neighbourhood of fixed points. However, very little was known about the global behaviour except in very simple cases. Early this century Pierre Fatou demonstrated (by giving an example) that the global behaviour can be surprisingly complicated. After the first world war, the subject was studied more extensively, in particular by Fatou [Fat19], Gaston Julia [Jul18], and others, and in recent years this area has enjoyed a period of growth and renewed interest, mainly due to developments in computer technology which allow any competent hacker to produce pictures of the stunningly intricate and beautiful sets that arise. A good general reference in this area is Beardon [Bea91].

1.3.1 Julia Sets and Fatou Components

The *Fatou set* (or *stable set*) for a complex map f is defined to be the (open) set of points that have a neighbourhood on which (the restriction of) the iterates $\{f^n\}$ of f form a *normal family*:

Definition 6 (Normal Family)
A family F of maps from a metric space (X, d) to a metric space (X', d') is said to be normal *in X if every infinite sequence of functions from F contains a (locally uniformly) convergent subsequence.*

The idea is that for points in the Fatou set the family of iterates of the function f *preserves the proximity of points.* The formal notion is that of *equicontinuity.* Recall, first, the definition of *continuity*:

Definition 7 (Continuity)
A map f from (X, d) to (X', d') is continuous *at a point $x_0 \in X_0 \subseteq X$ if, for every $\varepsilon > 0$, there exists a constant $\delta(f, \varepsilon, x_0) > 0$ such that for all $x \in X_0$:*

$$d(x_0, x) < \delta \quad \textit{implies} \quad d'(f(x_0), f(x)) < \varepsilon.$$

Notice that δ generally depends on the function f. If a constant δ can be found that works both for all x *and* all f in some family of functions, then the family is called *equicontinuous.* In other words:

Definition 8 (Equicontinuity)
A family F of maps f from (X, d) to (X', d') is equicontinuous *at a point $x_0 \in X_0 \subseteq X$ if, for every $\varepsilon > 0$, there exists a constant $\delta(\varepsilon, x_0) > 0$ such that for all $x \in X_0$ and all $f \in F$:*

$$d(x_0, x) < \delta \quad \textit{implies} \quad d'(f(x_0), f(x)) < \varepsilon.$$

i.e. every function in F maps the open ball $B_\delta(x_0)$ (i.e. $\{x : d(x_0, x) < \delta\}$) into a ball $B_{\varepsilon'}(f(x_0))$ with $\varepsilon' \leq \varepsilon$. It turns out (by the Arzelà-Ascoli Theorem, e.g. [Rud82]) that equicontinuity and normality are equivalent for continuous complex maps from a subset of the sphere $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ to $\hat{\mathbb{C}}$.

The *Julia set* (or *unstable set*) is then the complement of the Fatou set, i.e. the (compact) set of points that do not have a neighbourhood on which the iterates of f form a normal family. An equivalent (and perhaps more intuitive) definition is to say that the Julia set is the *closure of the set of repelling periodic points* of the map f.

(Intuitively, a point is in the Fatou set if it has a neighbourhood on which the dynamics of f is "tame" in the sense that f^n preserves the proximity of points; the point is in the Julia set if the dynamics is more "wild", driving points apart.)

Although the results concerning Julia sets are simpler in the case where the map f is a polynomial, they are more properly considered in the context of *rational* maps, i.e. those that may written as the ratio of two polynomials $f(z) = p(z)/q(z)$ on the Riemann sphere $\hat{\mathbb{C}}$. (In fact, it is a fundamental result that any map that is analytic on the whole sphere must be a rational map, [Bea91]. Thus, maps that are analytic on the whole sphere may have only a *finite* number of poles and zeros.)

Here are a few general results:

- The Julia set is typically a complicated object, often a fractal exhibiting self-similarity.
- The Julia set of the iterates of f is the same as that of f.
- All attractive periodic orbits, along with their basins of attraction, are contained within the Fatou set. The Julia set is the boundary of the basin of attraction of each attractive periodic orbit (perhaps including infinity).
- If a Julia set has an interior point, then it must be equal to the entire Riemann sphere.
- If f has degree 2 or more, then the Julia set has no isolated points. (Note that this is different to saying that the Julia set is connected.)
- For a generic point in the Julia set, the forward orbit of the point is everywhere dense in the Julia set.

It turns out that the dynamics of the map is controlled by the dynamics of the orbits of its *critical points*. There are many possible kinds of Julia set, even for very simple families of mapping. This is well illustrated by the quadratic family defined by

$$f_c : z \mapsto z^2 + c,$$

where c is complex. The famous *Mandelbrot set* M is the set of values of the parameter c for which the orbit of the critical point $z_{\text{crit}} = 0$ is bounded (corresponding to maps that have a connected Julia set). Much deep and beautiful mathematics has come from the study of this set, [Dou86], [Lei90]. (The set is shown in figure 1.5. The labelled points correspond to the Julia sets shown in figure 1.6 and figure 1.7.)

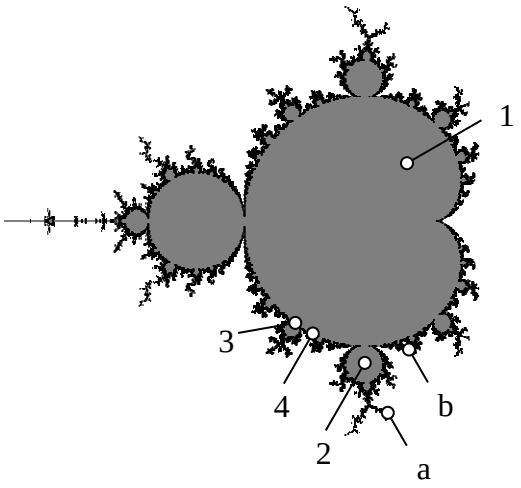

Figure 1.5: The Mandelbrot Set

The type of Julia set that is observed for a given value of the parameter c depends on the position of c in relation to the Mandelbrot set. The Julia set is connected if c lies within M and disconnected if not, so the boundary of M is of particular interest. In addition, each part of M corresponds to a different qualitative character of the Julia sets associated with it.

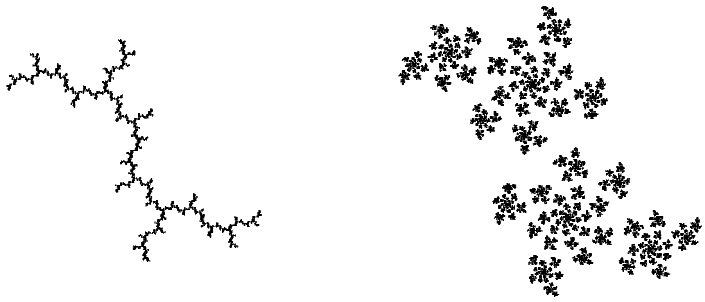

(a) Dendrite $c = -i$

(b) Disconnected Dust $c = 0.11031 - 0.67037i$

Figure 1.6: Two kinds of Julia set having a single connected Fatou component.

The maximal connected subsets of the Fatou set are referred to as *Fatou components*. There are cases where the Fatou set consists of a single (unbounded) connected component. For example, some Julia sets are tree-like structures called *dendrites*. This happens when an iterate of the critical point is periodic, which corresponds to the case where c is on one of the "hairs" of M (see figure 1.6(a)). Other Julia sets are totally disconnected "dusts",

as occurs when f has no attractive orbits and c lies outside M (see figure 1.6(b)).

The cases where the Fatou set has more than one connected component (so that parts of the Julia set surround finite regions of the plane) may be classified as follows:

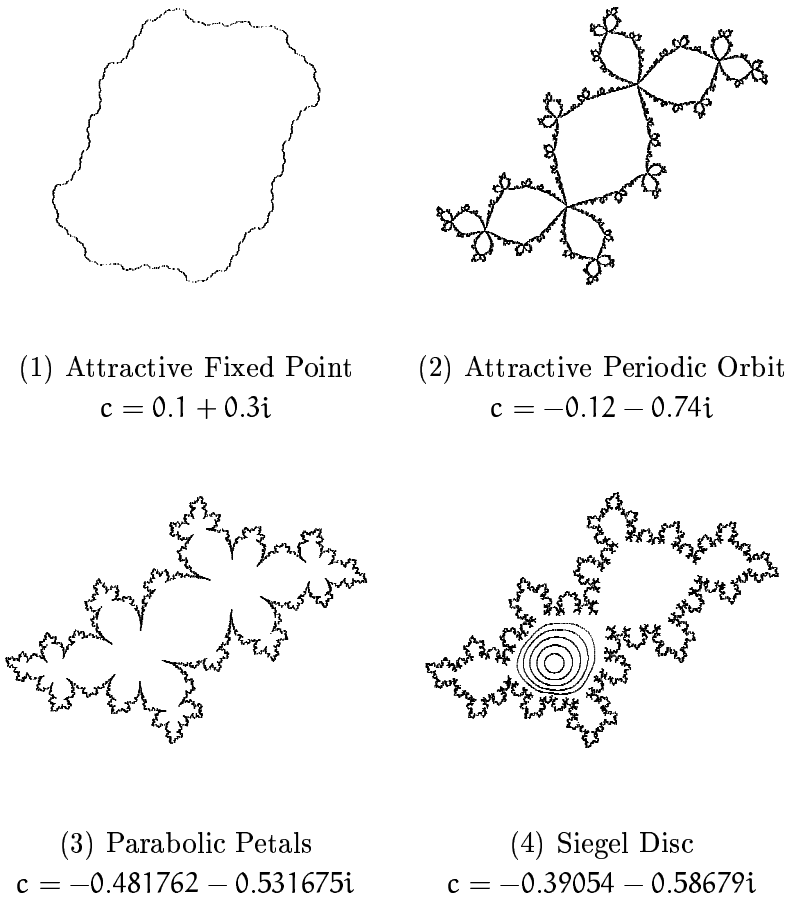

(1) Attractive Fixed Point
$c = 0.1 + 0.3i$

(2) Attractive Periodic Orbit
$c = -0.12 - 0.74i$

(3) Parabolic Petals
$c = -0.481762 - 0.531675i$

(4) Siegel Disc
$c = -0.39054 - 0.58679i$

Figure 1.7: Four Kinds of Julia Set with Bounded Fatou Components

1. When c lies in the main cardioid-shaped body of M, then the Julia sets are fractally-deformed quasi-circles, surrounding a single attractive fixed point, as shown in figure 1.7(1).

2. If c lies in one of the "buds" that is attached to the main body then f has an attractive periodic orbit and the Julia set consists of an infinite number of fractally deformed circles that are grouped in an intricate hierarchy around the points of the attractor and their pre-images, as shown in figure 1.7(2).

3. As the parameter value c passes from the main body to a bud (or from a bud to a smaller bud) through a narrow "neck", the dynamics of the Julia set undergoes a period n-tupling (on the real axis this corresponds to the Feigenbaum period-doubling scenario described earlier in section 1.2.1). Once again, as in the one-dimensional real case, aspects of these transitions are *universal*. At such a neck,

where buds join, the Julia set is *parabolic* and consists of a number of "petals". See figure 1.7(3).

4. For certain values of the parameter c, there may be a *Siegel disc*. In this case, the Julia set consists of an infinite number of curves bounding open regions, with f mapping each region into a larger one, and eventually into the region containing the fixed point. Once points land inside this region, which is known as a Siegel disc, they rotate on invariant circles around the fixed point. Figure 1.7(4) shows both the Julia set, and some of the quasi-circular curves in the interior of the Siegel disc.

(For general maps there is also a fifth case, known as a *Herman ring*, which does not occur for polynomial maps.) Sullivan [Sul85] proved that all of the Fatou components are eventually periodic. This settled a long-standing conjecture that the Fatou set has *no wandering domains*.

It is the fourth case given above, namely the Siegel disc, that this thesis focusses on. The following sections will introduce the relevant theory to clarify what is meant by a Siegel disc, will discuss the conditions under which the structure occurs, and will discuss the nature of the boundary of Siegel discs, with which our results are concerned.

1.3.2 Linearization and the Schröder Equation

In order to examine the dynamics of maps near to fixed points, it is natural to look at functions $z \mapsto f(z)$ of the form

$$f(z) = \lambda z + \sum_{j \geq 2} f_j z^j. \tag{1.14}$$

(i.e. $f(z)$ has a fixed point of multiplier λ at $z = 0$) where the power series on the right is convergent on some neighbourhood of the origin.

The question arises of when such a function will be locally *linearizable*. That is, under what conditions does a (smooth, invertible) function φ exist that conjugates f to its linear part, $w \mapsto \lambda w$, on some neighbourhood of the origin? In other words, the question is whether there exists some neighbourhood S of the origin within which there is a smooth invertible transformation $\varphi : w \mapsto z$ such that for all $w \in S$:

$$(\varphi^{-1} \circ f \circ \varphi)(w) = \lambda w. \tag{1.15}$$

Equivalently, the following equation holds for the function φ:

$$\varphi^{-1} \circ f - \lambda \varphi^{-1} = 0. \tag{1.16}$$

This equation and its variants are known as the *Schröder* (functional) *equation* after E Schröder who studied conjugacy functions in the 19th century. (Schröder was investigating methods for computing the iterates of rational functions [Sch71] in connection with algorithms for approximating roots [Sch70].)

For the case $|\lambda| \neq 0, 1$, G Koenigs (1884) established that there is indeed a local analytic change of coordinates $w = \varphi(z)$, with $\varphi(0) = 0$, such that $\varphi \circ f \circ \varphi^{-1}$ is just the linear map

$w \mapsto \lambda w$ for all w in some neighbourhood of the origin (the conjugacy being unique up to multiplication by a non-zero constant). Analogous results were established by Poincaré.

In the case where λ is a root of unity, say $\lambda^q = 1$, then (provided that $f^{\circ q}$ is not the identity) the fixed point is said to be *parabolic*. The dynamics of parabolic fixed points consists of alternating attracting and repelling *petals* and is not linearisable (see, for example, Beardon [Bea91]).

Conjugacy to Rotation and Small Divisors

It remains to see what happens when $|\lambda| = 1$ but λ is *not* a root of unity. To this end, let

$$\lambda = e^{2\pi i\Omega}, \quad \Omega \text{ real, irrational.} \tag{1.17}$$

i.e. the origin is an *irrationally indifferent* fixed point. The question of the existence of a function φ conjugating f to the irrational rotation $w \mapsto \lambda w$ is known as the *centre problem*; the existence of such a function would imply that families of concentric circles about the origin in the w-plane give rise to invariant neighbourhoods of the fixed point for f in the z-plane.

Consider the Schröder equation for this problem,

$$\varphi(\lambda w) = f(\varphi(w)).$$

By writing φ as a series with undetermined coefficients, a *formal* solution is easily obtained for the case where $\lambda^q \neq 1$. Putting

$$\varphi(w) = \varphi_1 w + \sum_{k\geq 2} \varphi_k w^k,$$

and then substituting this into the Schröder equation gives

$$\begin{aligned} \sum_{k\geq 2} (\lambda^k - \lambda)\varphi_k w^k &= \sum_{j\geq 2} f_j(\varphi(w))^j \\ &= \sum_{j\geq 2} f_j\big(\varphi_1 w + \sum_{\ell\geq 2} \varphi_\ell w^\ell\big)^j. \end{aligned}$$

Comparing coefficients yields a recursive formula for the φ_k. On the left, the term in w^k involves only φ_k, whereas on the right the corresponding terms involve φ_j with $2 \leq j \leq k-1$:

$$\varphi_k = \frac{1}{\lambda^k - \lambda} \sum_{j=2}^{k} f_j \sum_{\ell_1+\cdots+\ell_j=k} \varphi_{\ell_1} \cdots \varphi_{\ell_j}.$$

This formal solution to the equation is known as the *Schröder series*. Of particular interest, then, is the question of when this series will converge.

Notice that the expressions for the coefficients have a denominator $\lambda^k - \lambda$, which will be small if λ^k is ever close to λ, giving a so-called *small divisor*. (This indicates why no solution can be found in the case where λ is a root of unity.) Although for irrational Ω the

denominator is never zero, the presence of small divisors will hinder convergence, leading to a *small divisor problem.*

In 1927, after much argument about whether linearization was possible, Cremer [Cre27] proved that the series does *not* converge for values of Ω that are well approximable by rationals:

Definition 9 (Cremer Condition)
An irrational Ω is said to satisfy a Cremer condition of degree d *if $\lambda = e^{2\pi i\Omega}$ satisfies*

$$\limsup_{q\to\infty} \frac{\log\log(1/|\lambda^q - 1|)}{q} > \log d. \tag{1.18}$$

(Intuitively, $|\lambda^q - 1|$ must tend to zero extremely quickly as q increases. In other words, Ω must be close to being a root of unity.)

Theorem 10 (Cremer's Non-linearization Theorem)
Let z_0 be a fixed point of multiplier $\lambda = e^{2\pi i\Omega}$ for an arbitrary rational function of degree two or more. If Ω satisfies a Cremer condition, then z_0 is the limit of an infinite sequence of periodic points and therefore no linearization is possible in a neighbourhood of z_0.

In fact, it turns out that a *generic* real number Ω satisfies this condition. The notion of genericity being used here is a topological one (rather than a measure-theoretic one), namely:

Definition 11 (Genericity)
A property of a complex number $\lambda = e^{2\pi i\Omega}$ on the unit circle (corresponding to the real number Ω) is said to be generic *if the set of λ for which the property holds contains a countable intersection of dense open subsets of the circle.*

To recap: under the above definition, Cremer had established that no linearization is possible for a *generic* choice of Ω.

The question of whether this was true for *all* irrational numbers on the unit circle remained open until 1942, when it was solved by Siegel [Sie42]. Siegel showed that a linearization *is* in fact possible for *almost all* values of λ (i.e. for all λ lying outside a set having one-dimensional Lebesgue measure zero). Thus there is a striking difference here between the behaviour for *generic* λ (in the topological sense), and that for *almost all* λ (in the measure-theoretic sense).

Theorem 12 (Siegel's Theorem)
If Ω is Diophantine (recall definition 2), then the corresponding Schröder series has a nonzero radius of convergence.

Thus linearization is possible if Ω is not "too close" to being rational. In the cases where linearization is possible, the neighbourhood of the fixed point on which the conjugacy holds is known as a *Siegel disc* or *Siegel domain*; it corresponds to a rotating disc in the "linear" (w) coordinates.

An alternative definition (to that of equation 1.5) of the Diophantine (or *Siegel*) numbers may be given in terms of the properties of the convergents p_n/q_n of the continued fraction expansion for Ω (recall Definition 3), as follows:

Definition 13 (Siegel Condition)
A number Ω having n-th convergent p_n/q_n is said to satisfy the Siegel condition *if*

$$\sup \frac{\log q_{n+1}}{\log q_n} < \infty. \tag{1.19}$$

In fact, *almost all* real numbers are Diophantine (Siegel) numbers, i.e. the complement of the Diophantine numbers has Lebesgue measure zero. It follows that given a randomly chosen real number Ω then, with probability one, any map having a fixed point of multiplier $e^{2\pi i\Omega}$ will have a Siegel disc.

Siegel's work gave the first proof of convergence for a small divisor series, making use of the majorant series method and some delicate number-theoretic lemmas. (The method involves constructing a "majorising series", which converges and whose coefficients exceed those of the small divisor series. The comparison test for series then establishes convergence.)

Siegel's theorem has also been proved more recently using an iterative method that constructs a sequence of functions, ultimately converging to the conjugacy function itself (see, for example, Beardon [Bea91]). Essentially the same method was used by Arnold to prove a version of the KAM theorem (recall section 1.2.3).

In 1972, Brjuno proved that the Diophantine condition could be weakened, giving a much sharper version of Siegel's theorem.

Definition 14 (Brjuno Condition)
A real number Ω with n-th convergent p_n/q_n is said to be a Brjuno number*, satisfying the* Brjuno condition*, if*

$$\sum_n \frac{\log q_{n+1}}{q_n} < \infty. \tag{1.20}$$

Theorem 15 (Brjuno's Theorem)
Let Ω be any Brjuno number, then any map of the form

$$f(z) = e^{2\pi i\Omega}z + \sum_{j\geq 2} f_j z^j,$$

is (locally) analytically conjugate to its linear part.

The question of whether this condition is optimal was answered recently (in 1987) by Yoccoz [Yoc88], who showed that if the Brjuno condition is not satisfied (i.e. if Ω is such that $\sum \log(q_{n+1})/q_n = \infty$) then the quadratic map

$$f(z) = e^{2\pi i\Omega}z + z^2, \tag{1.21}$$

has an infinite number of periodic points in every neighbourhood of the origin and, therefore, is not linearisable. (Indeed, there is strong evidence, particularly in the light of this proof, that *all* maps for which the Brjuno condition is violated are not linearisable.) Brjuno's result is therefore the best possible.

In addition, a special case of Siegel's theorem (namely the case where the rotation number is any integer multiple of the golden mean) has been proved by Stirnemann [Sti93a] using the *necklace construction* of [Sti93c]. (The proof demonstrates one way in which renormalization may be used to tackle a small divisors problem; it utilises the so-called "simple fixed point" of a suitable operator.) The necklace construction actually constructs a parameterisation for the invariant curve directly, and is a key tool used for establishing many of the results presented in this thesis (see Chapters 4 and 5).

The next section describes the universality observed for Siegel discs, and its corresponding renormalization explanation, after which a proof of the existence of a critical fixed point for the relevant renormalization operator will be discussed.

1.4 Universality for Siegel Discs

Recall that the neighbourhood S of the fixed point on which the conjugacy to rigid rotation holds is called a *Siegel domain*, or *Siegel disc*. Every point in the Siegel disc lies on an invariant curve with the same rotation number. Having discussed the conditions for the existence of such an object, the main interest now lies in determining the nature of any scaling behaviour present and, particularly, the nature of the *boundary* ∂S, where the conjugacy breaks down.

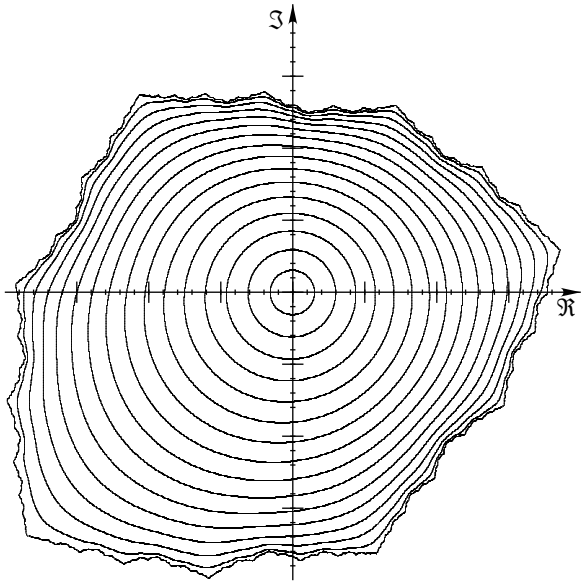

Figure 1.8: A Siegel Disc.

Figure 1.8 shows the boundary curve for a Siegel disc, along with some of the smooth curves contained in the interior.

In their paper [MN83], Manton and Nauenberg made a numerical study of the behaviour both inside the Siegel disc and on the boundary for a variety of maps and found results analogous to those discussed for circle maps in section 1.2.2 (and twist maps, section 1.2.3).

For golden mean rotation number, they found that the Schröder function appears to be

differentiable inside the disc (i.e. there is a *smooth* conjugacy to rigid rotation). So there is a scaling behaviour in the high numbered Fibonacci iterates, which is such that:

$$\lim_{n\to\infty} \frac{f^{Q_{n+1}}(z) - z}{f^{Q_n}(z) - z} = -\omega.$$

(The curves being smooth, the iterates approach the starting point asymptotically along a straight line.)

They then proceeded to examine the boundary ∂S. In particular they found that it appears to be a fractal curve passing through a critical (i.e. stationary) point z_c of the map, and having some interesting universal scaling properties. On the boundary the conjugacy is not differentiable, so the behaviour is different from the "simple" scaling observed above and, as before, will be referred to as "critical" scaling behaviour. They found that the odd and even numbered Fibonacci iterates approach the critical point asymptotically along two straight lines separated by an angle $2\theta_+$ ($\approx 107.3\ldots$ degrees), with the distance between successive odd or even iterates reducing like α^{-n}. More formally,

$$\lim_{n\to\infty} \frac{f^{Q_{n+1}}(z_c) - z_c}{f^{Q_n}(z_c) - z_c} = \alpha e^{(-1)^{n+1} 2i\theta_+},$$

where z_c is the critical point, and α is a real scale factor ($\alpha = -0.7419\ldots$). In other words, the scaling from one Fibonacci iterate to the next may be expressed as a reflection in some straight line passing through the critical point, followed by a real linear scaling.

Since the maps considered are analytic, they are also conformal except at the critical point, and it follows that there is a similar scaling at the pre-images of the critical point. It also follows, provided $f''(z_c) \neq 0$, that the same kind of scaling is shared by the forward iterates of z_c. In particular,

$$\lim_{n\to\infty} \frac{f^{Q_{n+1}}(f(z_c)) - z_c}{f^{Q_n}(f(z_c)) - z_c} = \alpha e^{(-1)^{n+1} 4i\theta_+}.$$

(Notice that the angle is now twice the value that is observed for scaling at the critical point.)

The images and pre-images of the critical point are dense on the boundary of the Siegel disc, which means that the scaling behaviour described above prevents the boundary from being even piecewise differentiable. A similar scaling behaviour occurs for the Fibonacci-numbered pre-images of the critical point. By taking a quadratic map, for example, and choosing the square root so that the branch cut is directed outside the boundary ∂S, we produce an inverse function f^{-1} that has the same invariant curves as the original f, with a critical scaling behaviour for the Fibonacci pre-images given by

$$\lim_{n\to\infty} \frac{f^{-Q_{n+1}}(z_c) - z_c}{f^{-Q_n}(z_c) - z_c} = \alpha e^{(-1)^{n+1} 2i\theta_-},$$

where θ_- is another constant angle ($2\theta_- \approx 119.6\ldots$ degrees).

Figure 1.9 shows the invariant curve passing through a point z_s in the "simple" scaling regime for an example map with golden rotation number. The Fibonacci iterates of z_s are marked; notice that they approach z_s asymptotically along a straight line. Shown

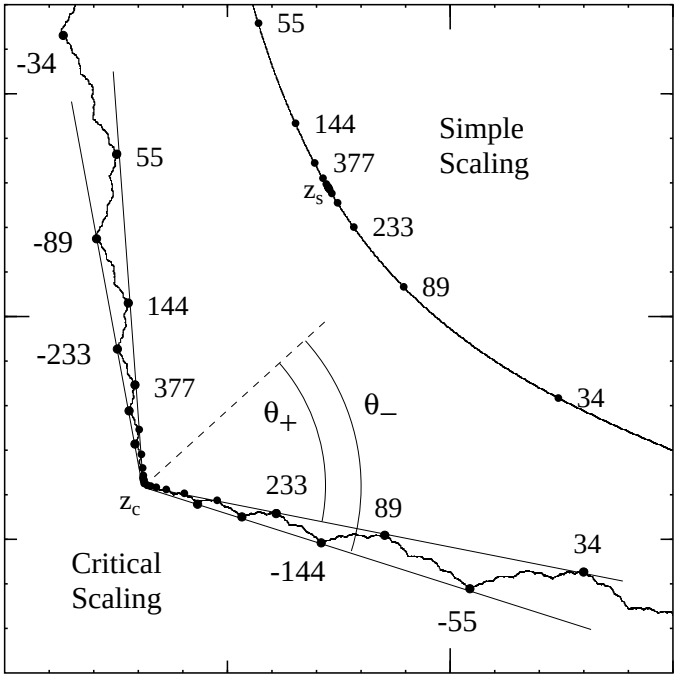

Figure 1.9: Simple and Critical Scaling for Golden Siegel Discs.

also is the critical invariant curve, along with the Fibonacci iterates and the Fibonacci pre-images of the critical point z_c, demonstrating the critical scaling mentioned above.

Manton and Nauenberg used the asymptotic self-similarity observed above to define a new curve which is non-closed (in fact, unbounded) and *exactly* self-similar. Intuitively, this is done by repeatedly magnifying the region around the critical point so that high numbered Fibonacci iterates are mapped onto lower ones. They conjectured that this curve, which will be called the *critical curve*, is *universal* for a large class of maps with golden mean rotation number (it will therefore sometimes be called the *universal curve*).

By choosing a suitably transformed coordinate system having the origin as the critical point and having the positive real axis as the symmetry line, it is possible to introduce two "scaling functions" which describe the the behaviour observed above, and which are conjectured to be universal. (In this thesis we adopt a slightly different approach: it turns out to be more convenient to use a *complex* scale factor α, rather than a real scaling and a rotated coordinate system.)

Based on these experimental findings, Manton and Nauenberg formulated the following conjectures supposed to hold for a "large class of maps":

1. The Schröder series converges on its circle of convergence. (The motion on the Siegel disc boundary having rotation number Ω for maps with multiplier $e^{2\pi i\Omega}$ at the fixed point.) The image of the circle of convergence, i.e. the boundary ∂S, is a Jordan curve passing through the critical point of the map that is closest to the fixed point.

2. The formal derivative of the Schröder series does not converge on the boundary, which suggests that the boundary curve is nowhere differentiable. (This contrasts

with the fact that, since an analytic function is infinitely differentiable within its circle of convergence, *every* curve *inside* the disc must be smooth.)

3. The Hausdorff dimension of the boundary curve is strictly greater than 1. (This means that the curve is a fractal.)

4. In the case of golden mean rotation number ($\Omega = \omega$), the boundary curve exhibits a universal self-similar structure that may be expressed by means of a scaling function, which is a simultaneous solution of two functional equations. It is essentially unique.

(It should be noted that the first few conjectures are for general rotation numbers having a Siegel disc, whereas the last refers to the special case of the golden mean.)

Widom's fixed point

The last conjecture was clarified by Widom [Wid83], who used renormalization group methods to numerically solve the functional equations for the special case of the quadratic map (equation 1.21), obtaining a very accurate approximation of the scaling function. Further, Widom demonstrated how to extend the result to families of mappings by considering perturbations of the scaling function. As before (section 1.2.2), it turns out to be easier to express the functional equations as a transformation on a space of *pairs* (E, F). The scaling function for the Siegel disc boundary then corresponds to a *critical fixed point* of the relevant operator that is found, empirically, to be a pair of *even* maps (i.e. maps of z^2). Iterating the renormalization operator numerically on the quadratic map, for example, yields a sequence of maps which appears to converge to the critical fixed point.

(Widom also presented strong evidence that the critical point lies on the boundary of the Siegel disc. This was later proved for polynomial maps by Rogers [Rog95].)

By considering a slightly modified operator (which will be given below), Stirnemann [Sti93b] was able to give a computer-assisted proof for the existence of the critical fixed point conjectured by Widom. (Further details are presented in chapter 2.) The proof was repeated, both in order to verify and adapt the computational part, and to obtain (a ball around) the fixed point for use in proving the later results presented in this thesis.

The remainder of this thesis deals with the remaining conjectures in the case of golden rotation number. Chapter 4 proves that in the domain of attraction of Widom's fixed point the boundary value of the Schröder series is an asymptotically self-similar Hölder continuous Jordan curve through a stationary point of the map. Further, it is not differentiable on a dense set of points. We obtain these results by verifying that Widom's fixed point satisfies the hypotheses of the *necklace construction*. The construction was developed in [Sti93c], where the objects of interest were twist maps of the cylinder (recall section 1.2.3). Part of chapter 4 is concerned with modifying the necklace construction so that we may use it for Widom's fixed point.

The conjecture concerning Hausdorff dimension will be dealt with in chapter 5. It turns out that the Hausdorff dimension is completely determined by the dimension of the invariant curve of the fixed point itself, i.e. the universal curve. It is very probable that the

generalised (dynamical) dimensions can be obtained along the same lines. All of them are renormalization invariants.

In order to characterise the "large class of maps" mentioned above, information would be needed concerning the spectral properties of the renormalization operator. Unfortunately, it has not even been possible to prove that Widom's fixed point is hyperbolic (the problem is that a "contracted" matrix containing the derivative of the operator is too far from diagonal with respect to the basis that we were using). However, the numerical evidence strongly suggests that it is, and that its stable manifold has real codimension-two (in a subspace of even commuting pairs of maps). The results should, therefore, hold for an *open* set of maps having a golden mean Siegel disc.

1.4.1 The Siegel Disc Operator

This section will define the relevant operator and will set the stage for discussing the existence of a critical fixed point.

Definition 16 (Rescaling Factor)
Let D_E *and* D_F *be non-empty connected open subsets of the complex plane, containing the origin. Consider the space of pairs* (E, F) *of analytic functions in one variable, defined on* D_E *and* D_F *respectively. The complex number*

$$\alpha = \alpha(E, F) = F(0), \tag{1.22}$$

will be called the rescaling factor *of the pair* (E, F).

In this thesis, two different forms of the Siegel disc operator will be used. The difference between them lies in the *order of composition* of the maps E and F.

The following form (which is closest to that originally considered by Widom [Wid83]) uses the so-called *inverse* composition order FE, as used in MacKay's operator (recall section 1.2.3). Chapter 2 establishes the existence of a fixed point for this operator.

In what follows, C denotes complex conjugation ($C : z \mapsto \bar{z}$) and juxtaposition of symbols will be used to indicate composition of the corresponding maps.

Definition 17 (Renormalization Operator — Inverse Order)
The (inverse-order) Siegel disc renormalization operator *is defined as follows. The domain of definition of the operator is the set of pairs* (E, F) *such that*

$$\alpha C(D_E) \subset D_F \tag{1.23}$$

$$\alpha C(D_F) \subset D_E \tag{1.24}$$

$$E(\alpha C(D_F)) \subset D_F. \tag{1.25}$$

The operator acts on pairs (E, F) *by means of* $(E, F) \mapsto (\tilde{E}, \tilde{F})$, *where*

$$\tilde{E}(z) = C\alpha^{-1}F\alpha C(z), \tag{1.26}$$

$$\tilde{F}(z) = C\alpha^{-1}FE\alpha C(z). \tag{1.27}$$

(Here, α, respectively α^{-1}, stands for the function $z \mapsto \alpha z$, respectively $z \mapsto \alpha^{-1} z$.)

Note: the operator $f(z) \mapsto CfC(z)$ has the effect of complex conjugation of the coefficients of power series. Suppose that

$$f(z) = \sum_{j=0}^{\infty} f_j z^j,$$

then

$$\begin{aligned} CfC(z) &= f_\star(z) = \overline{f(\overline{z})} \\ &= \sum_{j=0}^{\infty} \overline{f_j} z^j. \end{aligned}$$

(In the above, bars represent complex conjugation.) This means that for an analytic map f, the map CfC is also analytic.

Remark 18 (Complex Rescaling Factor)
The above definition of the operator differs from that used by Widom [Wid83]: Recall that he considered a coordinate system rotated through an angle θ which enabled the use of a real rescaling factor. As introduced earlier, the value of θ is the asymptotic angle (measured from the critical point) between successive Fibonacci iterates. As such, it is not rigorously known, a priori. It is computationally more convenient to consider a complex rescaling together with a conjugation of the plane.

(In addition, in order to establish the existence of a locally unique fixed point of this operator it is necessary to ensure that the Jacobian of a certain operator, namely a variant of Newton's method, has no unit eigenvalues. By fixing the *modulus* of $F(0)$, we remove one such eigenvalue corresponding to a dilation symmetry. However, a line of fixed points to the operator still remains, corresponding to a rotation symmetry of the scaling pictures. We remove this eigenvalue by also fixing the *argument* of $F(0)$. It is also necessary to take a projection to a subspace of commuting maps in order to establish local uniqueness. The details are presented in the next chapter.)

The second form of the operator uses the composition order EF, which will be called the *accretive* composition order, for reasons that will become clear in Chapter 4 where this version of the operator is used to characterise the Siegel disc boundary.

Definition 19 (Renormalization Operator — Accretive Order)
The ('accretive-order') Siegel disc renormalization operator *has, as its domain of definition, the set of pairs* (E, F) *such that*

$$\alpha C(D_E) \subset D_F \tag{1.28}$$

$$\alpha C(D_F) \subset D_E \tag{1.29}$$

$$F(\alpha C(D_F)) \subset D_E. \tag{1.30}$$

The operator acts on pairs (E, F) *by means of* $(E, F) \mapsto (\tilde{E}, \tilde{F})$, *where*

$$\tilde{E}(z) = C\alpha^{-1} F \alpha C(z), \tag{1.31}$$

$$\tilde{F}(z) = C\alpha^{-1} E F \alpha C(z). \tag{1.32}$$

(A fixed point for this operator is obtained from the fixed point for the 'inverse' operator by taking an *analytic continuation* of the maps. This is done in chapter 4.)

Remark 20 (Normalization Condition)
Notice that, for both composition orders, the form of the operator means that the definition $\alpha = F(0)$ *enforces the* normalization condition

$$E(0) = 1. \tag{1.33}$$

It is straightforward to verify that the operators given above have a *simple fixed point* given by the (commuting) pair of maps

$$E_s : z \mapsto z + 1, \tag{1.34}$$

$$F_s : z \mapsto z - \omega. \tag{1.35}$$

(Thus $\alpha_s := F_s(0) = -\omega$.) This fixed point accounts for the simple scaling scenario observed within the Siegel disc. It turns out [Sti93a] that when viewed in a suitable space of functions the fixed point is hyperbolic, having a one-dimensional unstable manifold whose direction corresponds to a change in rotation number, and a codimension-one stable manifold. The pair corresponding to rigid rotation is contained in the stable manifold, and changing rotation number results in a curve that intersects the stable manifold transversally.

If we consider a map f which has a golden mean Siegel disc and transform the coordinate system to give a new map f', in such a way that the origin is an interior point of the Siegel disc for f', then it follows that the pair of maps $(f', \mathbf{I})$ (where $\mathbf{I}$ indicates the identity map) is attracted to the simple renormalization fixed point and the corresponding simple asymptotic scaling scenario will be observed for the original map f.

The main result proved by Stirnemann in [Sti93b] is that the conjectured critical fixed point also exists:

Theorem 21
Widom's operator has a fixed point defined by a commuting pair (E, F) *of even analytic maps, defined on domains* D_E *and* D_F *respectively. The fixed point is locally unique in the subspace of commuting pairs.*

We shall refer to this pair (E, F) as *Widom's fixed point.* Thus E and F satisfy the *fixed point equation*:

$$E(z) = C\alpha^{-1}F\alpha C(z) \qquad \text{on } D_E, \tag{1.36a}$$

$$F(z) = C\alpha^{-1}FE\alpha C(z) \qquad \text{on } D_F. \tag{1.36b}$$

Remark 22 (Pairs of fixed points related by $\star$)
There is actually a pair *of critical fixed points, related to each other by the* $\star$ *operator that sends* $f(z)$ *to* $CfC(z)$.

The next chapter indicates how the proof for the existence of the critical fixed point was carried out. It also demonstrates how a ball around the fixed point in the space of functions is obtained, which will be used in the rest of this thesis.

Chapter 2

Existence of the Renormalization Fixed Point

2.1 Introduction

This chapter describes a computer-assisted proof, due to Stirnemann [Sti93b], that establishes the existence of the critical fixed point for the Siegel disc renormalization operator with the inverse composition order. Rigorous bounds on the fixed point are required for the results presented in this thesis (intuitively, we need to know the centre and radius of a ball of functions that is guaranteed to contain the true fixed point). For this reason (and in order to verify the earlier proof) the computer-assisted portion of the proof was repeated. This involved adapting the programs to a different machine architecture and programming language compiler, which required alterations to the code (the core routines of the rigorous code are *hardware specific*, see chapter 3).

The fixed point thus obtained was then used to prove the results presented in subsequent chapters. (In particular, it will be used to verify the so-called *necklace hypotheses* in Chapter 4, and to estimate the fractal dimension of Siegel disc boundaries in Chapter 5.) This chapter and the next one also serve to introduce several important concepts (in particular, the rigorous computer framework and how it relates to the corresponding space of functions) that will be used later.

2.1.1 Results

Recall the definition of the Siegel disc operator with the inverse composition order, $\mathsf{N} : (\mathsf{E}, \mathsf{F}) \mapsto (\tilde{\mathsf{E}}, \tilde{\mathsf{F}})$ where

$$\tilde{\mathsf{E}}(z) = \mathsf{C}\alpha^{-1}\mathsf{F}\alpha\mathsf{C}(z) \quad \text{on } \mathsf{D}_{\mathsf{E}}, \tag{2.1a}$$

$$\tilde{\mathsf{F}}(z) = \mathsf{C}\alpha^{-1}\mathsf{F}\mathsf{E}\alpha\mathsf{C}(z) \quad \text{on } \mathsf{D}_{\mathsf{F}}, \tag{2.1b}$$

where D_{E} and D_{F} are the domains of E and F, and $\alpha = \mathsf{F}(0)$.

Remark 23 (Use of the Inverse Composition Order)
The existence proof uses the l_1-estimate for the range of values of a function on its domain.

There do *exist domains such that this property is satisfied for the fixed point with the* accretive *composition order (*$\mathsf{F} = \mathsf{C}\alpha^{-1}\mathsf{E}\mathsf{F}\alpha\mathsf{C}$*). However, the* l_1*-estimate obtained for that case is not tight enough to give a rigorous proof. Hence the use of the inverse composition order here. Later, in chapter 4, a fixed point will be needed that satisfies the other (accretive) composition order. Existence of this fixed point will be deduced from the inverse one by taking analytic continuations of the maps.*

The existence proof established the following theorems:

Theorem 24 (Existence of the Critical Fixed Point)
The Siegel disc operator has a commuting fixed point, given by a pair of even functions. The rescaling factor α *satisfies*

$$\Re\alpha \in [-0.220265974,\ -0.220265961], \tag{2.2a}$$

$$\Im\alpha \in [-0.708481721,\ -0.708481708], \tag{2.2b}$$

$$|\alpha| \in [\ 0.741932224,\ \ 0.741932240]. \tag{2.2c}$$

(Note: compare the bounds on the interval containing the magnitude of α with the value of the real rescaling observed previously in section 1.4.)

Theorem 25 (Family of Non-commuting Fixed Points)
There exists a smooth real one-parameter family of fixed points through the commuting fixed point. The elements of this family are even and commute to zeroth order, but (with the obvious exception of the commuting fixed point) do not *commute strictly.*

(The definition of zeroth-order commutativity will be given in Section 2.5.1.)

The strategy used here was first developed by Lanford in [Lan82], and subsequently used to solve other fixed point problems of this kind (for example, [EKW84], [EW87], and [Mes85]).

2.2 Overview of proof

The aim of this section is to give an overview of the proof of Theorem 24.

The basic strategy (c.f. [Lan82]) is to demonstrate that a variant of Newton's method for the fixed point problem is contractive on a small ball around a good approximate fixed point in a suitable space of (pairs of) functions. The contraction mapping theorem then yields the existence of a true fixed point within the ball. Section 2.3 introduces the method.

The first step (Section 2.4) is to define a suitable space in which to work, in this case the problem is reduced to a space of pairs of *even* functions. The functions in each pair are expanded as power series defined on suitable domains. The method used for finding an approximate fixed point, along with corresponding domains, is given in Section 2.6.

The presence of the one-parameter family of non-commuting fixed points (see theorem 25, above) introduces a difficulty in applying Newton's method to the problem: in the space

of pairs of even functions, the (commuting) critical fixed point is not locally unique. This prevents convergence of the standard Newton method. Section 2.5 shows how this problem is overcome by preceding the renormalization operator (denoted N) by a smooth non-linear projection P to a space of *approximately commuting pairs*.

Section 2.7 then shows how the existence of a fixed point for the modified operator NP is proved with the assistance of rigorous computer estimates. (The computational framework used to obtain these estimates is discussed in greater detail in the next chapter.) The existence of a fixed point for the modified operator NP finally implies the existence of a fixed point for the original operator N, proving Theorem 24.

Given the existence of a critical fixed point, Stirnemann [Sti93b] also proved the existence of the real one-parameter family of non-commuting fixed points, using the so-called *invariant modulus* [MO89]. No details are given here, since they will not be needed in subsequent chapters (we only need the critical fixed point). The interested reader is referred to [Sti93b].

2.3 Newton's method for fixed point theorems

2.3.1 The Contraction Mapping Principle

A conceptually simple way of establishing the existence of a fixed point for a mapping is to use the Contraction Mapping Principle (arguably one of the most powerful results in mathematics), as expressed in the following theorem.

Theorem 26 (Contraction Mapping Principle)
Let $\mathcal{X}$ be a Banach space, *i.e. a complete normed space (with norm $\|\cdot\|$), and let f be a mapping of a set $\mathsf{X} \subseteq \mathcal{X}$ to itself such that, for some $\kappa < 1$,*

$$\|\mathsf{f}(x_1) - \mathsf{f}(x_2)\| \leq \kappa \|x_1 - x_2\| \quad \textit{for all } x_1, x_2 \textit{ in } \mathsf{X}, \tag{2.3}$$

then f has exactly one fixed point $x_* \in \mathsf{X}$ and, for any $x_0 \in \mathsf{X}$,

$$\mathsf{f}^n(x_0) \to x_* \quad \textit{as } n \to \infty. \tag{2.4}$$

A mapping f with the properties stated above is called a *contraction mapping* or, simply, a *contraction*. (See, for example, [Bar88].)

The theorem also yields an estimate for the distance of a given point from the fixed point. To see this, note that

$$\left\| \mathsf{f}^{n+1}(x_0) - \mathsf{f}^n(x_0) \right\| \leq \kappa^n \|\mathsf{f}(x_0) - x_0\|, \tag{2.5}$$

from which it may be shown that

$$\|x_0 - x_*\| \leq (1 - \kappa)^{-1} \|x_0 - \mathsf{f}(x_0)\|. \tag{2.6}$$

Remark 27
Furthermore, it may be shown [Lan91] that the fixed point varies continuously with the map. *This fact is vital to the success of computer-assisted proofs, which must necessarily deal with finite-dimensional versions (truncations) of operators.*

For a differentiable mapping on a *convex* set, contractivity may be verified by estimating the *norm of the derivative*:

Remark 28
Let x_1, x_2 be two points of a Banach space X, and let $[x_1, x_2]$ represent the straight line segment joining them. Let $f : X \to X$ be defined and differentiable on $[x_1, x_2]$. Then

$$\|f(x_1) - f(x_2)\| \leq \sup_{x \in [x_1, x_2]} \{\|Df(x)\|\} \cdot \|x_1 - x_2\|. \tag{2.7}$$

To establish the existence of a fixed point for Siegel disc renormalisation, we work in a Banach space of analytic functions (the full details will be given in Section 2.4). In order to apply the contraction mapping principle to a differentiable mapping on a ball in a Banach space, it is not enough to know that the norm of the derivative is less than unity in the ball; it is also required that the ball is mapped into itself. The following lemma demonstrates how this can be checked. (Basically, we need to know how far the approximate fixed point moves and by how much the ball around it shrinks.)

Lemma 29 (Contraction Mapping on a Ball)
Let $f : X \to X$ be differentiable on the (closed) ball $\overline{B_\rho}(x_0)$ of radius ρ about x_0 in the Banach space X, and suppose that there exists a constant $\kappa < 1$ such that

$$\|Df(x)\| \leq \kappa \quad \textit{for all } x \in \overline{B_\rho}(x_0). \tag{2.8}$$

Then, provided that

$$\|f(x_0) - x_0\| \leq (1 - \kappa)\rho, \tag{2.9}$$

f has exactly one fixed point in $\overline{B_\rho}(x_0)$.

Proof: We already have that f is a contraction for points in the ball, by the bound on the derivative. It remains to show that f maps the ball into itself. Let x be any point in the ball $\overline{B_\rho}(x_0)$. By definition

$$\|x - x_0\| \leq \rho,$$

and we have

$$\begin{aligned} \|f(x) - x_0\| &\leq \|f(x) - f(x_0)\| + \|f(x_0) - x_0\| \\ &\leq \kappa\rho + (1 - \kappa)\rho \\ &= \rho. \end{aligned}$$

□

This is illustrated graphically by figure 2.1.

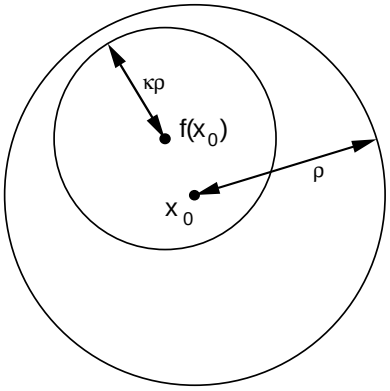

Figure 2.1: Proving the existence of a fixed point in a ball.

The contraction mapping theorem has the disadvantage that it only works for contractions: it cannot be used *directly* to establish the existence of fixed points for operators that are not contractive. However, in combination with Newton's method (which is ordinarily used for finding roots of functions), it may be successfully applied to fixed point problems of the kind faced here.

2.3.2 Newton's Method

Suppose that $f : A \to Y$, where A is an open subset of a Banach space X, and Y is a Banach space. (It is further needed that f is of class C^2 and that $Df(x)$ is invertible for each x.) The idea of Newton's method is to find a zero of f (i.e. a point x_* with $f(x_*) = 0$) by starting with an approximate zero of f, say x_0, and improving the estimate by iteration. Suppose that x_0 is close to a zero. Then for x close to x_0, $f(x)$ may be approximated by

$$f(x_0) + Df(x_0)(x - x_0). \tag{2.10}$$

Notice that this quantity has a zero at

$$x_1 = x_0 - (Df(x_0))^{-1} f(x_0), \tag{2.11}$$

Iteration proceeds by replacing the original guess x_0 with x_1 (in the hope that the latter will be a better approximation to the true zero). In other words, iteration is used to attempt to find a fixed point of the operator

$$\phi_{\text{zero}} : x \mapsto x - (Df(x))^{-1} f(x). \tag{2.12}$$

In order to see that the iteration will converge for x_0 close enough to x_*, observe that the derivative of ϕ_{zero} is given by

$$\begin{aligned} D\phi_{\text{zero}}(x) &= \mathbf{I} - D((Df(x))^{-1}) f(x) + (Df(x))^{-1} Df(x) \\ &= -D((Df(x))^{-1}) f(x), \end{aligned}$$

(where $\mathbf{I}$ is the identity) and that this is zero at x_*. By continuity there must be a small ball around the fixed point on which ϕ_{zero} is a contraction. Provided that the initial guess

x_0 is inside this ball, we are entitled to the hope that the iteration will converge to the fixed point.

In order to find *fixed points* of f (rather than zeros), notice that these correspond to zeros of $f(x) - x$, and the method is adapted to the following

$$\phi_{\text{fixed}} : x \mapsto x - (Df(x) - \mathbf{I})^{-1}(f(x) - x), \tag{2.13}$$

This method is inconvenient to apply computationally, because it requires a bound on the derivative of f. Instead, the following simplified method is typically used

$$\Phi : x \mapsto x + (\mathbf{I} - \Gamma)^{-1}(f(x) - x), \tag{2.14}$$

where Γ is a constant linear operator. Provided that Γ is a sufficiently good approximation to $Df(x)$ for $x \in \overline{B_\rho}(x_0)$, then this simplified operator should be a contraction on the ball and the Contraction Mapping Principle may be applied.

The above idea was used to establish the existence of the critical Siegel disc fixed point. The next section describes a suitable space in functions in which to work.

2.4 The Function Space

Empirically the critical fixed point is found to be a pair of *even* maps. It is therefore reasonable to work in a space of pairs of even maps. This was built into the proof by writing

$$E(z) = U(z^2) = UQ(z) \tag{2.15a}$$
$$F(z) = V(z^2) = VQ(z). \tag{2.15b}$$

Here, and in the following, Q denotes squaring, i.e. $Q : z \mapsto z^2$.

In terms of U and V, the operator looks as follows:

$$\tilde{U}(z) = C\frac{1}{\alpha}V\alpha^2C(z), \tag{2.16a}$$
$$\tilde{V}(z) = C\frac{1}{\alpha}VQU\alpha^2C(z), \tag{2.16b}$$

where $\alpha = V(0)$. In what follows, the operator will be denoted N. (Note that the above equation is not actually a fixed point equation for either (UQ, VQ), or (QU, QV), although such fixed point equations are readily deduced from it by squaring.

Remark 30 (Uncles of Fixed Point Equations)
In fact, equation 2.16 will be called an uncle *of the fixed point equation — a terminology borrowed from Hofstadter [Hof87]. This distinction will become more important in Chapter 5, where a number of different fixed-point equations are deduced from the existence of the corresponding uncles, for now it is merely a curiosity.*

For convenience, the *domains* of U and V are taken to be discs, denoted D_U and D_V respectively. The centre and radius of D_U (respectively of D_V) will be denoted c_U and r_U

(respectively c_V and r_V). (The 'original' domains D_E and D_F are then found by 'taking the square root' of these domains, in a suitable sense.)

U and V are then expanded as power series of the form

$$U(z) = \sum_{j=0}^{\infty} U_j \left(\frac{z - c_U}{r_U}\right)^j, \tag{2.17a}$$

$$V(z) = \sum_{j=0}^{\infty} V_j \left(\frac{z - c_V}{r_V}\right)^j, \tag{2.17b}$$

and the (complex) coefficients U_j and V_j are used as coordinates in the space of pairs (U, V).

Since the choice of the domains for the above expansion determines the underlying Banach space, it is crucial to the success of the proof. In particular, the computer proof necessarily deals with a *truncated* (i.e. finite-dimensional) version of the operator. The accuracy of this truncation is determined by the domains chosen. (The procedure that was used for finding the domains is summarised in Section 2.6.)

The space is normed by

$$\|(U, V)\| = \sum_{j=0}^{\infty} (|U_j| + |V_j|), \tag{2.18}$$

which corresponds to the sum of the ℓ_1 norms of U and V.

The presence of complex conjugation in the operator means that the derivative of N is only real-linear (in fact, complex *anti*-linear, i.e. it is the composition of a complex-linear map with complex conjugation). As such, it must be considered as an operator on the *realification* of the space. (The space $\mathbb{C}^n$, for example, has the realification $\mathbb{R}^{2n}$.) When dealing with the derivative, it is therefore more convenient to work with the alternative norm

$$\|(U, V)\| = \sum_{j=0}^{\infty} (|\Re U_j| + |\Im U_j| + |\Re V_j| + |\Im V_j|). \tag{2.19}$$

In fact, matters are somewhat more complicated than indicated above: the presence of the family of non-commuting fixed points mentioned in Theorem 25 creates an additional difficulty which is solved by taking a projection to a space of *approximately commuting pairs*. The details of this projection will be given in the next section.

2.5 Projection to the Commuting Subspace

The presence of the one-parameter family of non-commuting fixed points mentioned in Theorem 25 implies that the (commuting) critical fixed point is not locally unique. Correspondingly, the derivative of the renormalization operator is known (empirically) to have an indifferent (unit) eigenvalue at the critical fixed point, which means that the usual implementation of Newton's method cannot be expected to converge to it.

It turns out that the critical fixed point *is* locally unique in a subspace of *approximately commuting* pairs, which will be defined below. By preceding the renormalization operator N by a smooth non-linear projection P to this subspace, Newton's method can be made to converge successfully.

2.5.1 Approximately Commuting Pairs

The concept of approximate commutativity is now introduced. Commutativity of a pair (E,F) means that $\mathsf{EF}(z) = \mathsf{FE}(z)$. In the notation adopted for the space of even functions this means that $\mathsf{UQVQ}(z) = \mathsf{VQUQ}(z)$, i.e. $\mathsf{UQV}(z) = \mathsf{VQU}(z)$. This motivates the following definition.

Definition 31 (Commutator)
The commutator $\mathcal{C}(\mathsf{U},\mathsf{V})$ *of a pair* (U,V) *is the function*

$$\mathcal{C}(\mathsf{U},\mathsf{V})(z) = \mathsf{UQV}(z) - \mathsf{VQU}(z), \tag{2.20}$$

defined wherever the right hand side makes sense.

Exact commutativity appears to be difficult to implement in a rigorous computational framework. By restricting the form of the commutator for pairs that do not commute exactly, the useful concept of *approximate* commutativity is introduced.

Definition 32 (Approximate Commutativity)
Let m *be an integer greater than or equal to* 0. *A pair* (U,V) *is said to* commute approximately to order m *if the commutator* $\mathcal{C}(\mathsf{U},\mathsf{V})$ *is analytic at the origin and has the form*

$$\mathcal{C}(\mathsf{U},\mathsf{V})(z) = bz^{m+1} + O(z^{m+2}), \tag{2.21}$$

where b *may take the value zero.*

The computer assisted portion of the proof verifies that the renormalization operator is well defined in an open set B in the space of pairs, within which the commutator $\mathcal{C}(\mathsf{U},\mathsf{V})$ is well defined and analytic in a neighbourhood of the origin. Furthermore, it is shown that the pairs (U,V) in B satisfy the condition

$$\mathsf{U}'(0) \neq 0, \quad \mathsf{V}'(0) \neq 0, \tag{2.22}$$

and that the absolute value of the rescaling factor $|\alpha| = |\mathsf{V}(0)|$ is different from 1.

The above properties have two important consequences:

- The renormalization operator preserves approximate commutativity: if a pair (U,V) commutes to order m, then the renormalized pair $(\tilde{\mathsf{U}},\tilde{\mathsf{V}})$ also commutes to order m.
- Any fixed point of the Siegel disc operator that commutes approximately to order $m \geq 1$ automatically commutes strictly, i.e. the commutator is zero in a neighbourhood of the origin.

These facts are proved for the case of even maps in [Sti93b] (Lemma L1 and L2).

2.5.2 The Projection

The definition of the projection to the commuting subspace is now given. How existence of a commuting fixed point of N itself can be deduced from existence of a fixed point of NP is shown in full in [Sti93b].

The range of the projection is the space of pairs approximately commuting to first order. For a pair (U, V), this means that the following conditions hold

$$\begin{aligned} UQV(0) &= VQU(0), \\ (UQV)'(0) &= (VQU)'(0). \end{aligned}$$

For the sake of numerical simplicity, it will be assumed that the pair $p = (U, V)$ is normalized (i.e. that $U(0) = 1$). This is justified because at the fixed point $p = NP(p)$, and N enforces normalization (its output pairs are always normalized).

The idea is to enforce the above commutativity conditions by adding a polynomial to V as follows:

$$V \mapsto V + b_1 z^2 + b_2 z^3. \tag{2.23}$$

(Where the values of b_1 and b_2 are to be determined.) Note that the normalization condition $U(0) = 1$, the value $\alpha = V(0)$, and the value $V'(0)$ are all left invariant under this substitution. The result of the above substitution is a linear system for b_1 and b_2, the corresponding matrix having integer coefficients which makes the projection convenient to implement and leads to small roundoff error in the computer calculations. (Further details were given in [Sti93b].)

The projection P thus defined is well-defined and smooth (i.e. C^∞) in the ball B of functions used in the proof, and has the following properties: for any pair p in the domain of P, then $P(p)$ approximately commutes to first order; and for any pair p in the domain of P which approximately commutes to first order, then $P(p) = p$.

The projection requires evaluation of the derivative of U at the points 0 and α^2 and of the derivative of V at the points 0 and 1. For the domains that are used in the proof presented here, it turns out that 0 is close to the boundary of D_U and that both 0 and 1 are close to the boundary of D_V. This leads to a large truncation error, which means that a comparatively high truncation degree (160) has to be used. (Compare with the critical circle maps fixed point [Mes85], where 80 was enough.)

2.6 Finding Suitable Domains

The aim of this section is to describe the method by which the domains of U and V used in the existence proof were found, along with a suitable approximate fixed point p. Since the maps will be expanded on these domains, the choice for the domains is crucial.

Remark 33 (Domains for the Accretive Fixed Point)
Recall that the existence proof uses the inverse composition order (FE*, equivalently* VQU*). A similar method to the one presented here was also used to find domains for the fixed*

point with the accretive composition order (EF, equivalently UQV), which will be used in Chapter 4 and Chapter 5.

The process of finding good domains, along with a good approximate fixed point, is ultimately a question of guesswork, guided by experiment. There are many parameters that may be varied, making it difficult to come up with a general algorithm.

The first step is to find a *numerical* approximate fixed point. This can be done by iterating a truncated form of the renormalization operator a few times on a simple pair that is known (empirically) to be attracted. For example, the pair (f, f) where f is the quadratic map given by

$$f(z) = \frac{\lambda}{2}z^2 + 1 - \frac{\lambda}{2}, \tag{2.24}$$

where $\lambda = e^{2\pi i\omega}$ and ω is the golden mean. (Note that this pair automatically commutes.)

In order that the operator is well-defined (see, for example, equation 2.16), the following equations will be required to hold for all pairs (U, V) in a ball around an approximate fixed point:

$$\alpha^2 C\overline{D_U} \subset D_V, \tag{2.25a}$$
$$\alpha^2 C\overline{D_V} \subset D_U, \tag{2.25b}$$
$$QU\alpha^2 C\overline{D_V} \subset D_V. \tag{2.25c}$$

(In the above, and throughout this thesis, bars denote the topological closure. Complex conjugation will always be indicated by C unless otherwise stated.) These equations are called the *domain extension conditions*. The importance of domain extension is indicated in [Mac82] (section 2.3.8), where it is referred to as the *analyticity improving* property: the linearization of a renormalization operator which satisfies the domain extension conditions on complex domains (at at fixed point) will be a *compact operator*.

Remark 34 (Compact Operators)
Compactness is desirable, since compact operators behave somewhat like finite-dimensional ones (see, for example, [Rud82]). In particular, compactness (along with the fact that the fixed point of a contraction mapping varies continuously with the map) plays an important role as it implies that fixed points of truncated versions of an operator (with high enough truncation degree) and their spectra will be close enough to those of a fixed point of the true operator. If this were not the case, then a computer assisted proof (in which one necessarily deals with truncated operators and functions) would be hopeless.

In order to find suitable domains, a method which was applied to circle maps in [Mes85] was used. Firstly, the truncation degree (i.e. the degree of the polynomials used to approximate the maps) is fixed, (at, say, 20). Then rough initial domains are found graphically, by trial and error, using the corresponding numerical fixed point. These domains are then used to determine both the (truncated form of) the operator, and the space in which it lives.

Given the initial domains, and an approximate fixed point, an iterative procedure is used to improve them. Firstly, the numerical fixed point is expanded on the chosen domains. This

involves expressing the maps as polynomials in $(z - c_U)/r_U$ and $(z - c_V)/r_V$, respectively (as indicated in equation 2.17). A better approximation to the fixed point of the truncated operator is then found by using Newton's method.

Remark 35 (Use of the Projection)
Recall that the derivative of the renormalization operator is known to have a unit eigenvalue. This means that (for high enough truncation degree) the derivative of the truncated operator will have an eigenvalue close to unity, which would cause Newton's method to break down. For this reason the renormalisation operator is preceded by the projection defined in the previous section when using Newton's method to improve the approximate fixed point.

Given an improved approximate fixed point, the next step is to find improved domains. Newton's method may then be used once more to obtain an approximate fixed point defined on the new domains, and the procedure can be repeated if necessary. The truncation degree is typically increased after a few iterations.

Given an approximate fixed point, the idea is to use the computer to estimate certain quantities and adjust the domains accordingly. The most important quantities are those affecting the domain extension conditions (equation 2.25). The norm of the inverse matrix involved in Newton's method and the decay rate for the size of the coefficients of the polynomials may also be considered. The trade-offs between the importance of the various quantities are complicated, and balancing them usually involves a great deal of trial and error.

For example, assume that the pair (U, V) is close to the renormalization fixed point. The third domain extension condition (Equation 2.25c) indicates that the function $z \mapsto QU\alpha^2 C(z)$ should map the closure of D_V into its interior and, therefore, that it should have a fixed point there. Notice that V is then composed with this function in the fixed point equation

$$V = C\alpha^{-1}VQU\alpha^2 C \quad \text{on } D_V, \tag{2.26}$$

and the result must be expanded around the centre of D_V. By choosing the centre c_V of D_V to be exactly this fixed point, the truncation error in the above composition is minimised. In particular, this corresponds to composing V with a power series having *zero constant term.* (This consideration is discussed more fully in Section 3.5.2, which describes the relevant part of the computer implementation of the proof.) Thus the parameter c_V can be fixed. Similar considerations may be used to fix c_U.

With the new centres for D_U and D_V fixed, it remains to find suitable radii. In order to satisfy the second domain extension condition (Equation 2.25b), the radius r_U will be chosen so that D_U contains the (concentric) disc $\alpha^2 C D_V$, with room to spare. The remaining degree of freedom, the radius r_V, may then be chosen so that the other domain extension conditions (2.25c) and (2.25a) are also satisfied.

j	$\Re U_j$	$\Im U_j$
0	$7.57550776023832140 \cdot 10^{-01}$	$8.11605469289232578 \cdot 10^{-02}$
1	$2.68059524988865039 \cdot 10^{-01}$	$1.78509805751320472 \cdot 10^{-01}$
2	$3.39653220162117048 \cdot 10^{-02}$	$-4.16254831382228477 \cdot 10^{-02}$
3	$-7.19318178151169465 \cdot 10^{-03}$	$-9.28624696084178608 \cdot 10^{-03}$
4	$-2.35070094130932416 \cdot 10^{-03}$	$1.03193644752552316 \cdot 10^{-03}$
5	$9.58731545862152699 \cdot 10^{-05}$	$5.47952840191413717 \cdot 10^{-04}$
6	$1.20531652113259281 \cdot 10^{-04}$	$7.27726890033071134 \cdot 10^{-06}$
7	$7.56547167155382664 \cdot 10^{-06}$	$-2.52759993215326129 \cdot 10^{-05}$
8	$-5.02709387087134926 \cdot 10^{-06}$	$-2.89057468903472061 \cdot 10^{-06}$
9	$-8.76121094939487690 \cdot 10^{-07}$	$9.31782283057971695 \cdot 10^{-07}$
⋮	⋮	⋮

j	$\Re V_j$	$\Im V_j$
0	$2.78608590300148426 \cdot 10^{-11}$	$4.61196766688199912 \cdot 10^{-11}$
1	$8.39214060501707215 \cdot 10^{-01}$	$5.60684927121525956 \cdot 10^{-01}$
2	$2.87126352294469311 \cdot 10^{-01}$	$-6.05778383746779059 \cdot 10^{-01}$
3	$-5.11645411029023056 \cdot 10^{-01}$	$-1.90279663857481185 \cdot 10^{-01}$
4	$-1.17152962603158384 \cdot 10^{-01}$	$4.29977350992223339 \cdot 10^{-01}$
5	$3.59490075590482128 \cdot 10^{-01}$	$6.15974910136376300 \cdot 10^{-02}$
6	$2.16862203742139402 \cdot 10^{-02}$	$-2.99369957116730423 \cdot 10^{-01}$
7	$-2.48012417686435865 \cdot 10^{-01}$	$6.00587464417352227 \cdot 10^{-03}$
8	$2.46452281568751985 \cdot 10^{-02}$	$2.04140693796179329 \cdot 10^{-01}$
9	$1.66790746839407172 \cdot 10^{-01}$	$-3.65950060805623617 \cdot 10^{-02}$
⋮	⋮	⋮

Table 2.1: Coefficients of U (expanded on D_U) and V (expanded on D_V).

The following values were used in the proof

$$c_U = -0.2188497414079558 + i0.2328147240271490, \tag{2.27a}$$
$$r_U = 0.3640985354093064, \tag{2.27b}$$
$$c_V = 0.5672961438978619 + i0.1229664702397770, \tag{2.27c}$$
$$r_V = 0.636. \tag{2.27d}$$

The first few (complex) coefficients in the expansion for the maps of the approximate fixed point are given in the following tables (2.1 and 2.1). (A truncation degree of 160 was used in the actual proof.)

Notice, in particular, that the constant term of V (in table 2.1) is small for the expansion on the chosen domains. This makes for better numerics during the composition VQU.

2.7 Implementing Newton's Method

Given an approximate fixed point p (in this case, a pair of polynomials of degree 160) expanded on suitable domains, it remains to implement the modified version of Newton's method for N under the projection P. (The calculations described below are carried out by computer using a rigorous framework that will be described in Chapter 3.)

The modified form of Newton's method for NP looks as follows (c.f. equation 2.14):

$$\Phi : p \mapsto p + (\mathbf{I} - \Gamma)^{-1}(NP(p) - p). \tag{2.28}$$

Recall that Γ is to be a fixed linear operator that approximates $D(NP)(p)$. Since the basis of the space is fixed (it is induced by the choice of domains), Γ is simply a matrix. (It is chosen by calculating the images of the monomials $w^j = (\frac{z-c}{r})^j$ under NP.)

The computer program then proves that $\mathbf{I} - \Gamma$ is invertible. This is done by computing the inverse matrix $(\mathbf{I} - \Gamma)^{-1}$ numerically by Gaussian elimination, and then applying Hansen's method [Moo66], which entails adding a small symmetrical interval around each of the entries in the computed inverse matrix to produce a matrix (of intervals) Δ which is guaranteed to contain $(\mathbf{I} - \Gamma)^{-1}$. (This is done by ensuring that $\Delta(\mathbf{I} - \Gamma)$ contains the identity matrix.) A more thorough treatment of interval arithmetic and the rigorous computer framework is given in Chapter 3.

Provided that both the approximate fixed point p and the matrix Γ are good enough approximations, it is possible to find a small ball $B_\rho(p)$ around p (specified by the radius ρ), on which the operator NP is well defined and differentiable. (For the proof given here, the radius of B was chosen to be $6 \cdot 10^{-9}$.) This is confirmed by rigorous computer calculations, which verify that the *domain extension conditions* (equation 2.25) hold in the ball B. (In fact, the program verifies domain extension as a by-product of computing the renormalization of the approximate fixed point.)

Domain extension in B implies that the operator is *well defined* in B, that it is *differentiable* there, and that its derivative $D(NP)$ is a *compact* operator (see [Mac82, §2.3.8]).

Further calculations give a rigorous upper bound κ for the norm of the derivative of Φ in the ball. This is done by evaluating $D\Phi = \Delta(DN - \Gamma)$ on all base vectors (monomials) to produce a matrix of column vectors, and then calculating the (maximum column) norm.

Finally, the program calculates an upper bound for the quantity

$$\|\Phi(p) - p\| = \|\Delta(NP(p) - p)\|,$$

which is known as the *residue*.

Remark 36 (Real and Complex Norms)
The program works out the real l_1*-norm by taking the realification of* $NP(p) - p$. *Note that since* $|x| + |y| \leq \sqrt{2}|z|$ *for any complex number* $z = x + iy$, *a real bound may be produced from a complex one by multiplication by* $\sqrt{2}$. *Also, the triangle inequality* $|z| \leq |x| + |y|$ *demonstrates that the real* l_1*-norm is stronger than the complex* l_1*-norm. This is convenient because it allows the program to work with a complex ball around the approximate fixed point, which is easier to implement in the rigorous framework (see Chapter 3 for more details).*

The above bounds are then used to confirm the crucial inequality

$$\|\Phi(p) - p\| \leq \rho(1 - \kappa), \tag{2.29}$$

(recall Lemma 29) which proves that there is a unique fixed point of Φ inside the ball. The success of this part of the proof depends on a careful balance between the choice of domains, the quality of the approximate fixed point (and the corresponding matrix Γ), and the radius of the ball.

Existence of a fixed point for Φ implies the existence of a fixed point for NP. (Notice that both Φ and NP share the same fixed points.) It follows (see [Sti93b, §4]) that the fixed point of NP is also a fixed point of the original operator N and that it commutes strictly, completing the existence proof.

The following program output illustrates the proof (an approximate fixed point, and corresponding matrices, had already been computed and are read from files):

Wed Apr 10 20:09:44 BST 1996
  Reading approximate fixed point p...done.
  Reading matrix Γ...done.
  Reading matrix Δ...done.
  Setting error bound, $\rho = 6.00000 \cdot 10^{-09}$.
  Estimating residue...
    Upper bound on residue, $\|\Phi(p) - p\|$, is $1.78563 \cdot 10^{-09}$.
  Renormalizing approximate fixed point...
    Operator is well defined.
  Estimating derivative of Newton's method...
    Upper bound on κ is $6.67132 \cdot 10^{-01}$.
  Upper bound on residue $\leq$ lower bound on $\rho(1 - \kappa)$.
Proof succeeded.
Thu Apr 11 08:54:59 BST 1996

The next chapter discusses in more detail the rigorous computer framework that was used for this proof. This framework will then be extended and used (along with the fixed point obtained here) in subsequent chapters to establish new results.

Chapter 3

Computer Assisted Proofs

3.1 Introduction

As demonstrated in the previous chapter, the type of computer assisted proof used here is composed of two parts: (1) a traditional (conceptual) proof of the theorems in question, typically involving some assumptions on the numerical values taken by certain quantities, and (2) a computer program written with the aim of verifying those assumptions by performing rigorous numerics. For example, the existence proof of the previous chapter relies on the assumption that a particular inequality holds (namely Equation 2.8 in Lemma 29) in order to prove that a variant of Newton's method is a contraction. The aim of this chapter is to demonstrate how such assumptions may be verified *rigorously*, by using a computer.

Rigorous computer calculations of this kind have been used in a variety of problems, including:

- Proofs concerning the Feigenbaum conjectures on the period-doubling scenario (Lanford [Lan82], Eckmann and Wittwer [EW85] and [EW87], and Koch, Schlegel and Wittwer [KSW]).
- Existence of solutions to functional equations arising in area-preserving maps (Eckmann, Koch, and Wittwer [EKW84]).
- The non-existence of invariant circles for area-preserving maps (MacKay and Percival [MP85]).
- The stability of a semi-relativistic quantum mechanical model for matter (Fefferman and de la Llave [FL85]).
- The existence of stable quasi-periodic motions for a three-body model of the Sun-Jupiter-Ceres system (Celletti and Chierchia [CC96]).

(Other references of interest are [EKW82], [EMO91], and [KW86].)

The following sections serve to describe the (analytical and computer) framework behind proofs of this nature. The final section (3.6) contains a few philosophical musings on the

validity of these proofs.

3.2 Floating Point Arithmetic and Rounding

It might seem that the use of computers for rigorous calculations must fail at the very first step: they cannot, in general, deal with *exact* representations of numbers and functions. Moreover, arithmetical operations performed by computer do not in general produce exact answers, even in cases where the operands *are* represented exactly.

In order to cope with these problems some guarantees are needed from the manufacturers of computer hardware concerning the representation of numbers in computer memory. Further, it must be possible to systematically take into account all rounding errors that may occur during calculations, so that reliable bounds may be kept on the true (exact) results.

In 1985, the `IEEE` established a standard (Standard 754-1985, [IEEE85]) for the operation of computer hardware that addresses the above points. It specifies both a format for representing numbers in the computer memory, and certain guarantees on the accuracy of arithmetical operations (i.e. rules concerning how the computer must perform rounding after operations). The version of the standard which will be described here deals with a finite subset of the real numbers, each represented by the states of 64 bits of computer memory, and will be referred to as '64 bit `IEEE` arithmetic'.

The following section describes how floating point arithmetic is performed under hardware that conforms to this standard. The subsequent sections then demonstrate how these features may be used to obtain rigorous bounds on calculations with real (and complex) numbers and functions.

3.2.1 Representable Numbers

Physical computers (as opposed to mathematical idealisations) are, of course, limited by having *finite* resources and by the demand that their calculations finish in finite (indeed, *reasonable*) time. This means that, in general, it is impossible to represent real numbers in computer hardware and to perform operations with them. Conventional computer hardware uses a fixed-size binary 'floating point' representation for numbers in memory. By using this method only a finite subset of the real numbers, typically consisting of certain rational numbers with terminating binary expansions, may be represented.

Other approaches are possible. For example, variable-length storage techniques may be used to implement 'arbitrary precision arithmetic'. Of course, such approaches cannot overcome the fundamental difficulty of representing a general real number by the data held in a *finite* portion of computer memory. In addition, techniques involving variable-length storage typically require many of the machine's basic built-in operations to be re-implemented. In situations where extensive calculations must be carried out the time taken for such a program to execute can become prohibitive.

Instead, the approach that is taken here is to use both the usual fixed-length (in this case 64

bit) binary floating-point representation *and* the computer's in-built arithmetic operations, under the assurances provided by the `IEEE` standard concerning their reliability.

Storage format Under the 64 bit `IEEE` standard, numbers are represented in the following format. Each number takes up 64 bits of memory, these bits will be numbered 0–63. Bit zero is called the *sign bit* and specifies the sign of the number (positive or negative). Bits 1–11 represent the *exponent* of the number (the power of 2 in a certain binary expansion). Finally, bits 12–63 represent the *mantissa* of the number (also known as the 'fractional part').

Certain patterns of the states of bits are reserved to indicate special conditions. For example, the bit pattern with all of the 64 bits set to zero is used to represent the number zero. Some of these conditions are indicated in the following table (3.1).

Condition	Bit 0 (Sign)	Bits 1–11 (Exponent)	Bits 12–63 (Mantissa)
Zero	Zero	All Zeros	All Zeros
`isnan`		All Zeros	Not All Zeros
`isinf`		All Ones	All Zeros

Table 3.1: Format of `IEEE` double precision numbers.

The condition `NaN` or `isnan` is used to indicate that the result of an operation is 'not a number', this may arise when the result is not mathematically defined. The condition `isinf` is used to represent plus or minus infinity. 'Infinite' numbers are given a value to facilitate the definition of rounding (see Section 3.2.2).

With the exception of the special conditions indicated above, the following paragraphs describe how a real number is represented in this format.

Firstly, there is a convention that identifies the state of a bit with the integers zero and one in the binary expansion of a number. So, for example, the configuration of the 11 bits for the 'exponent' represents (under the standard base 2 representation for non-negative integers) an integer e in the range

$$0 \leq e \leq 2^{11} - 1 = 2047. \tag{3.1}$$

The values $e = 0$ and $e = 2047$ are reserved for special conditions, such as those given above. The remaining values are used to define a quantity E, with the value

$$E = e - (2^{10} - 1). \tag{3.2}$$

(Thus E ranges from -1023 to 1023.) This quantity will be used as the power of 2 in a certain binary expansion of the real number being represented.

Similarly, the sequence of values held in the 52 bits of the 'mantissa' are used to specify an integer

$$0 \leq m \leq 2^{52} - 1 = 4503599627370495. \tag{3.3}$$

This is used to define the quantity

$$M = 1 + (m \cdot 2^{-52}). \tag{3.4}$$

Thus M corresponds symbolically (i.e. formally) to the binary number $1.m$, i.e. the binary number with integer part 1 and with the mantissa making up the first 52 bits of the fractional part.

Finally, the value s associated with the state of the 'sign-bit' is used to indicate whether the number being represented is positive ($s = 0$) or negative ($s = 1$).

Under the above definitions, any pattern of 64 bits for which $0 < e < 2047$ is taken to represent a real number r defined by the following equation

$$\begin{aligned} r &= (-1)^s \cdot 1.m \cdot 2^{e-(2^{10}-1)} && (3.5a)\\ &= (-1)^s \cdot M \cdot 2^E. && (3.5b) \end{aligned}$$

Definition 37 (The Representable Numbers)
The set of real numbers of the form given in equation 3.5b, along with the number zero, will be called the representable numbers *and will be denoted* $\mathcal{R}$.

Under this scheme, the smallest positive representable number (corresponding to $s = 0$, $e = 1$, and $m = 0$) is 2^{-1022}. The largest positive representable number (corresponding to $s = 0$, $e = (2^{11} - 1) - 1$, and $m = 2^{52} - 1$) is $(1 - 2^{-53})2^{1024}$.

3.2.2 Accuracy of operations

In addition to providing a method for representing real numbers, the `IEEE` standard also specifies the accuracy to which arithmetical operations must be performed by the hardware. The key idea is that *the computer is required to perform all arithmetical operations to as many places as possible, with a minimum of rounding error, among the numbers it can represent.*

In order to make these specifications precise, the following definitions will be used.

Definition 38 (The Operations `up` and `down`)
Given a representable number $r \in \mathcal{R}$, *the operations* `up` *and* `down` *will be defined as follows:*

$$\begin{aligned} \texttt{up}(r) &= \inf\{s \in \mathcal{R} : s > r\} && (3.6a)\\ \texttt{down}(r) &= \sup\{s \in \mathcal{R} : s < r\} && (3.6b) \end{aligned}$$

In the case where the corresponding set is empty (i.e. where r *is the largest or smallest representable number), the symbol* `undefined` *will be used.*

In terms of the above definitions, the computer hardware must satisfy the following requirements on the way that operations are performed. (In what follows, arithmetical operators having a subscript c indicate the *calculated* values produced by computer, as opposed to the exact mathematical ones.) Let $r, s \in \mathcal{R}$ be representable numbers, then

- the computer correctly finds $-\mathtt{r} \in \mathcal{R}$,
- the computer correctly implements the comparisons $\mathtt{r} = \mathtt{s}$, $\mathtt{r} < \mathtt{s}$, $\mathtt{r} \leq \mathtt{s}$, $\mathtt{r} > \mathtt{s}$, and $\mathtt{r} \geq \mathtt{s}$,

Subject to the additional constraint that the exact result of the operations concerned is not either greater or smaller than *all* of the representable numbers (a situation which is termed 'overflow') it is also required that

- the calculated product $\mathtt{r} *_\mathtt{c} \mathtt{s}$ is such that

$$\mathtt{down}(\mathtt{r} *_\mathtt{c} \mathtt{s}) \leq \mathtt{rs} \leq \mathtt{up}(\mathtt{r} *_\mathtt{c} \mathtt{s}), \tag{3.7}$$

- the calculated quotient $\mathtt{r} \div_\mathtt{c} \mathtt{s}$ is such that

$$\mathtt{down}(\mathtt{r} \div_\mathtt{c} \mathtt{s}) \leq \mathtt{r}/\mathtt{s} \leq \mathtt{up}(\mathtt{r} \div_\mathtt{c} \mathtt{s}), \tag{3.8}$$

- the calculated sum $\mathtt{r} +_\mathtt{c} \mathtt{s}$ is such that

$$\mathtt{down}(\mathtt{r} +_\mathtt{c} \mathtt{s}) \leq \mathtt{r} + \mathtt{s} \leq \mathtt{up}(\mathtt{r} +_\mathtt{c} \mathtt{s}). \tag{3.9}$$

In addition, *if the exact result of a calculation happens to be a representable number, then the calculated answer must be identical to the exact one.*

3.2.3 Directed rounding modes

As stated above, when an operation is performed on `IEEE` standard hardware the computed result must be exact whenever the exact answer is representable. On the other hand, if the exact answer is not representable (but lies between the greatest and smallest representable numbers) then the computer must give either the nearest representable number that is greater than the exact answer, or the nearest representable number that is smaller than the exact answer. In fact, the hardware contains facilities for choosing which of these two answers will be given. Use of this feature is via so-called *directed rounding modes*. By choosing a particular mode, the computer can be instructed to consistently round either upward, downward, or to the nearest representable number.

More precisely, the following rounding modes are available:

- **Round up**: if the exact answer is representable, then the calculated answer is exact, otherwise the calculated answer is the next representable number *greater* than the exact answer.
- **Round down**: if the exact answer is representable, then the calculated answer is exact, otherwise the calculated answer is the next representable number *smaller* than the exact answer.
- **Round near**: (this is the default rounding mode) if the exact answer is representable, then the calculated answer is exact, otherwise the calculated answer is the *nearest* representable number to the exact answer.

With careful systematic use of these rounding modes, it is possible to implement standard mathematical operations in such a way as to get rigorous (i.e. genuine mathematical) upper and lower bounds on the exact results.

In order to use directed rounding modes in a computer program, software subroutines must be written which access the corresponding features of the machine's hardware. It suffices to write three routines, each of which puts the machine hardware into one of the directed rounding modes. It is important to note that the corresponding operations are *hardware specific*, i.e. the method for accessing rounding modes typically varies from machine to machine. This means that, in order to write code that may be used on several different platforms, it is necessary to write a different version of the module that accesses the hardware for each machine.

It is also important to note that some programming language compilers include *optimization* options, which may re-arrange the order of execution of code in order to increase speed. When compiling such hardware-dependent modules, it is advisable to switch these features off. To ensure that calculations are rigorous, it is also advisable to include a test to verify that the request for a particular rounding mode was successful.

In addition to providing subroutines to select rounding modes, it is useful to write subroutines to test the sign-bit of a number, and to check for the conditions `isnan` (indicating that the result of an operation is 'not a number') and `isinf` (indicating that a result is 'infinite'). These routines are useful for detecting conditions like overflow (which occurs when the result of an operation lies outside the range of representable numbers).

3.2.4 Rigorous Input and Output

When writing a computer program to perform rigorous arithmetic, there is an important additional consideration. Namely, the question of how numbers are to be input and output under the software language being used. Careful use of features like rounding modes during operations will be rendered pointless if it turns out that the computer does not input or output numbers in a reliable fashion. In particular, the internal representation of a number is in terms of a *binary* expansion (as indicated in the preceding sections), whereas *decimal* representations for numbers, being more convenient for the human programmer, are typically used in program code and for input/output operations such as printing results to a computer screen. To ensure that a program is rigorous, it is necessary to implement 'safe' versions of input/output operations that perform the necessary conversion between formats in a reliable way.

The next section demonstrates how rigorous bounds may be obtained for calculations involving real numbers by using *interval arithmetic* in which *intervals* are used instead of numbers. A computer implementation of interval arithmetic is then described, which involves using only representable numbers as the end-points of intervals and choosing suitable rounding modes during all calculations.

3.3 Interval arithmetic

This section gives a brief summary of the subject of *interval analysis*, also known as interval arithmetic (see Moore [Moo66]). The strategy is to decompose complicated mathematical formulae into elementary operations, and for each operation the aim is to find an *interval* which contains the exact result. The idea of interval arithmetic is then to work exclusively with algebra on these intervals, rather than with algebra on real numbers.

These notions will now be formalised.

Definition 39 (Intervals)
An interval *is defined to be a finite, closed interval of* $\mathbb{R}$,

$$[\ell, u] = \{r \in \mathbb{R} : \ell \leq r \leq u\} \quad \textit{where } \ell, u \in \mathbb{R},\ \ell \leq u. \tag{3.10}$$

The set of all such intervals will be denoted $\mathbb{J}$.

Operations on these intervals should be defined in such a way that the result of any *pointwise* operations with numbers is guaranteed to lie within the resulting interval. For example, given two intervals $A, B \in \mathbb{J}$, the result of the interval addition operation $A +_i B$ should be such that

$$a + b \in (A +_i B) \quad \text{for all } a \in A, b \in B. \tag{3.11}$$

3.3.1 Elementary operations

Interval arithmetic versions of standard arithmetical operations are given below. The interval version of operations are marked with a subscript i to distinguish them from the usual operations.

1. Addition:

$$[\ell, u] +_i [\ell', u'] = [\ell + \ell', u + u']. \tag{3.12}$$

2. Unary Minus:

$$-_i[\ell, u] = [-u, -\ell]. \tag{3.13}$$

3. Inverse:

$$1 \div_i [\ell, u] = \begin{cases} [1/u, 1/\ell] & \text{if } \ell > 0 \text{ or } u < 0, \\ \texttt{undefined} & \text{otherwise.} \end{cases} \tag{3.14}$$

4. Product:

$$[\ell, u] *_i [\ell', u'] = \begin{array}{l} [\min(\ell\ell', \ell u', u\ell', uu'), \\ \max(\ell\ell', \ell u', u\ell', uu')]. \end{array} \tag{3.15}$$

5. Absolute value (interval):

$$\mathrm{abs_i}[\ell,\mathtt{u}] = \begin{cases} [0, \max(|\ell|,|\mathtt{u}|)] & \text{if } \ell\mathtt{u} \le 0, \\ [\min(|\ell|,|\mathtt{u}|), \max(|\ell|,|\mathtt{u}|)] & \text{otherwise.} \end{cases} \tag{3.16}$$

6. Upper bound on real absolute value:

$$\mathrm{abs_{ir}}[\ell,\mathtt{u}] = \max(|\ell|,|\mathtt{u}|). \tag{3.17}$$

3.3.2 Computer implementation

Interval arithmetic, as described above, uses *real* numbers for the end-points ℓ and $\mathtt{u}$ of each interval $[\ell,\mathtt{u}]$. By working with the subset of intervals that have *representable* numbers as their end-points, and by using suitable directed rounding modes during all calculations, interval arithmetic may be performed rigorously by computer. That is, the intervals produced by computer operations can be guaranteed to contain the exact mathematical results. In particular, the computed intervals will be *supersets* of those produced by conventional (exact) interval arithmetic.

The first step is to restrict attention to only those intervals that are representable.

Definition 40 (Representable Intervals)
An interval $\mathtt{i} \in \mathtt{I}$ *is said to be* representable *if and only if its end-points are representable. The set of representable intervals will be denoted* $\mathcal{I}$.

$$\mathcal{I} \subset \mathtt{I} \cap (\mathcal{R} \times \mathcal{R}). \tag{3.18}$$

In order to see how the implementation of reliable computer versions of interval operations is accomplished, consider the example of *interval addition*. The preceding description of the `IEEE` standard means that if $[\ell,\mathtt{u}]$ and $[\ell',\mathtt{u}']$ are representable intervals, then

$$\begin{aligned} [\ell,\mathtt{u}] +_\mathrm{i} [\ell',\mathtt{u}'] &= [\ell+\ell', \mathtt{u}+\mathtt{u}'] \\ &\subseteq [\mathtt{down}(\ell +_\mathrm{c} \ell'), \mathtt{up}(\mathtt{u} +_\mathrm{c} \mathtt{u}')] \end{aligned}$$

(recall that $+_\mathrm{c}$ indicates the *calculated* value) i.e. instructing the computer to round *down* when calculating the sum of the lower bounds (ℓ and ℓ') and to round *up* when computing the sum of the upper bounds ($\mathtt{u}$ and $\mathtt{u}'$) guarantees that the resulting calculated (representable) interval will be a superset of the exact mathematical one.

By choosing suitable rounding modes during all calculations, it is possible to implement all of the operations of interval arithmetic rigorously on a computer.

Object-oriented languages like `C++` (detailed in Stroustrup [Str93]) give the programmer the ability to define new data-types (in `C++` these are called *classes*), and to endow them with suitable versions of standard arithmetical operations (by a technique known as 'operator overloading'). This makes such languages particularly suitable for this kind of work. The data-structure used to represent intervals is quite simple. Each interval is simply of pair of double precision (i.e. 64 bit `IEEE`) machine numbers, one for the lower bound of the interval and one for the upper bound. The following piece of code (program fragment 3.1) is part of a `C++` declaration that was used for this purpose:

Program Fragment 3.1 Structure of the `Interval` class.

```
class Interval
{
  private:
    double     low;   //lower bound
    double     high;  //upper bound
  public:
    //Read-only access to the attributes:
    double     LowerBound() const {return low;}
    double     UpperBound() const {return high;}
    ...
    //Operations:
    Interval&  operator+=(const Interval&);
    ...
};
```

(Code after the symbol `//` indicates a comment.)

Encapsulation: A few words of explanation are in order. It may be helpful to think of a `C++` object, i.e. a single *instance* of a `C++` class (for example, an `Interval`), as being somewhat like an organic cell. The inner-workings of the cell are hidden away inside a membrane, protecting them from the outside world. At its surface, the cell presents interfaces (for example, receptors for certain chemicals) to other cells. The analogous concept in `C++` (and other object-oriented languages) is known as *encapsulation*. In the code above, for example, the values of the lower and upper bound are shielded from outside interference by declaring them to be `private`. The interactions of any `Interval` with the "outside world", i.e. with the rest of the program, are mediated by the operations and functions that it presents at its "surface": these are declared to be `public`. For example, in order for another object to find out what the lower bound on the `Interval` is, it must use the public function `LowerBound()` — it cannot access the lower bound directly (further, we use the declaration `const`, which is a way of ensuring that this request for information is not allowed to change the `Interval`, i.e. it is kept *constant*).

This "semi-permeable membrane" of access restriction is extremely useful when writing programs: when used carefully it ensures that objects keep their integrity by insisting that all attempts to alter the object can only use the object's own set of interfaces. For example, suppose that all programming took place at a level where we accessed the upper and lower bounds on the `Interval` directly and manipulated them. This would result in much repeated code as the common operations (e.g. adding `Intervals`) were used again and again. A mistake in the programming at this level would be disastrous. (A worse case-scenario would be a piece of code that resulted in an `Interval` whose lower bound was greater than its upper bound.) By implementing a definite set of "standard" `Interval` operations which act on both the upper and lower bounds in a consistent and mathematically correct way, and by insisting that `Intervals` may only be altered by using these operations, we preclude the possibility of many different types of error from

ever occurring. (This technique also makes the program simpler and much easier to read, which is a bonus in itself.)

Implementing Interval Addition: To see how operations are implemented rigorously consider, for example, the addition operation described above. This involves the following steps:

1. instruct the computer to round up
2. add together the upper bounds
3. instruct the computer to round down
4. add together the lower bounds
5. return the result

A corresponding fragment of program code is the following, which increments an existing interval by adding another (`b`) to it. This is accomplished by *overloading* the `C++` "`+=`" operator:

Program Fragment 3.2 Interval addition.

```
Interval& Interval::operator+=(const Interval& b)
{
  RoundUp();           //select upward rounding
  high += b.high;      //add upper bounds
  RoundDown();         //select downward rounding
  low += b.low;        //add lower bounds
  RoundNear();         //resume default rounding mode
  return *this;        //output result
}
```

Careful use of features like operator-overloading helps the programmer to write transparent code, whose purpose is self-evident. This is desirable, as it makes the computer portion of a proof much easier to verify. For example, the following fragment of code:

Program Fragment 3.3 Interval arithmetic example.

```
  Interval one(1.0,1.0);
  Interval three(3.0,3.0);
  Interval third = one/three;   //interval division
```

gives the result (to low accuracy) of `[ 3.33333e-01, 3.33334e-01]` for the variable `third`, i.e. $\frac{1}{3} \in [0.333333, 0.333334]$.

Remark 41 (Operator Overloading as Syntactic Sugar)
From a programming perspective it should be borne in mind, however, that operator overloading is really "syntactic sugar" for using function calls. For example, the above

could equally well be expressed by something of the form `third.Assign( Quotient( one,three))` *which is, however, less pleasing to the mathematical eye — syntactic sugar can be very sweet!*

3.4 Rectangle arithmetic

In order to prove the theorems in this thesis, it is necessary to deal with *complex* numbers and with operations on complex functions. Two extensions to interval arithmetic come to mind (1) complex arithmetic using discs, and (2) using rectangles parallel to the axes (i.e. using the Cartesian product of two intervals, one for the real part and another for the imaginary part of the number). The latter seems to be the most straightforward extension for handling complex numbers; a pair of intervals will be used to define a *rectangle* in the complex plane. This reduces the problem of implementing complex arithmetic to that of successfully combining existing (real) interval operations. Since real interval arithmetic is needed for the proofs anyway, this approach was taken.

The following `C++` data-structure indicates the computer implementation.

Program Fragment 3.4 Structure of the `rectangle` class.

```
class Rectangle
{
  private:
    Interval         re;  //real part
    Interval         im;  //imaginary part
  public:
    //Read-only access to the attributes:
    const Interval&  Real() const {return re;}
    const Interval&  Imag() const {return im;}
    //Operations:
    Rectangle        operator*(const Rectangle&, const Rectangle&);
    ...
};
```

Rectangle multiplication, for example, reduces to interval multiplication and interval addition: for rectangles $\mathtt{a} = (\mathtt{a}_{\mathrm{real}}, \mathtt{a}_{\mathrm{imag}})$ and $\mathtt{b} = (\mathtt{b}_{\mathrm{real}}, \mathtt{b}_{\mathrm{imag}})$ (where the components of each pair are intervals) the rectangle multiplication operation $\mathtt{a} *_{\mathrm{r}} \mathtt{b}$ is defined by

$$\mathtt{a} *_{\mathrm{r}} \mathtt{b} = (\mathtt{a}_{\mathrm{real}} *_{\mathrm{i}} \mathtt{b}_{\mathrm{real}} -_{\mathrm{i}} \mathtt{a}_{\mathrm{imag}} *_{\mathrm{i}} \mathtt{b}_{\mathrm{imag}},\ \mathtt{a}_{\mathrm{real}} *_{\mathrm{i}} \mathtt{b}_{\mathrm{imag}} +_{\mathrm{i}} \mathtt{a}_{\mathrm{imag}} *_{\mathrm{i}} \mathtt{b}_{\mathrm{real}}), \tag{3.19}$$

corresponding to the `C++` code:

Program Fragment 3.5 Rectangle multiplication.

```
Rectangle Rectangle::operator*(const Rectangle& a, const Rectangle& b)
{
    Interval x = a.Real()*b.Real() - a.Imag()*b.Imag();
    Interval y = a.Real()*b.Imag() + a.Imag()*b.Real();
    return Rectangle(x,y);
}
```

Unfortunately, there is a serious problem associated with using the class of rectangles parallel to the axes: it is not rotation invariant. This means that multiplication operations involving rectangles (as described above) often produce inefficient bounds. In particular, rectangle multiplication is not associative, which should be apparent because interval arithmetic is not distributive. In fact, Cauchy's law

$$|zw| \leq |z| \cdot |w| \quad \text{and, in particular,} \quad |zw| = |z| \cdot |w|, \tag{3.20}$$

is violated.

For example, consider the rectangle product $Z *_{\mathrm{r}} W$ with $Z = ([-1,1],[-1,0])$ and $W = \frac{1}{2}([1,1],[1,1])$. For $z \in Z$ and $w \in W$ it may be seen that $|z| \leq \sqrt{2}$ and $|w| \leq \frac{1}{\sqrt{2}}$, giving $|zw| \leq 1$. However, rectangle arithmetic (using interval arithmetic on the real and imaginary parts) gives the result

$$Z *_{\mathrm{r}} W = ([-\frac{1}{2},1],\ [-1,\frac{1}{2}]), \tag{3.21}$$

with maximal absolute value $\sqrt{2} > 1$.

This consideration means that great care must be taken when, for example, calculating high powers of complex numbers using rectangle arithmetic. The 'naive' method of repeated multiplication typically leads to appalling bounds which are not tight enough to be used in a proof. In these kinds of situations it is necessary to use more efficient methods, such as repeated squaring, rather than straightforward multiplication (see, for example, Knuth [Knu81]).

3.5 Functional analysis

The previous sections demonstrated that rigorous estimates are possible for operations on sets of real numbers (intervals) and sets of complex numbers (rectangles). This section demonstrates that these techniques may be extended to yield analogous estimates for sets of complex *functions*, also known as *standard function balls*. The method closely follows that given in Eckmann, Koch, and Wittwer [EKW84] and is designed to allow efficient computer implementation.

Complex analytic functions in one complex variable, defined on the unit disc, are considered. By preceding such functions with an affine map (taking a general disc to the unit disc), the framework is then extended to allow functions defined on general discs (*general*

function balls). The key concept is to represent a set of analytic functions by the coefficients of a Taylor series expansion (truncated to some degree), along with rigorous bounds on the high-order part of the expansion.

By using only representable numbers to represent function balls, and using analogous techniques to those of interval (and rectangle) analysis, the framework may be programmed rigorously on a computer.

3.5.1 Standard Function balls

The first step is to decide on a suitable space of functions. In what follows, let D denote the open unit disc

$$D = \{x \in \mathbb{C} : |x| < 1\}. \tag{3.22}$$

The space that will be used is the set $\mathcal{A}$ of complex analytic functions on D that are continuous and bounded on the closed unit disc $\overline{D}$, and of finite ℓ_1-norm. For a function f written as

$$f(z) = \sum_{j=0}^{\infty} f_j z^j. \tag{3.23}$$

the ℓ_1-norm is given by

$$\|f\|_1 = \sum_{j=0}^{\infty} |f_j|. \tag{3.24}$$

Equipped with the above norm, $\mathcal{A}$ forms a (complex) Banach space (see [Rud82]).

In order that a computer may be used, a representation of these functions is needed in terms of a *finite* number of quantities. The techniques developed earlier may then be used to produce bounds on the results of operations involving these functions. To this end the concept of a *boundary* is defined. A boundary serves to define a *closed*, *convex* set in $\mathcal{A}$ which will be called a *standard function ball*. Operations on boundaries are then implemented in such a way that the result of pointwise operations (on functions in the corresponding ball) are always guaranteed to lie within the ball represented by the resulting boundary (c.f. interval arithmetic, section 3.3).

The boundaries used here take the form of triples,

$$\nu = (\nu_P, \nu_H, \nu_G), \tag{3.25}$$

where $\nu_P = (\nu_j; 0 \le j \le N)$ is a *vector* of $N+1$ rectangles, and $\nu_H, \nu_G > 0$ are positive real numbers. The set of all such boundaries will be denoted $\mathcal{B}$.

Each boundary $\nu \in \mathcal{B}$ serves to define a *standard function ball* $\mathcal{A}(\nu) \subset \mathcal{A}$ containing all the functions that may be written as follows:

$$\mathcal{A}(\nu) = \{f \in \mathcal{A} : f = f_P + f_H + f_G\}, \tag{3.26}$$

where the following restrictions are made on f_P, f_H, and f_G.

1. f_P is called the *polynomial part* and represents a truncated Taylor series, whose coefficients f_j lie within the corresponding rectangles ν_j (of ν_P). Formally,

$$f_P(z) = \sum_{j=0}^{N} f_j z^j \tag{3.27}$$

   where $f_j \in \nu_j$. The quantity $N \in \mathbb{N}$ will be fixed throughout a computer-assisted proof, and is called the *truncation degree*. (Increasing N generally leads to sharper bounds.)

2. $f_H \in \mathcal{A}$ is called the *high-order part* and is a 'high-order' function (which may be thought of as representing the 'neglected' terms in the truncated Taylor expansion). More formally, the quantity

$$\lim_{s \to 0} \frac{f_H(sz)}{s^{N+1}} \tag{3.28}$$

   is finite for all $z \in \overline{D}$, and f_H is bounded by

$$\|f_H\|_1 \leq \nu_H. \tag{3.29}$$

3. $f_G \in \mathcal{A}$ is called the *general part*, and may contain terms of any order. It is bounded by

$$\|f_G\|_1 \leq \nu_G. \tag{3.30}$$

   The bound ν_G (sometimes called the *error bound*) may be thought of as defining the radius of a ball in the space of functions. It provides a convenient means of performing calculations with balls of functions (for example, in the existence proof described in chapter 2 it was necessary to estimate the norm of the derivative of an operator on a small ball of functions around an approximate fixed point). This quantity is also used to improve the efficiency of bounds obtained during functional *composition* (this is discussed later in this section). The general idea is that the norm of f_G should be very small.

It remains to demonstrate how operations on function balls are implemented in terms of operations on their boundaries.

The following operations are among those that may be implemented (in what follows f and g are function balls and a, z are complex scalars):

- multiplication by a scalar (rectangle) af
- addition and subtraction $f + g$, $f - g$
- norm $\|f\|$
- multiplication $f \cdot g$
- functional composition $f \circ g$

- differentiation followed by composition $\mathtt{f}' \circ \mathtt{g}$
- evaluation on a given rectangle $\mathtt{f}(z)$
- evaluation of the derivative on a given rectangle $\mathtt{f}'(z)$

Remark 42 (Derivative of a Function Ball)
Note that while the evaluation of a derivative and differentiation followed by composition are implemented, it is not generally meaningful to consider differentiating a ball of functions. (To see why this is so, notice that the formal power series obtained by term-by-term differentiation of an existing series will not necessarily converge.)

Further details of two fundamental operations are given below to illustrate the technique. The first involves calculating the norm of a function ball (i.e. calculating an interval that contains the norms of all functions in the function ball), the second involves composing two function balls together (i.e. calculating a new function ball which contains the composition of every possible pair of functions consisting of one from each ball). The discussion of functional composition also serves to better explain the role of the 'general term' $\mathtt{f}_\mathrm{G}$ in these calculations.

Norm

Given a boundary $\nu \in \mathcal{B}$, it is required to find an *interval* that is guaranteed to contain the norm of every function $\mathtt{f}$ within the corresponding function ball $\mathcal{A}(\nu)$.

Take $\nu = (\nu_\mathrm{P}, \nu_\mathrm{H}, \nu_\mathrm{G})$ and consider a corresponding function $\mathtt{f} \in \mathcal{A}(\nu)$ with $\mathtt{f} = \mathtt{f}_\mathrm{P} + \mathtt{f}_\mathrm{H} + \mathtt{f}_\mathrm{G}$. Then it follows that

$$\|\mathtt{f}\|_1 \in \mathtt{norm}_\mathrm{b}(\nu), \tag{3.31}$$

where

$$\mathtt{norm}_\mathrm{b}(\nu) = \sum_{\mathrm{j}=0}^{\mathrm{N}} \mathrm{abs}_\mathrm{i}(\nu_\mathrm{j}) +_\mathrm{i} [0, \nu_\mathrm{H}] +_\mathrm{i} [-\nu_\mathrm{G}, \nu_\mathrm{G}]. \tag{3.32}$$

(Recall that the ν_j are bounds on the coefficients of the polynomial part, and that $\nu_\mathrm{H}, \nu_\mathrm{G}$ are bounds on the norm of the high-order and general parts, respectively. Intuitively, the polynomial and high order parts, $\mathtt{f}_\mathrm{P}$ and $\mathtt{f}_\mathrm{H}$, both contribute independently to the norm, whereas the general part $\mathtt{f}_\mathrm{G}$ may affect all of the coefficients in the Taylor expansion, and so may contribute negatively.)

Correspondingly, a positive upper bound on the norm is defined as follows

$$\|\mathtt{f}\|_1 \leq \mathtt{norm}_\mathrm{br}(\nu) = \sum_{\mathrm{j}=0}^{\mathrm{N}} \mathrm{abs}_\mathrm{ir}(\nu_\mathrm{j}) + \nu_\mathrm{H} + \nu_\mathrm{G}. \tag{3.33}$$

3.5.2 Composition

In addition to providing a useful way of performing calculations for balls of functions centred on some polynomial, the general part f_G also plays an important role in producing efficient (i.e. tight) bounds. The following discussion aims to clarify this point.

For implementing standard vector space operations, the general part f_G is not required. However, the advantages of using this term become clear when the implementation of functional *composition* is considered. Take two functions $f, g \in \mathcal{A}$ and suppose that

$$f = f_P + f_H, \tag{3.34}$$

i.e. that there is no general part, f_G. If, in addition, $\|g\|_1 \leq 1$ then the composition

$$f \circ g, \tag{3.35}$$

is a well defined member of the space. In particular,

$$\|f \circ g\|_1 \leq \|f\|_1. \tag{3.36}$$

Using linearity to separate the composition into two parts gives

$$f \circ g = f_P \circ g + f_H \circ g. \tag{3.37}$$

Computation of bounds on the first term is straightforward as it consists of the composition of a polynomial with a function ball, a procedure which involves only elementary algebraic operations. However, estimates are also required on the second term $f_H \circ g$. Unfortunately, no information is available concerning the coefficients of f_H, other than the fact that

$$\|f_H\|_1 \leq \nu_H, \tag{3.38}$$

which results in the bound

$$\|f_H \circ g\|_1 \leq \|f_H\|_1 \cdot \|g\|_1^{N+1} \leq \nu_H \cdot \|g\|_1^{N+1}. \tag{3.39}$$

In the case where g has no constant term, $f_H \circ g$ generates only terms of order strictly greater than N (remember that f_H consists only of terms of degree greater than N). In this special case, then, the term $f_H \circ g$ simply contributes to the high order bound of the result.

In general, however, g will contain a non-zero constant term. This means that $f_H \circ g$ will also contribute to terms of order smaller than N in the result. i.e. it will contribute to the polynomial part $(f_H \circ g)_P$. One way to cope with this (in addition to adjusting the high-order bound of the result) would be to enlarge each of the rectangles bounding the the coefficients of the polynomial part of the result. However, this approach dramatically increases the number of elementary operations that must be performed and typically leads to very inefficient bounds.

An alternative approach is to use a 'general part' which may contain terms of any order. The contribution from terms like $f_H \circ g$ may then be absorbed into the bound on the general part $(f_H \circ g)_G$ of the result. Of course, using this extra term is done at the expense of

having to carry it through all of the Banach algebra operations, where it is not actually needed. However, this strategy does not lead either to a drastic performance penalty or a loss of precision.

The actual strategy taken during functional composition is somewhat more complicated than indicated above, and involves removing the constant term of g_P from the rest of the polynomial part and dealing with its contribution separately. In this way, not too much is lost if (by chance) g happens to have a zero, or very small, constant term. This is very important for the estimates used in the existence proof where, correspondingly, the domains are chosen so as to make the most 'important' functions used in compositions have a zero constant term (see Section 2.6).

3.5.3 Computer Implementation

The following fragment of C++ code illustrates the implementation of function balls on the unit disc. From a programming perspective, it is useful to first implement two extra classes: a `Vector` consists of an array of rectangles and is used to bound the coefficients of the polynomial part, and a `Pub` represents a *positive upper bound* upon which arithmetical operations are defined which consistently round *upward*.

Program Fragment 3.6 Structure of the `Ball` class.

```
class Ball
{
  private:
    static int       n;        //truncation degree for all balls
    Vector           poly;     //bounds on polynomial part
    Pub              high;     //high-order bound
    Pub              general;  //general bound
    int              pp;       //indicates pure power
  public:
    //Read-only access to the attributes:
    static int       Degree()               {return n;}
    const Vector&    Coefficients()   const {return poly;}
    Pub              ErrorBound()     const {return general;}
    Pub              HighOrderBound() const {return high;}
    bool             IsPurePower()    const {return (pp>=0);}
    //Operations:
    friend Interval  Norm(const Ball& b);              //norm
    Ball             operator()(const Ball& b) const;  //compose
    ...
};
```

For the sake of efficiency, an integer `pp` is also included in the structure. This is used to indicate whether the function represented is a *pure power*, i.e. if it takes the form $z \mapsto z^j$ where j is a non-negative integer. Treating pure powers as a special case is useful when performing *evaluation* of functions: ordinarily one would use Horner's algorithm (nested

multiplication), instead a more efficient algorithm is used for pure powers which reduces the number of multiplication operations by using repeated squaring. This technique is essential for producing good estimates for high values of j.

The implementation of operations for function balls then builds on the implementations of intervals, rectangles, and vectors of rectangles. The following program fragment (3.7) illustrates how the implementation of the norm looks in this framework (c.f. equation 3.32):

Program Fragment 3.7 Norm of a Standard Function Ball

```
Interval Norm(const Ball& b)
{
    return Norm(b.poly)
           + PositiveInterval(b.high)
           + SymmetricInterval(b.general);
}
```

(The keyword `friend` in program fragment 3.6 was used to declare that the `Norm` function should be allowed access to the internal structure of the function ball, which is apparent in the above because the `poly`, `high` and `general` parts of `b` are accessed without using corresponding public access functions. This is purely for the sake of efficiency and aesthetics: `Norm` is a straightforward operation for a function ball which requires a single access to each data member and it makes sense to grant this access directly for such a simple function, rather than insisting that it go through a public interface. An alternative might be to use "inlined" access functions. See, for example, [Str93] and [CL94])

Another advantage of operator overloading occurs for the implementation of function balls on computer: we can use the convenient notation `f = g(h)` to represent composition by overloading the `C++` "function call operator", `operator()`.

3.5.4 General analytic functions

The previous subsections described the use of *standard* function balls; sets of analytic functions defined on the *unit disc*. In general, however, there is a need to deal with functions defined on other domains. The extension of the mathematical framework to cope with functions defined on arbitrary discs is simple. Let f be an analytic function defined on an open disc D_f,

$$D_f := \{z : |z - c_f| < r_f\}. \tag{3.40}$$

(where f is analytic on D_f and continuous and bounded on $\overline{D_f}$). f is expanded into a power series of the form

$$f(z) = \sum_{j=0}^{\infty} f_j \left(\frac{z - c_f}{r_f}\right)^j. \tag{3.41}$$

This is equivalent to preceding an analytic function $F \in \mathcal{A}$ (defined on the open unit disc) by the affine map which takes the disc D_f to the unit disc, as follows:

$$f(z) = F\left(\frac{z - c_f}{r_f}\right). \tag{3.42}$$

The norm $\|f\|$ is then defined to be $\|F\|$, i.e. the norm of corresponding function on the unit disc.

To this end, balls of functions defined on general domains (*general* function balls) are implemented in the computer framework by a structure composed of a circle (defining the extent of the domain) and a corresponding ball of functions defined on the unit disc (a *standard* function ball). This is indicated by the following program fragment (3.8).

Program Fragment 3.8 Structure of the `Function` class.

```
class Function
{
  private:
    Circle  domain;    //domain boundary
    Ball    standard;  //standard (unit-disc) ball
  public:
    friend Interval Norm(const Function&);  //norm
    ...
};
...
Interval Norm(const Function& f)
{
    return Norm(f.standard);  //use norm of standard ball
}
...
```

(Of course, the above depends on first implementing a suitable `Circle` class.)

For example, to obtain bounds on the evaluation of a general function ball on a rectangle of complex numbers, bounds are first obtained on the affine map which takes the general domain to the unit disc, then the corresponding standard function ball is applied to give the final bounds.

Operations, such as composing these function balls together, can then be implemented in terms of operations on the appropriate standard function balls (subject to suitable restrictions on the domains of the operands).

3.5.5 Pairs of functions

In order to implement the renormalisation operator, it is helpful to design a class to represent *pairs* of analytic functions, each equipped with a (disc) domain of definition. This is fairly straightforward, since the typical operations to be performed on the pair can be built-up from their counterparts for single functions. In addition, facilities to support the renormalisation operator, for example routines to compute the coefficients for the

projection to the commuting subspace (c.f. section 2.5.2), can also be provided as part of the class.

The basic outline is as follows (program fragment 3.9).

Program Fragment 3.9 Structure of the `Pair` class.

```
class Pair
{
  private:
    Function         U;
    Function         V;
    ...
  public:
    const Function&  FirstFunction()   const {return U;}
    const Function&  SecondFunction()  const {return V;}
    const Circle&    FirstDomain()     const {return U.Domain();}
    ...
    Pub              ErrorBound()      const;
    Pub              HighOrderBound()  const;

    //Elementary operations:
    Pair&            operator=(const Pair&);   //assign value
    Pair&            operator+=(const Pair&);  //addition
    Pair&            operator*=(const Rectangle&);  //scalar mult
    ...

    //Rescaling factor:
    Rectangle        Alpha()           const {return V(0);}

    //Routines for renormalisation:
    Pair&            Project();      //project to subspace
    Pair&            Conjugate();    //conjugate functions
    Pair&            Compose();      //composition
    Pair&            Renormalize();  //renormalise (using the above)
    ...
};
```

For example, the error and high-order bounds can be constructed from their `Function` counterparts.

3.6 Validity of Computer-assisted Proofs

Some remarks are perhaps in order concerning the status of computer-assisted proofs. In particular, to what extent is a computer-assisted proof considered to be valid? A few of the most prominent points will be discussed here.

One view is that the only kind of proof that is really acceptable is a conceptual proof whose computational part involves only calculations that may be verified by another mathemati-

cian. This is certainly a reasonable attitude; indeed, it is always desirable to have a conceptual proof that involves little or no computation. However, it is often the case that failure to find an analytical proof blocks progress in an area for considerable time. In the interim, a computer-assisted proof may encourage further progress in the general area. In addition, there is a case for believing that certain aspects of computer-assisted proofs, for example the lack of ambiguity that results from having to break the calculation into explicit program steps, might make them at least as easy to verify as some of the more lengthy conceptual proofs.

Below, we briefly discuss some of the other considerations.

Hardware and Operating System Reliability: Computers are complicated devices, consisting of a great many interacting parts. In addition, powerful workstations, of the sort needed to run more difficult proofs, typically have an operating system that caters for several users at once. If a proof requires large amounts of computer resources, then the operating system will typically shuffle its resources around (for example, using areas of a hard-disk as a virtual memory). In this kind of environment, it is not possible to give a concrete guarantee that nothing will go wrong.

However, it seems possible to obtain a very low probability of error by taking care with the way that programs are designed. As much as possible, they are designed to be hardware and compiler independent, using features of the programming language that are stable, and taking special care with memory allocation and de-allocation.

The Mathematics of Error Bounds One very important point that is sometimes misunderstood is that the error bounds are designed to be true mathematical bounds; they are not obtained by heuristics, but are based on solid mathematics. Thus the computer-part of the proof is designed to be much more than just a sophisticated numerical experiment. (The details presented in the previous sections of this chapter were intended in part to clarify this point.)

Communication and Reproducibility Many conceptual proofs have some computational aspect (though, typically, one that can be verified by hand); they vary in the extent to which they are computational. One thing that distinguishes a computer-assisted proof is that there is usually a clean and definite separation between the conceptual and computational parts. Furthermore, the computational part is completely and unambiguously specified by a definite sequence of steps, i.e. the program, that is verified syntactically and, to some extent, semantically by a strict mathematical process (the action of the programming language compiler). The fact that the computational part may be communicated accurately and unambiguously and then verified by another mathematician in a systematic way is certainly attractive. (In addition to a separation between the conceptual and computational parts, there is an additional distinction which arises from the assumption that the computer hardware is really carrying out the steps specified. One can verify that the steps themselves are correct and consider independently the issue of whether they are truly embodied in a particular piece of executable code.)

Summary Perhaps it is best to accept the "computer part" of a computer-assisted proof for what it is: At worst, it is a rather sophisticated numerical experiment that yields information of at least the same reliability as any non-rigorous counterpart. At best, it is a proof in the traditional sense and, further, one that may potentially be communicated efficiently to other mathematicians with very great accuracy and lack of ambiguity. (In either case, any intuition gained into the problem as a result may be regarded as a bonus.)

Chapter 4

Properties of Boundary Curves

4.1 Introduction

The aim of this chapter is to prove that the critical Siegel disc renormalization fixed point whose existence was established using a rigorous computer-assisted proof in Chapter 2, and any pair in its domain of attraction, has a Hölder continuous invariant curve through a critical point (see chapter 1). This means that any complex analytic map attracted by the fixed point has a Siegel disc bounded by a Hölder continuous Jordan curve passing through a critical point of the map. It is further proved that the motion on the curve is conjugate to a pure rotation, and that the conjugation function is not differentiable on a dense set of points. It is known empirically that the stable manifold of the fixed point has real co-dimension two. This would imply that the above holds for an *open* set of maps having a golden mean Siegel disc. (These results were published in Burbanks and Stirnemann [BS95].)

The proof is, once again, computer assisted and uses the *necklace construction* of [Sti93c]. In order to carry the necklace construction out for the critical fixed point, it is necessary to rigorously verify certain topological assumptions concerning the action of the maps of the critical fixed point on their domains. In order to perform this part of the proof, the rigorous framework presented in the previous chapter was extended to allow function balls to be evaluated on domain boundaries (in this case circles and their images under certain maps).

The necklace construction was developed with twist maps of the cylinder in mind (recall section 1.2.3) and designed for the corresponding renormalization operator (generally known as *MacKay's operator*, recall section 1.2.3). Fortunately, the Siegel disc case is sufficiently similar that the construction may be adapted and used here. In fact, the necklace construction had already been applied to the *simple* fixed point of Siegel disc renormalization in [Sti93a], in order to prove a special case of Siegel's theorem.

4.1.1 Organization

This chapter is divided into four main sections. The first gives the main hypotheses and results that are proved for the Siegel disc case. The original necklace construction is then summarised in order to lay the foundations for the main proofs. Next, the construction is translated into the framework of Siegel disc renormalization. Details are then given of the computer-assisted proof that verifies both the translated hypotheses of the necklace construction and also the additional constraints needed for the Siegel disc case. Finally, the chapter concludes with a short discussion of the relevance of these results to complex dynamics (in particular, the current state of knowledge concerning the spectrum of the operator is given).

4.1.2 Hypotheses and Results

First, the results established in this chapter will be stated. Then it will be necessary to summarise the *necklace construction* of Stirnemann [Sti93c] which is adapted for the proofs.

The accretive fixed point

In this chapter, the fixed point equation with the *accretive* composition order, i.e. EF, will be needed. Note that this is the opposite order to that used in the existence proof. However, it turns out that is is possible to construct analytic continuations of the maps E and F of the existence proof to new domains on which the accretive fixed point equation holds. Formally:

Proposition 43 (Existence of an Accretive Fixed Point)
There exist analytic continuations of the maps E *and* F *(which will again be denoted by* E *and* F*), to domains* Δ_{E} *and* Δ_{F}*, such that*

$$\mathsf{E}(z) = \mathsf{C}\frac{1}{\alpha}\mathsf{F}\alpha\mathsf{C}(z) \quad \text{on } \Delta_{\mathsf{E}}, \tag{4.1a}$$

$$\mathsf{F}(z) = \mathsf{C}\frac{1}{\alpha}\mathsf{EF}\alpha\mathsf{C}(z) \quad \text{on } \Delta_{\mathsf{F}}. \tag{4.1b}$$

This result is proved in Section 4.3.2, essentially by using the commutativity of E and F. In this chapter, the pair (E,F) on the *new* domains will be referred to as the *accretive fixed point* or, simply, the *fixed point*.

Invariant Curve of the Fixed Point

In this chapter, the following theorem for the fixed point (E,F) is established:

Theorem 44 (Invariant Curve for the Fixed Point)
Let (E,F) *denote the fixed point. There exists a Hölder continuous injective curve (see definition below)*

$$\chi : [-\omega, 1-\omega] \to \mathbb{C},$$

such that

$$\begin{aligned} E\chi(t) &= \chi(t-1) \quad \text{for } 1-\omega \le t \le 1, && (4.2a)\\ F\chi(t) &= \chi(t+\omega) \quad \text{for } -\omega \le 1 \le 1-\omega, && (4.2b)\\ \alpha C\chi(t) &= \chi(-\omega t), && (4.2c)\\ \chi(0) &= 0. && (4.2d) \end{aligned}$$

This curve is not differentiable on a dense set of points.

(As usual, in the above ω denotes the golden mean $\frac{1}{2}(\sqrt{5}-1)$.)

Definition 45 (Hölder Continuity)
A function $f : X \to Y$ *is said to be* Hölder continuous *(or a* Hölder function of exponent α*) if for all* $x, y \in X$ *there exits a constant* c *with*

$$|f(x) - f(y)| \le c|x - y|^{\alpha}.$$

(If α *may be taken to be* 1*, then* f *is said to be* Lipschitz *and* c *is known as a* Lipschitz constant *for* f*.)*

Invariant Curve of Pairs Attracted

The above result continues to hold for pairs that are attracted by the fixed point in the following sense.

Definition 46 (Asymptotic Self-similarity)
A pair of C^1 *functions* (E_0, F_0) *is called* asymptotically self-similar *if there exists a sequence of pairs* (E_j, F_j) *and a sequence of (complex anti-linear) rescaling maps* B_{j+1} *such that*

$$\begin{aligned} E_{j+1} &= B_{j+1} F_j B_{j+1}^{-1}, && (4.3a)\\ F_{j+1} &= B_{j+1} E_j F_j B_{j+1}^{-1}, && (4.3b) \end{aligned}$$

where the sequence of maps converges to (E, F) *with respect to the* C^1 *norm on the domains* Δ_E*, respectively* Δ_F *and where the rescaling maps are restricted as given below.*

The above notation has been adopted to be consistent with [Sti93c], as the next section will indicate. Each B_{j+1} may be written in the form

$$B_{j+1}(z) = C\frac{1}{\alpha_{j+1}}(z),$$

where each α_{j+1} is a complex rescaling factor of absolute value strictly *less* than 1. (In fact, the proofs could still be carried through if *finitely* many of the rescaling factors have absolute value *greater* than (or equal to) 1. It is only important that the condition holds eventually.) Note that, in the above definition, the rescaling maps α_j are also required to converge.

In terms of the above definition, the following result is established.

Theorem 47 (Invariant Curve for Pairs Attracted)
Let (E_0, F_0) *be an asymptotically self-similar pair. Then there exists a Hölder continuous injective curve* $\chi_0 : [-\omega, 1-\omega] \to \mathbb{C}$ *such that*

$$E_0\chi_0(t) = \chi_0(t-1) \quad \text{for } 1-\omega \le t \le 1, \tag{4.4a}$$

$$F_0\chi_0(t) = \chi_0(t+\omega) \quad \text{for } -\omega \le 1 \le 1-\omega, \tag{4.4b}$$

$$\chi_0(0) = 0. \tag{4.4c}$$

This curve is not differentiable on a dense set of points.

Invariant Jordan Curve of Functions Attracted

Finally, the above theorems are extended to functions (rather than pairs) that are *attracted* by the fixed point in the following sense:

Definition 48 (Attraction of functions to the fixed point)
Consider a single analytic function $f(z)$. *The function* f *is said to be* attracted by the fixed point *if the pair* $(\mathbf{I}, f)$ *is asymptotically self-similar.*

(Here, $\mathbf{I}$ is the identity function.) It is worth noting that asymptotic self-similarity is a weaker condition than convergence to the fixed point under iteration of the renormalisation operator, since there is no explicit normalisation condition imposed on the rescalings (the necklace construction was designed with portability in mind).

Theorem 49 (Invariant Curve for Functions Attracted)
Let f *be an analytic function, defined on an arbitrary simply connected open domain. If* f *is attracted by the fixed point, then* f *has an invariant Jordan domain. On this domain, the map is a homeomorphism and, moreover, it is analytically conjugate to a rigid rotation by the angle* $2\pi m\omega$, *where* m *is an integer different from zero. The natural parameterization of the boundary curve is Hölder continuous; it is not differentiable on a dense set of points.*

In particular, if an analytic function is attracted by the fixed point, it has a Siegel disc bounded by a Hölder continuous non-smooth Jordan curve.

Corollary 50
(cf. [MN83, Conjecture 1]) For the class of maps in the above theorem, the associated Schröder series converges on its circle of convergence, and the image of the circle of convergence is a non-smooth Jordan curve passing through a critical point of the map.

(It had already been proved by Herman [Her85] that for rational functions with degree at least two, the boundary of Siegel discs with Diophantine rotation number always contains a stationary point of the map.)

The next section will summarise the necklace construction, on which the proofs of the above results are based. Then the necklace construction will be adapted for use here in section 4.3, where some modifications will be given. The conceptual parts of the proofs for the above results will be given in section 4.4. Finally, section 4.5 will give the computer part of the proof (which shows that the necessary hypotheses actually hold).

4.2 The Necklace Construction

The aim of this section is to give a brief summary of the *necklace construction* presented in [Sti93c]. The construction was developed for twist maps and used the corresponding renormalization operator (recall section 1.2.3). By making suitable ammendments and verifying some extra conditions, it will be used in the context of the Siegel disc renormalization operator to establish the results given in the previous section. In order that the proofs of these results make sense, it will be necessary to give some details of the method of proof of the necklace construction. We will then show how these details are translated into the framework of Siegel disc renormalization.

The key idea is that by verifying some (purely topological) conditions concerning the action of the maps of the fixed point, it is possible to deduce the existence of a transitive invariant curve for all functions attracted by it.

The proof is constructive: it proceeds iteratively, beginning with the so-called *zeroth generation*, which corresponds to the initial domains of the maps of the fixed point. The various topological conditions ensure, for example, that subsequent 'generations' are contained within their predecessors and that the set constructed at each stage contains no 'gaps'. This iterative method of proof is essentially equivalent to treating the invariant set as the attractor of an *iterated function system* (IFS), and this approach will be taken in chapter 5 in order to yield results concerning the dimension of Siegel disc boundaries.

4.2.1 Hypotheses and Results

In what follows, the notation adopted for twist maps in [Sti93c] will be used, with one modification: the calligraphic letters $\mathcal{U}$, $\mathcal{T}$, and $\mathcal{B}$ will be used to distinguish the maps and rescalings used there from those that will be used in the Siegel disc case (namely U, V, and B). The clash of notation is unfortunate, but we retain it to be consistent with the existing literature. (Occasionally, when it is clear from context which maps are being used, and in computer output, we may drop the calligraphic style.)

Definition 51 (c.f. [Sti93c])
Let $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{\mathcal{T}}$ be connected open subsets of the plane. A pair $(\mathcal{U}, \mathcal{T})$ of homeomorphisms (defined on $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{\mathcal{T}}$, respectively), will be called a fixed point *or* self-similar *if it satisfies the* fixed point equation

$$\mathcal{U} = \mathcal{B}\mathcal{T}\mathcal{B}^{-1}, \tag{4.5a}$$

$$\mathcal{T} = \mathcal{B}\mathcal{U}\mathcal{T}\mathcal{B}^{-1}, \tag{4.5b}$$

where $\mathcal{B}$ is a (real) linear-diagonal map with diagonal elements of absolute value strictly greater than 1.

Note the composition order, $\mathcal{U}\mathcal{T}$, of the maps in the above definition. This is the so-called *accretive composition order*; it is the opposite of that used in the existence proof for Siegel disc renormalization given in chapter 2. (The reason for using this order will be made clear in section 4.2.3.) In order to apply the necklace construction to the Siegel disc case,

it is necessary to find a fixed point that works with this composition order. Fortunately, this can be done by taking an *analytic continuation* of the maps of the fixed point found earlier to new domains, such that the above form of the fixed point equation is satisfied. The details of this are given in section 4.3.2

Four groups of assumptions are then placed on the fixed point in order to obtain the main results of [Sti93c]. These assumptions will be referred to as *domain extension*, *connectedness*, *contractivity*, and *disjointness*.

The first two groups (domain extension and connectedness) are sufficient to guarantee the existence of an invariant set:

Definition 52 (Domain Extension Conditions)

$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{U}}} \subset \mathcal{D}_{\mathcal{T}}, \tag{4.6a}$$
$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \subset \mathcal{D}_{\mathcal{T}}, \tag{4.6b}$$
$$\mathcal{T}\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \subset \mathcal{D}_{\mathcal{U}}. \tag{4.6c}$$

Verification of the corresponding conditions for the Siegel disc operator (at the accretive fixed point) is discussed in 4.5.5.

Note that the second domain extension condition (equation 4.6b) implies that the right-hand side of the fixed point equation for $\mathcal{U}$ (equation 4.5a) is well-defined on $\mathcal{D}_{\mathcal{T}}$. Thus, $\mathcal{U}$ may be extended to a homeomorphism on the union $\mathcal{D}_{\mathcal{U}} \cup \mathcal{D}_{\mathcal{T}}$.

Definition 53 (Connectedness Conditions)

$$\mathcal{D}_{\mathcal{U}} \cap \mathcal{D}_{\mathcal{T}} \neq \emptyset, \tag{4.7a}$$
$$\mathcal{U}\mathcal{D}_{\mathcal{T}} \cap \mathcal{D}_{\mathcal{T}} \neq \emptyset, \tag{4.7b}$$
$$\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}} \cap \mathcal{B}^{-1}\mathcal{D}_{\mathcal{U}} \neq \emptyset. \tag{4.7c}$$

Intuitively, the connectedness assumptions ensure that there will be no 'gaps' in the invariant set constructed.

Using the above conditions, the following result is obtained ([Sti93c, Theorem 5.17]):

Theorem 54
If a homeomorphism of the cylinder is attracted by a fixed point that satisfies the assumptions concerning domain extension and connectedness, then it has a compact, connected, separating invariant set.

(The term *separating* is taken to mean that the complement of the set on the cylinder has exactly two unbounded vertical components.) A single map f is said to be attracted if the pair $(\mathrm{f}, \mathbf{I})$ is attracted (where, if we are talking in terms of the lift of the map from the cylinder to the plane, we insist that f commutes with backward rotation).

In order to establish *transitivity* of the resulting invariant set, it is necessary to impose the third assumption, namely that of *contractivity*:

Definition 55 (Contractivity Condition)
The map

$$\mathcal{T}\mathcal{B}^{-2}, \tag{4.8}$$

is required to be a uniform contraction on $\mathcal{D}_{\mathcal{U}} \cup \mathcal{D}_{\mathcal{T}}$.

Finally, the fourth set of assumptions, namely *disjointness*, enables a parameterisation of the invariant curve to be constructed. From this it is also possible to deduce that the conjugation function is Hölder continuous (where the Hölder exponent is obtained from bounds on the contractivity condition).

Definition 56 (Disjointness Conditions)

$$\mathcal{U}\overline{\mathcal{D}_{\mathcal{T}}} \cap \overline{\mathcal{D}_{\mathcal{U}}} = \emptyset, \tag{4.9a}$$
$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \cap \mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} = \emptyset, \tag{4.9b}$$
$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \cap \mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{U}}} = \emptyset, \tag{4.9c}$$
$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{U}}} \cap \mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{U}}} = \emptyset. \tag{4.9d}$$

Yielding the following theorem ([Sti93c, Theorem 7.34]).

Theorem 57
If a homeomorphism of the cylinder is attracted by a fixed point that satisfies the assumptions concerning domain extension, connectedness, contractivity, and disjointness, then it has an invariant golden circle, and the induced mapping on the curve is continuously conjugate to a rigid rotation. In particular, the invariant curve is Hölder continuous. For maps attracted by the critical fixed point, the conjugator is not differentiable on a dense set of points.

In order to use these results in the context of Siegel disc renormalization, it is necessary to amend some of the proofs and to impose some extra conditions. The following subsections summarise the important aspects of the proofs that will be used.

4.2.2 Domain Pairs

The necklace construction works by constructing a pair of sets that is, in some sense, *invariant* for the fixed point. In [Sti93c, §1] these sets are given as the limit of a recursively defined sequence of pairs of sets that are referred to as *domain pairs*, namely

$$(\mathcal{M}_0, \mathcal{N}_0) = (\mathcal{D}_{\mathcal{U}},\ \mathcal{D}_{\mathcal{T}}), \tag{4.10a}$$
$$(\mathcal{M}_{j+1}, \mathcal{N}_{j+1}) = (\mathcal{T}\mathcal{B}^{-1}\mathcal{N}_j,\ \mathcal{B}^{-1}(\mathcal{M}_j \cup \mathcal{N}_j)). \tag{4.10b}$$

The above definition is then used to define a sequence of sets $(\mathcal{L}_j)$ where

$$\mathcal{L}_j = \mathcal{M}_j \cup \mathcal{N}_j.$$

(Note that the above definition will be used as the basis for an Iterated Function System which will be used to explore the fractal dimension of the necklace curve in chapter 5.)

From the domain extension conditions, it follows [Sti93c, Lemmas 1.1–1.3] that the sequences $(\mathcal{M}_j)$ and $(\mathcal{N}_j)$ are decreasing sequences of non-empty compact sets, and that $\mathcal{M}_j \cup \mathcal{N}_j \neq \emptyset$ for all non-negative integers j. Additionally, by imposing the connectedness conditions, it follows [Sti93c, Lemma 1.4] that the sets $\mathcal{M}_j$ and $\mathcal{N}_j$ are connected.

[Sti93c, Lemma 1.5] then established that the sets $\mathcal{L}_j$ satisfy the recurrence relation:

$$\mathcal{L}_{j+2} = \mathcal{B}^{-1}\mathcal{L}_{j+1} \cup \mathcal{T}\mathcal{B}^{-2}\mathcal{L}_j, \quad j \geq 0. \tag{4.11}$$

The *limit pair* of the sequence of pairs $(\mathcal{M}_j, \mathcal{N}_j)$ is defined by

$$(\mathcal{M}, \mathcal{N}) = \left(\bigcap_j \mathcal{M}_j, \bigcap_j \mathcal{N}_j \right).$$

It follows [Sti93c, Lemma 1.6] that, provided a fixed point of the twist map operator satisfies both the domain extension and the connectedness conditions, then the limit set of the domain pairs is non-empty, compact, and connected (in the sense that the union $\mathcal{L} = \mathcal{M} \cup \mathcal{N}$ is non-empty, compact, and connected).

4.2.3 Fibonacci Strings

In order to prove that the limit pair obtained above is (in a suitable sense) *invariant* for the maps of the fixed point, an explicit form was given [Sti93c, §2] for the sequence of domain pairs. The representation is in terms of *Fibonacci strings*, which are defined below. (The details of the Fibonacci strings will be needed in some of the amendments that we make to the proofs presented in [Sti93c].)

Iterating the fixed point equation formally with the accretive composition order gives rise to sequences (or *strings*) of the maps $\mathcal{U}$ and $\mathcal{T}$ that satisfy a simple recursion relation. For example, the second iterate of the fixed point equation gives

$$\begin{aligned} \mathcal{U} &= \mathcal{B}^3\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{B}^{-3} \\ \mathcal{T} &= \mathcal{B}^3\mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{B}^{-3}. \end{aligned}$$

The sequences of $\mathcal{U}$ and $\mathcal{T}$ on the right-hand side are called *Fibonacci strings*, and may be defined formally as follows:

Definition 58 (Fibonacci Strings)
Let Σ act on strings of symbols by performing the following substitutions:

$$\Sigma : \mathcal{U} \mapsto \mathcal{T}, \tag{4.12a}$$

$$\Sigma : \mathcal{T} \mapsto \mathcal{U}\mathcal{T}. \tag{4.12b}$$

Then the j-th Fibonacci string, $\mathcal{F}_j$, is given by

$$\mathcal{F}_{-1} = \mathcal{U}, \tag{4.13a}$$

$$\mathcal{F}_j = \Sigma(\mathcal{F}_{j-1}) \quad \text{for } j \geq 0. \tag{4.13b}$$

An equivalent definition is the following:

$$\begin{aligned} \mathcal{F}_{-1} &= \mathcal{U}, && (4.14a)\\ \mathcal{F}_0 &= \mathcal{T}, && (4.14b)\\ \mathcal{F}_j &= \mathcal{F}_{j-2}\mathcal{F}_{j-1}, \quad \textit{for } j \geq 1. && (4.14c) \end{aligned}$$

For example, the first few Fibonacci strings are as follows:

$$\begin{aligned} \mathcal{F}_{-1} &= \mathcal{U},\\ \mathcal{F}_0 &= \mathcal{T},\\ \mathcal{F}_1 &= \mathcal{U}\mathcal{T},\\ \mathcal{F}_2 &= \mathcal{T}\mathcal{U}\mathcal{T},\\ \mathcal{F}_3 &= \mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T}. \end{aligned}$$

One further definition will be needed, namely that of *Fibonacci tails*:

Definition 59 (Fibonacci Tails)
For a given Fibonacci string, $\mathcal{F}_j$, any string $\mathcal{X}$ such that

$$\mathcal{Y}\mathcal{X} = \mathcal{F}_j \quad \textit{for some non-empty string } \mathcal{Y},$$

will be called a Fibonacci tail *of $\mathcal{F}_j$. For a given $\mathcal{F}_j$, the set of all such tails will be denoted $\{\mathcal{F}_j\}$.*

Thus, the Fibonacci tails consist of the trailing symbols of a Fibonacci string, including the empty-string but excluding the full Fibonacci string itself. For example, the Fibonacci tails of the first few strings are as follows (where $\Diamond$ denotes the empty string, which is naturally regarded as the identity when the strings are viewed as compositions of maps):

$$\begin{aligned} \{\mathcal{F}_{-1}\} &= \{\Diamond\},\\ \{\mathcal{F}_0\} &= \{\Diamond\},\\ \{\mathcal{F}_1\} &= \{\Diamond, \mathcal{T}\},\\ \{\mathcal{F}_2\} &= \{\Diamond, \mathcal{T}, \mathcal{U}\mathcal{T}\},\\ \{\mathcal{F}_3\} &= \{\Diamond, \mathcal{T}, \mathcal{U}\mathcal{T}, \mathcal{T}\mathcal{U}\mathcal{T}, \mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T}\}. \end{aligned}$$

The reasons for using the accretive composition order $\mathcal{U}\mathcal{T}$ in the fixed point equation (rather than the 'inverse' order $\mathcal{T}\mathcal{U}$) are indicated by the following comments:

Remark 60 (One Fibonacci Tail of Each Length)
Any Fibonacci tail of $\mathcal{F}_j$ is automatically a Fibonacci tail of $\mathcal{F}_{j+1}$. There is exactly one tail of any given length.

It follows that for any two distinct Fibonacci strings (n.b. *strings* rather than *tails*) with $j \geq 0$, the shorter one is a tail of the longer one. (This is the reason behind choosing the

name "accretive" for this composition order; the Fibonacci strings always grow from the left.)

This definition allows the fixed point equations to be rewritten [Sti93c, Lemma 2.5] allowing an explicit form for the sequence of domain pairs to be derived. In particular, it turns out that

$$\mathcal{L}_j = \{\mathcal{F}_j\}\mathcal{B}^{-j}\mathcal{D}_\mathcal{T} \cup \{\mathcal{F}_{j-1}\}\mathcal{B}^{-j}\mathcal{D}_\mathcal{U}, \quad \text{for } j \geq 0. \tag{4.15}$$

Definition 61 (Fibonacci Patches)
The sets of the form $\mathcal{F}\mathcal{B}^{-j}\mathcal{D}_\mathcal{U}$ and $\mathcal{F}\mathcal{B}^{-j}\mathcal{D}_\mathcal{T}$ in the above equation (4.15), where $\mathcal{F}$ signifies any single Fibonacci tail, will be called Fibonacci patches.

For example, the *zeroth generation* of the construction is given by:

$$\mathcal{L}_0 = \{\mathcal{F}_0\}\mathcal{B}^{-0}\mathcal{D}_\mathcal{T} \cup \{\mathcal{F}_{-1}\}\mathcal{B}^{-0}\mathcal{D}_\mathcal{U} = \diamondsuit\mathcal{B}^{-0}\mathcal{D}_\mathcal{T} \cup \diamondsuit\mathcal{B}^{-0}\mathcal{D}_\mathcal{U} = \mathcal{D}_\mathcal{T} \cup \mathcal{D}_\mathcal{U}.$$

(Recall that the empty string corresponds to application of the identity map.) Thus the zeroth generation consists of the two Fibonacci patches $\mathcal{D}_\mathcal{T}$ and $\mathcal{D}_\mathcal{U}$, which agrees with the original definition of the domain pairs (equation 4.10). Similarly, the second generation is given by

$$\begin{aligned}
\mathcal{L}_2 &= \{\mathcal{F}_2\}\mathcal{B}^{-2}\mathcal{D}_\mathcal{T} \cup \{\mathcal{F}_1\}\mathcal{B}^{-2}\mathcal{D}_\mathcal{U} \\
&= \{\diamondsuit, \mathcal{T}, \mathcal{U}\mathcal{T}\}\mathcal{B}^{-2}\mathcal{D}_\mathcal{T} \cup \{\diamondsuit, \mathcal{T}\}\mathcal{B}^{-2}\mathcal{D}_\mathcal{U} \\
&= \mathcal{B}^{-2}\mathcal{D}_\mathcal{T} \cup \mathcal{T}\mathcal{B}^{-2}\mathcal{D}_\mathcal{T} \cup \mathcal{U}\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_\mathcal{T} \cup \mathcal{B}^{-2}\mathcal{D}_\mathcal{U} \cup \mathcal{T}\mathcal{B}^{-2}\mathcal{D}_\mathcal{U},
\end{aligned}$$

which consists of the five patches $\mathcal{B}^{-2}\mathcal{D}_\mathcal{T}$, $\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_\mathcal{T}$, $\mathcal{U}\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_\mathcal{T}$, $\mathcal{B}^{-2}\mathcal{D}_\mathcal{U}$, and $\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_\mathcal{U}$.

In fact, the explicit form for the domain pairs may be further elaborated, by defining the *next symbol* of a Fibonacci tail:

Definition 62 (Next Symbol ν)
Let $\mathcal{F}$ be any Fibonacci tail. Since there is exactly one tail of each length, it follows that there is a unique tail $\mathcal{G}$ that is longer than $\mathcal{F}$ by one symbol. The next symbol $\nu(\mathcal{F})$ *of $\mathcal{F}$ is then defined by*

$$\nu(\mathcal{F}) = \mathcal{T} \quad \text{if } \mathcal{G} = \mathcal{T}\mathcal{F} \tag{4.16a}$$

$$\nu(\mathcal{F}) = \mathcal{U} \quad \text{if } \mathcal{G} = \mathcal{U}\mathcal{F}. \tag{4.16b}$$

By defining $\{\mathcal{F}_j\}_\mathcal{U}$ to be the set containing those tails from $\{\mathcal{F}_j\}$ that have next symbol $\mathcal{U}$, and $\{\mathcal{F}_j\}_\mathcal{T}$ to be those which have next symbol $\mathcal{T}$, it is possible to give the sets $\mathcal{M}_j$ and $\mathcal{N}_j$ (whose union is $\mathcal{L}_j$) in the following explicit form:

$$\mathcal{M}_j = \{\mathcal{F}_j\}_\mathcal{U}\mathcal{B}^{-j}\mathrm{D}_\mathcal{T} \cup \{\mathcal{F}_{j-1}\}_\mathcal{U}\mathcal{B}^{-j}\mathrm{D}_\mathcal{U} \tag{4.17a}$$

$$\mathcal{N}_j = \{\mathcal{F}_j\}_\mathcal{T}\mathcal{B}^{-j}\mathrm{D}_\mathcal{T} \cup \{\mathcal{F}_{j-1}\}_\mathcal{T}\mathcal{B}^{-j}\mathrm{D}_\mathcal{U}. \tag{4.17b}$$

Thus, the set $\mathcal{M}_j$ consists of the Fibonacci patches whose next symbol is $\mathcal{U}$, and the set $\mathcal{N}_j$ consists of the Fibonacci patches whose next symbol is $\mathcal{T}$.

Armed with this explicit representation of the domain pairs in terms of Fibonacci patches, the next step is to find a set that is (in some sense) *invariant* for the fixed point. It turns out to be easier to work with invariant *pairs* of sets, as indicated by the following definition.

Definition 63 (Invariant Pairs of Sets)
A pair $(\mathcal{M},\mathcal{N})$ *of sets is said to be* invariant *for a pair* $(\mathcal{U},\mathcal{T})$ *of maps if the following equations hold,*

$$\mathcal{U}\mathcal{M} \subseteq \mathcal{M}\cup\mathcal{N}, \tag{4.18a}$$
$$\mathcal{T}\mathcal{N} \subseteq \mathcal{M}\cup\mathcal{N}. \tag{4.18b}$$

Using this definition, it was established in [Sti93c, §3] that the limit pair $(\mathcal{M},\mathcal{N})$ of the domain pairs is itself an invariant pair for the fixed point $(\mathcal{U},\mathcal{T})$ and that, furthermore, the set $\mathcal{L}$ contains the orbit of the origin under iteration of the maps of the fixed point. The 'orbit' of the origin is the point set obtained by applying (the maps corresponding to) the set of *all* Fibonacci tails to the point $(0,0)$.

In other words, a fixed point of the accretive renormalisation operator that satisfies the conditions of *domain extension* and *connectedness* has a connected compact invariant pair of sets containing the orbit of the origin.

4.2.4 Asymptotic Self-Similarity

[Sti93c, §4] extended the above result to hold in a weaker context: rather than considering the fixed point itself, the key concept is to look at *asymptotically self-similar* pairs of maps.

Definition 64 (Asymptotic Self-similarity)
A pair of maps $(\mathcal{U}_0,\mathcal{T}_0)$*, defined on domains* $\mathcal{D}_\mathcal{U}$ *and* $\mathcal{D}_\mathcal{T}$ *respectively, is called* asymptotically self-similar *if there exists a sequence of pairs* $(\mathcal{U}_j,\mathcal{T}_j)$ *(defined on the same domains) and a sequence of linear-diagonal rescaling maps* $\mathcal{B}_{k+1}$ *such that*

$$\mathcal{U}_{j+1} = \mathcal{B}_{j+1}\mathcal{T}_j\mathcal{B}_{j+1}^{-1}, \tag{4.19a}$$
$$\mathcal{T}_{j+1} = \mathcal{B}_{j+1}\mathcal{U}_j\mathcal{T}_j\mathcal{B}_{j+1}^{-1}, \tag{4.19b}$$

where the sequence of maps converges on the domains $\mathcal{D}_\mathcal{U}$*, respectively* $\mathcal{D}_\mathcal{T}$*.*

By *convergence*, it is meant that the sequences $(\mathcal{U}_j)$ and $(\mathcal{T}_j)$ converge with respect to the C^1 norm on their domains, and that the diagonal elements of the rescaling maps $\mathcal{B}_{j+1}$ also converge. In particular, it is required that the diagonal elements should converge to numbers of absolute value strictly *greater* than 1 so that the inverse rescalings are contractions.

Remark 65 (Correction)
Note that the corresponding definition in [Sti93c, §0,p. 371] is incorrect: the proof requires C^1 *convergence, not just* C^0 *convergence.*

The idea is that, provided the conditions of domain extension and connectedness are satisfied by the limit pair of such a sequence (which corresponds to a fixed point of renormalization), then they will also hold for all pairs $(\mathcal{U}_j, \mathcal{T}_j)$ in the sequence, where $j \geq k$ for a high enough index k. The key tool used to establish this is the concept of an *approximately invariant sequence* of pairs of sets. The limit pair of such a sequence is invariant for an asymptotically self-similar pair of maps.

Definition 66 (Approximately Invariant Pairs of Sets)
A decreasing sequence of pairs of sets $(\mathcal{M}_j, \mathcal{N}_j)$ *is called* approximately invariant *for a pair* $(\mathcal{U}, \mathcal{T})$ *of maps if, for each* j*, there exists a corresponding* $k \geq j$ *such that*

$$\mathcal{U}\mathcal{M}_k \subset \mathcal{M}_j \cup \mathcal{N}_j, \tag{4.20a}$$

$$\mathcal{T}\mathcal{N}_k \subset \mathcal{M}_j \cup \mathcal{N}_j. \tag{4.20b}$$

(By *decreasing* it is meant that the sequences $(\mathcal{M}_j)$ and $(\mathcal{N}_j)$ decrease separately.)

It turns out [Sti93c, Theorem 4.13] that an asymptotically self-similar pair of maps $(\mathcal{U}_0, \mathcal{T}_0)$, for which the conditions of domain extension and connectedness are eventually satisfied (in the sense defined above), has a connected compact invariant pair of sets containing the orbit of the origin.

4.2.5 Extending the Invariant Set

The above results yielded a compact, connected, invariant pair of sets for a certain iterated image of an asymptotically self-similar pair of maps under renormalisation. The question that now arises is what this means for the original maps themselves. [Sti93c, §5] extends the invariant set to a larger set that is invariant for the original maps. (In the case of twist maps of the cylinder, this invariant set goes around the cylinder.) This invariant set is the *necklace*.

A problem arises when using the extension procedure for maps attracted to the critical Siegel disc fixed point. This will be discussed in section 4.4.3.

4.2.6 Ordering the patches

[Sti93c, §7] uses the additional assumption of *contractivity* at the fixed point, to demonstrate that the orbit of the origin is dense in the limit set and, consequently, that the invariant curve is transitive.

Further, using the additional assumption of *disjointness*, it is possible to define an *ordering* on the Fibonacci patches making up the explicit representation of the domain pairs. This ordering is such that (1) only neighbouring patches (in the ordering) overlap (as sets in the plane), and (2) the ordering is refined in successively higher generations of the construction.

Each patch may be identified by a set of *attributes*, these include:

- The Fibonacci tail $\mathcal{F}$.

- The generation index j.
- A symbol, called the *tag*, which is either $\mathcal{U}$ (if the patch is of the form $\mathcal{F}\mathcal{B}^{-\mathsf{j}}\mathcal{D}_{\mathcal{U}}$) or $\mathcal{T}$ (if the patch is of the form $\mathcal{F}\mathcal{B}^{-\mathsf{j}}\mathcal{D}_{\mathcal{T}}$).

The first step is to define an ordering on the patches in terms of their Fibonacci tails:

Definition 67 (The Function ρ)
Let ω be the golden mean. For any Fibonacci tail $\mathcal{F}$, let q denote the number of times that the symbol $\mathcal{T}$ occurs and let p denote the number of times that the symbol $\mathcal{U}$ occurs. Then the map ρ is defined by

$$\rho(\mathcal{F}) = \omega\mathsf{q} - \mathsf{p}.$$

(The map ρ is injective, since ω is an irrational number.) Given two Fibonacci tails, $\mathcal{F}$ and $\mathcal{G}$, we say that $\mathcal{F} < \mathcal{G}$ if and only if $\rho(\mathcal{F}) < \rho(\mathcal{G})$. This relation will be used to define an ordering on the Fibonacci patches within each generation of the construction.

Notice that two patches within a generation may have the same Fibonacci tail, and thus the same value of ρ. For example, consider the first generation ($\mathsf{j} = 1$): this consists of the patches $\mathcal{T}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$, $\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$, and $\mathcal{B}^{-1}\mathcal{D}_{\mathcal{U}}$. Notice that the last two patches have the same Fibonacci tail (namely, the empty string) and thus have the same value for ρ (namely 0).

In order to remove this potential ambiguity in the ordering of the patches, the following definition is made:

Definition 68 (Natural Patch Ordering)
Let P and Q be distinct patches within the jth generation, having Fibonacci tails $\mathcal{F}$ and $\mathcal{G}$, respectively. If the tails are different, then $\mathsf{P} < \mathsf{Q}$ if and only if $\mathcal{F} < \mathcal{G}$. If the tails are the same, then the patch with tag $\mathcal{U}$ is smaller if j is odd, and the patch with tag $\mathcal{T}$ is smaller if j is even.

Examples of this will be given later (a computer program was written to calculate the ordering, along with other relations, and will be discussed in section 4.2.7).

Armed with the ordering defined above, the next step is to define two further relations between the Fibonacci patches:

1. A relation on the patches within each generation. This will be known as *succession*. The *successor* of a patch P will be denoted $\sigma(\mathsf{P})$.
2. A relation in which each patch in a generation is assigned a unique *parent* patch in the previous generation. This will be called *legal descendence*. The *legal parent* of a patch P will be denoted $\pi(\mathsf{P})$.

The reason for the epithet *legal* is because, as will be shown shortly, it is natural to consider two distinct forms of descendence. In order that two will not be confused, the first form will be called *legal* descendence, and the second form will be called *biological* descendence. (The patch referred to in the previous generation will be called the *legal parent*, or just *parent*, in the former case, and the *natural parent* in the latter.)

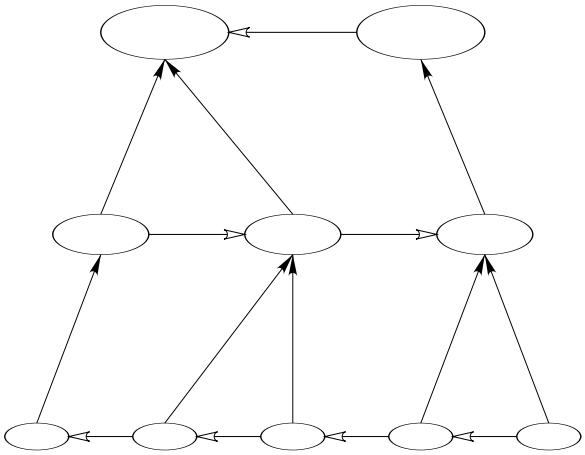

Figure 4.1: Succession and legal descendence in the necklace construction.

Figure 4.1 demonstrates the two main relations (succession and legal descendence) schematically. The patches within each generation are represented by a row of ellipses. The lines pointing upward (with filled arrows) indicate the legal parent of a particular patch. The horizontal lines (with empty arrows) indicate the successor of each patch.

Definition 69 (Succession)
Let P *be a patch of the* j*th generation, different from* $\hat{\mathcal{F}}_{\mathrm{j}}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$ *(where* $\hat{\mathcal{F}}_{\mathrm{j}}$ *is the maximal length tail of* $\mathcal{F}_{\mathrm{j}}$*). Then* P *has a unique successor* $\sigma(\mathsf{P})$*, defined by:*

$$\sigma(\mathcal{F}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{U}}) = \mathcal{F}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}} \tag{4.21a}$$

$$\sigma(\mathcal{F}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}) = \begin{cases} \mathcal{F}\mathcal{F}_{\mathrm{j}-1}\mathcal{B}^{-\mathrm{j}}\mathcal{D}_{\mathcal{T}}, & \text{if } \mathcal{F} \in \{\mathcal{F}_{\mathrm{j}-2}\} \\ (\mathcal{F}-\mathcal{F}_{\mathrm{j}-2})\mathcal{B}^{-1}\mathcal{D}_{\mathcal{U}}, & \text{if } \mathcal{F} \in \{\mathcal{F}_{\mathrm{j}}\}\backslash\{\mathcal{F}_{\mathrm{j}-2}\} \\ & \quad \text{and } \mathcal{F} \neq \hat{\mathcal{F}}_{\mathrm{j}}. \end{cases} \tag{4.21b}$$

Note [Sti93c, Lemma 7.9] that the succession agrees with the patch ordering (introduced in Definition 68) for odd generation index j and opposes it for even j. This fact, and other properties, will be illustrated in section 4.2.7.

The importance of this relation is apparent from the following: provided that the connectedness conditions (equation 4.7) hold, then any two patches within a given generation that are neighbours in the succession (i.e. one is the successor of the other) must overlap as sets.

Definition 70 (Legal Descendence)
For each patch P*, not in the zeroth generation, there exists a unique parent patch* $\pi(\mathsf{P})$ *in the previous generation, defined as follows:*

$$\pi(\mathcal{F}\mathcal{B}^{-(\mathrm{j}+1)}\mathcal{D}_{\mathcal{U}}) = \mathcal{F}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}} \tag{4.22a}$$

$$\pi(\mathcal{F}\mathcal{B}^{-(\mathrm{j}+1)}\mathcal{D}_{\mathcal{T}}) = \begin{cases} \mathcal{F}\mathcal{B}^{-\mathrm{j}}\mathcal{D}_{\mathcal{T}}, & \text{if } \mathcal{F} \in \{\mathcal{F}_{\mathrm{j}}\} \\ (\mathcal{F}-\mathcal{F}_{\mathrm{j}})\mathcal{B}^{-1}\mathcal{D}_{\mathcal{U}}, & \text{if } \mathcal{F} \in \{\mathcal{F}_{\mathrm{j}-1}\}\mathcal{F}_{\mathrm{j}}. \end{cases} \tag{4.22b}$$

Notice that the above definition means that patches tagged by $\mathcal{T}$ give rise to *two* legal children, whereas those tagged by $\mathcal{U}$ give rise to *one* legal child.

From the form of the descendence relation it follows that *each patch is contained (as a set) within its legal parent.* The descendence relation is important because it is used to prove that, under the additional assumption of *disjointness* (equation 4.9), any two patches of a given generation overlap if *and only if* they are neighbours in the succession; in fact, if they are *not* neighbours then even their closures are disjoint.

Remark 71 (Inverses of Fibonacci Tails)
In order to prove that only neighbouring patches overlap (in the sense given above) it is necessary to take the inverse of certain maps represented by Fibonacci tails (see [Sti93c, Lemma 7.15]). For the twist maps proof this was not a problem, since the maps concerned are composed of homeomorphisms. However, in the corresponding proof for Siegel discs, this is not the case. Section 4.5.9 demonstrates how this is overcome, by a finding a certain domain (called a fundamental domain*) on which the corresponding maps are locally homeomorphic, and which contains the Fibonacci patches used in the proof.*

The result is that the ordering relation is refined in successive generations and, in the limit, yields a parameterisation of the invariant curve defined by the $\mathcal{L}_j$. In more detail, consider a point $z \in \mathcal{L}_j$ (where $j \geq 0$): it follows that z is contained in the closure of at least one patch in generation j of the construction, say P, with Fibonacci tail $\mathcal{F}$. It follows from the arguments above that z lies in *at most two* patches and, by a suitable choice of P in the case where there is a tie, it is possible to define a mapping $\psi_j(z) = \rho(P)$ whose limit ψ is well defined as $j \to \infty$. Finally, the inverse of ψ, denoted ϕ, forms the parameterisation of the invariant curve $\mathcal{L}$.

In fact, the above construction yields a parameterization, ϕ, of a *piece* of the curve, namely the part corresponding to $\phi([-\omega, 1])$. In the twist maps proof [Sti93c], this set was extended into an invariant curve which, intuitively speaking, goes *around* the cylinder. In the case of the Siegel disc proof, the above construction yields a piece of the universal invariant curve (i.e. the invariant curve of the critical fixed point). Recall from Chapter 1 that this universal curve is unbounded in the complex plane. The whole curve may be constructed from the piece obtained here by exploiting self-similarity.

More formally, [Sti93c, Lemma 7.28–7.30] yields a continuous injective curve $\phi([-\omega, 1-\omega])$ for the critical fixed point, such that

$$\mathcal{U}\phi(t) = \phi(t-1) \quad \text{for } 1-\omega \leq t \leq 1, \tag{4.23a}$$

$$\mathcal{T}\phi(t) = \phi(t+\omega) \quad \text{for } -\omega \leq 1 \leq 1-\omega, \tag{4.23b}$$

$$\phi(0) = 0. \tag{4.23c}$$

(This shows that the mapping induced on the invariant curve is conjugate to a pure rotation.)

[Sti93c, Lemma 7.31] gives that the curve $t \to \phi(t)$ is Hölder continuous. To get some idea of how this is achieved, note that the necklace construction can be used to obtain a bound on the quantity $|\phi(s) - \phi(t)|$, where s and t are points in the interval $[-\omega, 1] = \phi^{-1}(\mathcal{L})$.

Intuitively, this follows from the fact that we can bound the maximum number of patches in a given generation that $\phi(s)$ and $\phi(t)$ can be separated by. (The corresponding proof in [Sti93c] contains an error, which we correct in section 4.3.6.)

Using the Hölder exponent thus obtained, [Sti93c, Lemma 7.33] demonstrates that the curve is not differentiable at the origin and thus, using transitivity, that it is not differentiable on a dense set of points.

These results are generalised to the case of asymptotically self-similar pairs in [Sti93c, Theorem 7.34].

The following subsection illustrates the necklace construction in more detail, using a computer program written for this project. Then we will demonstrate how the necklace construction, summarised above, is translated into the language of the Siegel disc case. Finally, the (computer-assisted) verification of the corresponding hypotheses will be given.

4.2.7 Illustration of the Necklace Construction

In order to visualise the necklace construction more clearly, a program was written to mimic the construction (including the descendence and succession relations) and to draw diagrams of the result.

In programming terms, we may think of each generation as a singly-linked list of patches, where the linkage is defined by the succession relation. The patches of each generation were constructed in the succession ordering and then linked to their parents in the previous generation using the definition of legal descendence.

In the program, each patch is represented by a `C++` data-structure like the following:

Program Fragment 4.1 Data-structure representing a Fibonacci patch

```
class Patch
{
  private:
    const Generation&  gen;      //reference to generation
    String             tail;     //Fibonacci tail
    char               tag;      //tag to indicate D_U or D_T
    char               nu;       //next symbol (nu)
    int                ord;      //position in the ordering

    Patch*             prev;     //previous patch in succession
    Patch*             next;     //next patch in succession
    Patch*             parent;   //legal parent

    Patch*             natural;  //natural parent
    String             lineage;  //path through M and N
  ...
};
```

(Note that the next symbol ν can be computed from the Fibonacci tail and is, therefore,

redundant. The same is true of the pointers to the parent patches and the position in the natural ordering. However, it makes sense to store these attributes for the sake of efficiency.) Notice the two extra attributes, a pointer to the so-called *natural parent* and the *lineage string*: these will be used to express a second form of descendence, later in this section.

Figure 4.2 shows a schematic view of the necklace construction from the zeroth generation to the fifth. The patches are shown from left to right in the ordering. Neighbouring patches are shown overlapping, and each patch is shown beneath its legal parent patch (note that, in reality, each patch would be *contained within* its legal parent). Arrows show the children of each patch (notice that no two arrows cross each other, which demonstrates graphically that the natural ordering is *refined* in each new generation).

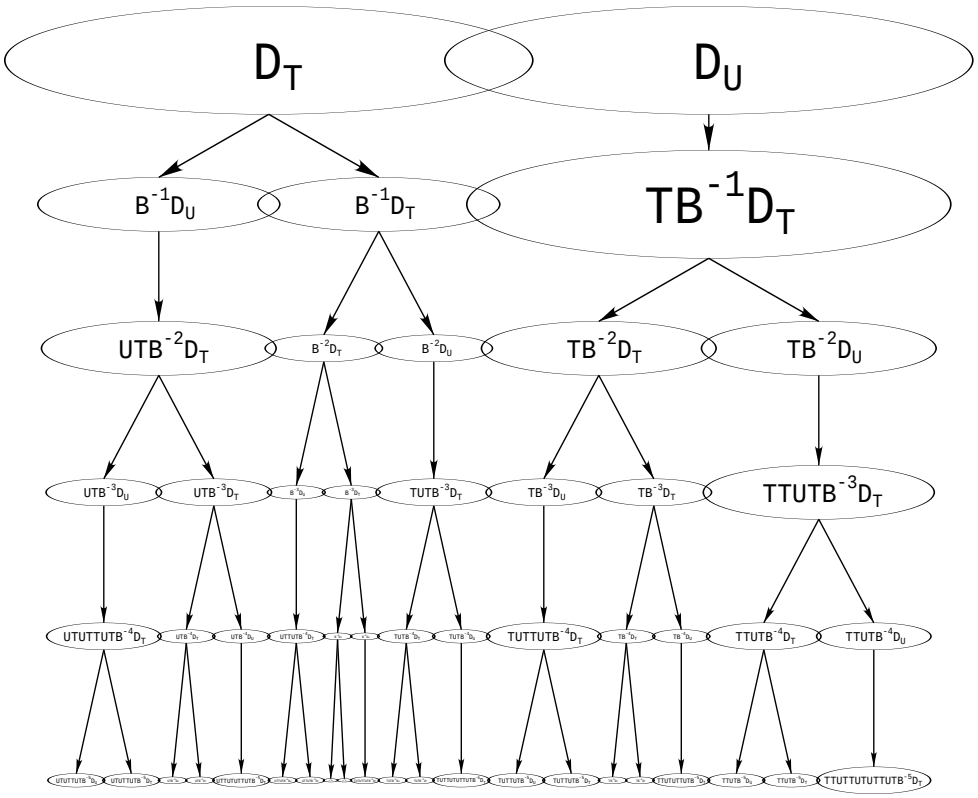

Figure 4.2: Schematic of the generations $j = 0$ to $j = 5$, demonstrating legal descendence.

It will now be useful to define the second form of descendence, before giving detailed output from the program to illustrate the necklace construction.

Definition 72 (Biological/Natural Descendence)
Observe from the equation defining the domain pairs (equation 4.10) that each of the Fibonacci patches is produced from a patch in the previous generation by the application of either the map $\mathcal{B}^{-1}$, or the map $\mathcal{TB}^{-1}$. The patch in the previous generation will be called the biological *(or* natural*) parent.*

From the definition of the domain pairs, patches in the previous generation that are part of $\mathcal{M}_{j-1}$ (having next symbol $\mathcal{U}$) have one natural child, whereas those in $\mathcal{N}_{j-1}$ have two natural children. Accordingly, patches in the current generation which are part of $\mathcal{N}$ have been produced by applying the map $\mathcal{B}^{-1}$ to their natural parent, whereas patches in the current generation that are part of $\mathcal{M}$ are the image of their natural parent under the map $\mathcal{TB}^{-1}$.

Remark 73 (Use of Natural Descendence)
The natural descendence relation is important computationally for calculating high-generation approximations to the necklace set: since application of the map $\mathcal{B}^{-1}$ is trivial (the map is simply a rescaling), there is only one computationally non-trivial map to apply (namely $\mathcal{T}$). Using natural descendence in this way is much more efficient than simply computing an approximation to a patch by applying the maps of its Fibonacci tail (along with suitable rescaling), as the latter would involve compositions of the two 'non-trivial' maps $\mathcal{U}$ and $\mathcal{T}$.

In order to see how the natural parent of a given patch may be calculated, consider a patch in generation $j-1$ of the form $\mathcal{FB}^{-(j-1)}\mathcal{D}$, where $\mathcal{D}$ is either $\mathcal{D}_{\mathcal{U}}$ or $\mathcal{D}_{\mathcal{T}}$. Suppose, on the one hand, that $\mathcal{B}^{-1}$ is applied to produce a patch in $\mathcal{N}_j$ (in what follows, the parentheses are used for clarity):

$$\begin{aligned}
\mathcal{B}^{-1}(\mathcal{FB}^{-(j-1)}\mathcal{D}) &= (\Sigma\Sigma^{-1})\mathcal{B}^{-1}\mathcal{FB}^{-(j-1)}\mathcal{D} && (4.24)\\
&= \Sigma\Sigma^{-1}(\mathcal{B}^{-1}\mathcal{FB}^{-(j-1)})\mathcal{D} && (4.25)\\
&= \Sigma(\Sigma^{-1}(\mathcal{B}^{-1}\mathcal{FB}))\mathrm{B}^{-j}\mathcal{D} && (4.26)\\
&= \Sigma\mathcal{FB}^{-j}\mathcal{D}. && (4.27)
\end{aligned}$$

(Where we have used the fact that application of the symbol substitution Σ is equivalent to conjugation by $\mathcal{B}$.) This shows that given a patch in generation j of the form $\mathcal{GB}^{-j}\mathcal{D}$ with next symbol $\mathcal{T}$, its natural parent can be found as follows:

1. Invert the symbol substitutions performed by the operator Σ to produce a 'DNA-sample' $\mathcal{F} = \Sigma^{-1}\mathcal{G}$.
2. Search the previous generation for a patch whose Fibonacci tail matches the sample, and whose tag matches that of the child. (Recall that the tag indicates whether $\mathcal{D}$ is $\mathcal{D}_{\mathcal{U}}$ or $\mathcal{D}_{\mathcal{T}}$.)

Suppose, on the other hand, that the map $\mathcal{TB}^{-1}$ is applied to $\mathcal{FB}^{-(j-1)}\mathcal{D}$ to produce a patch in $\mathcal{M}_j$, then we have

$$\mathcal{TB}^{-1}\mathcal{FB}^{-(j-1)}\mathcal{D} = \mathcal{T}\Sigma\mathcal{FB}^{-j}\mathcal{D}, \quad (4.28)$$

which follows by the same reasoning as above. Thus, when the next symbol of the child is $\mathcal{U}$, its natural parent can be found by first removing the leading $\mathcal{T}$ from its Fibonacci tail, and then using the same algorithm as above.

Once the natural parent is known it is straightforward to calculate the *lineage* of a patch: the lineage consists of a string containing *exactly* j symbols showing the natural ancestry

of the patch. It is constructed by tracing back up the natural family tree and at each stage adding a symbol L (for *left*) or R (for *right*) to indicate whether the ancestor is part of $\mathcal{N}_j$ or $\mathcal{M}_j$ and, therefore, whether the map $L = \mathcal{B}^{-1}$ or $R = \mathcal{T}\mathcal{B}^{-1}$ must be applied to produce that patch from its natural parent. The lineage string is read from the right to the left (just like the Fibonacci tail) and tells us how to compute the patch efficiently from one of the initial domains.

For an explicit example of how natural descendence and lineage strings relate to the patches of the Fibonacci necklace, consider a patch that has been produced from an initial domain by applying the map $\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{B}^{-5}$ (one of the patches of the 5th generation). We use the fact that for any any Fibonacci tail $\mathcal{F}$, the following identity holds:

$$\mathcal{F} = \Sigma\Sigma^{-1}\mathcal{F} = \mathcal{B}^{-1}\Sigma^{-1}\mathcal{F}\mathcal{B}. \tag{4.29}$$

This gives

$$\begin{aligned}
(\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T})\mathcal{B}^{-5} &= \mathcal{B}^{-1}(\mathcal{U}\mathcal{T}\mathcal{U}\mathcal{T}\mathcal{T}\mathcal{U}\mathcal{T})\mathcal{B}^{-4} \\
&= \mathcal{B}^{-2}\mathcal{T}(\mathcal{T}\mathcal{U}\mathcal{T})\mathcal{B}^{-3} \\
&= \mathcal{B}^{-2}\mathcal{T}\mathcal{B}^{-1}(\mathcal{U}\mathcal{T})\mathcal{B}^{-2} \\
&= \mathcal{B}^{-2}\mathcal{T}\mathcal{B}^{-2}\mathcal{T}\mathcal{B}^{-1} \\
&= \mathcal{B}^{-1}\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1} \\
&= \mathsf{LLRLR},
\end{aligned}$$

Where we have identified the lineage string LLRLR with the corresponding sequence of maps. The parentheses in the above working indicate the string of symbols to which the identity (equation 4.29) is applied at each stage to produce the next line. (Notice that, in the above, we repeatedly work from the left inverting the symbol substitution represented by Σ provided that the string of symbols does not begin with $\mathcal{T}\mathcal{T}$. In that case, inversion of Σ would give a string beginning with $\mathcal{U}\mathcal{U}$ which is not in the image of Σ. Instead, in this case we leave the leading $\mathcal{T}$ alone and continue from the next symbol along.)

j	Q_j	F_j
−1	0	U
0	1	T
1	1	UT
2	2	TUT
3	3	UTTUT
4	5	TUTUTTUT
5	8	UTTUTTUTUTTUT
6	13	TUTUTTUTUTTUTTUTUTTUT

Figure 4.3: Basic Fibonacci numbers and Fibonacci strings

Remark 74 (Length of the Lineage String)
It is important here to note that the relevant lineage string has length j. *The above identity may be applied in a different order to produce a string of different length that*

would not be relevant to the current generation. (This comment is important because each tail occurs in more than one generation.)

Shown in figure 4.4 is the output of the program for the zeroth, first, second and fifth generations ($j = 0, 1, 2, 5$) of the Fibonacci necklace. Firstly, the relevant Fibonacci numbers (Q_j) and strings (F_j) are given (see Figure 4.3). Then the patches are listed in the order of succession. Also shown is the value of ρ (to a few decimal places), the next symbol ν, the legal parent of the patch, the natural parent of the patch, and the lineage string.

Patch	ρ	ν	Legal parent	Natural parent	Lineage
D_U	0.000	$\underline{U}$	none	none	$\diamondsuit$
D_T	0.000	T	none	none	$\diamondsuit$
$B^{-1}D_U$	0.000	T	D_T	D_U	L
$B^{-1}D_T$	0.000	$\underline{T}$	D_T	D_T	L
$TB^{-1}D_T$	0.618	U	D_U	D_T	R
$TB^{-2}D_U$	0.618	U	$TB^{-1}D_T$	$B^{-1}D_U$	RL
$TB^{-2}D_T$	0.618	$\underline{U}$	$TB^{-1}D_T$	$B^{-1}D_T$	RL
$B^{-2}D_U$	0.000	T	$B^{-1}D_T$	$B^{-1}D_U$	LL
$B^{-2}D_T$	0.000	T	$B^{-1}D_T$	$B^{-1}D_T$	LL
$UTB^{-2}D_T$	−0.382	T	$B^{-1}D_U$	$TB^{-1}D_T$	LR
$UTUTTUTB^{-5}D_U$	−0.528	T	$UTUTTUTB^{-4}D_T$	$TTUTB^{-4}D_U$	LRLRL
$UTUTTUTB^{-5}D_T$	−0.528	T	$UTUTTUTB^{-4}D_T$	$TTUTB^{-4}D_T$	LRLRL
$UTB^{-5}D_U$	−0.382	T	$UTB^{-4}D_T$	$TB^{-4}D_U$	LRLLL
$UTB^{-5}D_T$	−0.382	T	$UTB^{-4}D_T$	$TB^{-4}D_T$	LRLLL
$UTTUTUTTUTB^{-5}D_T$	−0.292	T	$UTB^{-4}D_U$	$TUTTUTB^{-4}D_T$	LRLLR
$UTTUTB^{-5}D_U$	−0.146	T	$UTTUTB^{-4}D_T$	$TUTB^{-4}D_U$	LLLRL
$UTTUTB^{-5}D_T$	−0.146	T	$UTTUTB^{-4}D_T$	$TUTB^{-4}D_T$	LLLRL
$B^{-5}D_U$	0.000	T	$B^{-4}D_T$	$B^{-4}D_U$	LLLLL
$B^{-5}D_T$	0.000	T	$B^{-4}D_T$	$B^{-4}D_T$	LLLLL
$TUTUTTUTB^{-5}D_T$	0.090	T	$B^{-4}D_U$	$UTTUTB^{-4}D_T$	LLLLR
$TUTB^{-5}D_U$	0.236	T	$TUTB^{-4}D_T$	$UTB^{-4}D_U$	LLRLL
$TUTB^{-5}D_T$	0.236	T	$TUTB^{-4}D_T$	$UTB^{-4}D_T$	LLRLL
$TUTTUTUTTUTB^{-5}D_T$	0.326	$\underline{T}$	$TUTB^{-4}D_U$	$UTUTTUTB^{-4}D_T$	LLRLR
$TUTTUTB^{-5}D_U$	0.472	U	$TUTTUTB^{-4}D_T$	$TUTB^{-4}D_U$	RLLRL
$TUTTUTB^{-5}D_T$	0.472	U	$TUTTUTB^{-4}D_T$	$TUTB^{-4}D_T$	RLLRL
$TB^{-5}D_U$	0.618	U	$TB^{-4}D_T$	$B^{-4}D_U$	RLLLL
$TB^{-5}D_T$	0.618	U	$TB^{-4}D_T$	$B^{-4}D_T$	RLLLL
$TTUTUTTUTB^{-5}D_T$	0.708	U	$TB^{-4}D_U$	$UTTUTB^{-4}D_T$	RLLLR
$TTUTB^{-5}D_U$	0.854	U	$TTUTB^{-4}D_T$	$UTB^{-4}D_U$	RLRLL
$TTUTB^{-5}D_T$	0.854	U	$TTUTB^{-4}D_T$	$UTB^{-4}D_T$	RLRLL
$TTUTTUTUTTUTB^{-5}D_T$	0.944	U	$TTUTB^{-4}D_U$	$UTUTTUTB^{-4}D_T$	RLRLR

Figure 4.4: Details of patches from generations 0,1,2, and 5.

The following general features may be observed in the above output.

- There are Q_{j+2} patches in generation j.
- The first patch in the succession is $\hat{\mathcal{F}}_{j-1}\mathcal{B}^{-j}\mathcal{D}_{\mathcal{U}}$ (where $\hat{\mathcal{F}}_{j-1}$ is the maximal Fibonacci tail of $\mathcal{F}_{j-1}$). The final patch is $\hat{\mathcal{F}}_j\mathcal{B}^{-j}\mathcal{D}_{\mathcal{T}}$ (where $\hat{\mathcal{F}}_j$ is the maximal Fibonacci tail of $\mathcal{F}_j$).
- The natural ordering (indicated by increasing values of ρ) agrees with the succession ordering for odd generation index j, and opposes it for even generation index.
- In the limit, the first patch in the natural ordering corresponds to $\rho = -\omega$, the patches of the form $\mathcal{B}^{-j}\mathcal{D}_{\mathcal{U}}$ and $\mathcal{B}^{-j}\mathcal{D}_{\mathcal{T}}$ (which contain the origin) correspond to $\rho = 0$, and the final patch in the natural ordering corresponds to $\rho = 1$.
- Each generation is divided into two parts, $\mathcal{M}_j$ and $\mathcal{N}_j$. The patches with next symbol $\nu = \mathcal{U}$ are in the former, those with next symbol $\nu = \mathcal{T}$ are in the latter. The next symbol changes after the *pivotal patch* (marked in the program output by an underlined next symbol ν). In the limit, the pivotal patch corresponds to $\rho = 1 - \omega$. Thus, the piece $\mathcal{N}$ of the invariant curve is given by $\phi([-\omega, 1-\omega])$, and the piece $\mathcal{M}$ is given by $\phi([1-\omega, 1])$. (An intriguing additional observation is that the Fibonacci tail of the pivotal patch is palindromic.)

4.3 Translating the Necklace Construction

In this section, we demonstrate how the necklace construction, as summarised in the preceeding sections, may be successfully translated into the framework of Siegel disc renormalization. Once this has been done (in particular, when the corresponding hypotheses have been given), a computer-assisted proof will be given that the hypotheses of the necklace construction hold for the critical Siegel disc fixed point (it has already been confirmed that they hold for the simple fixed point, and this fact used to prove a special case of Siegel's theorem, in [Sti93a]). The proof will use (the ball containing) the fixed point that was obtained from the existence proof described in Chapter 2.

4.3.1 The Functions U and V

Recall that the components E and F of the fixed point are *even* functions, and that this was built into the existence proof (chapter 2) by writing $E = UQ$, and $F = VQ$ (recall that Q denotes squaring). The output from the existence proof was a pair of analytic functions U and V, defined on discs D_U and D_V, that satisfy the *inverse* fixed point equation, namely

$$U(z) = C\alpha^{-1}V\alpha^2 C \quad \text{on } D_U, \tag{4.30a}$$

$$V(z) = C\alpha^{-1}VQU\alpha^2 C \quad \text{on } D_V, \tag{4.30b}$$

where $\alpha = V(0)$.

Remark 75 (Uncle of the Fixed Point Equation)
The above equation is not really a fixed point equation for either $(E, F) = (UQ, VQ)$ *or* (QU, QV). *Just as for the corresponding equation in chapter 2 (namely equation 2.16), a fixed point equation can be readily deduced from it by squaring. As before, we will call the above equation the* uncle *of the fixed point equation. (Once again, this terminology will prove useful in chapter 5.)*

It is convenient, computationally, to continue to work with the maps U and V, rather than E and F.

4.3.2 Composition Order

A problem arises because the existence proof used the inverse composition order FE (respectively VQU), whereas for the necklace construction, the *accretive* composition order EF (respectively UQV) is required. Fortunately, it is possible to take analytic continuations of the maps U and V to new domains, Δ_U and Δ_V respectively, such that the fixed point equation with the accretive composition order is satisfied, namely

$$U = C\alpha^{-1}V\alpha^2C \quad \text{on } \Delta_U, \tag{4.31a}$$
$$V = C\alpha^{-1}UQV\alpha^2C \quad \text{on } \Delta_V. \tag{4.31b}$$

(As before, the above equation is an 'uncle' of the fixed point equation.) In what follows a slight abuse of notation will be adopted, in that U and V will again be used to denote the maps on the *new* domains Δ_U and Δ_V respectively.

Details of the analytic continuations will now be given (proofs that these continuations can actually be carried out will be given in the computer-assisted part of this chapter).

We will take two analytic continuations of V. They are obtained by substituting the equation (4.30a) for U into the one (4.30b) for V. With the VQU composition order this yields

$$V = C\alpha^{-1}VQC\alpha^{-1}V\alpha^2C\alpha^2C = C\alpha^{-1}VQC\alpha^{-1}V|\alpha|^4 \quad \text{on } D_V.$$

We denote the function on the right hand side by $\hat{V}$. Since equation 4.30b holds at least on D_V, then $\hat{V}$ agrees with V on D_V. Moreover by the domain extension property (which was already verified as a by-product of the existence proof), the natural domain of definition of $\hat{V}$ is larger than D_V (in fact, it is a connected open neighbourhood of the closure of D_V.) Therefore, $\hat{V}$ is an analytic extension of V.

On the other hand, with the fixed point equation (4.31) for the accretive composition order UQV we obtain

$$V = C\alpha^{-1}C\alpha^{-1}V\alpha^2CQV\alpha^2C = |\alpha|^{-2}V\alpha^2CQV\alpha^2C.$$

We denote the function on the right hand side by $\hat{W}$. It turns out the natural domain of definition of $\hat{W}$ is non-empty and connected and intersects the domain of definition of $\hat{V}$ in a connected neighbourhood of the origin. Since the fixed point is commuting, $\hat{W}$ and $\hat{V}$ coincide on the intersection. It follows that $\hat{W}$ is an analytic continuation of V.

To sum up: the three functions V, $\hat{V}$, and $\hat{W}$ agree with each other wherever two of them are defined. Together, they make up one single analytic function, which we denote by $\tilde{V}$. Then

$$\tilde{U} = C\alpha^{-1}\tilde{V}\alpha^2 C$$

is an analytic extension of the original U. (Observe that the natural domain of definition of $\tilde{U}$ is connected, being the pre-image of the domain of definition of $\tilde{V}$ under the map $z \mapsto \alpha^2 C(z)$, and that it contains D_U. $\tilde{U}$ and U agree on D_U, by the fixed point equation.)

The maps U and V on the new domains Δ_U and Δ_V are restrictions of $\tilde{U}$ and $\tilde{V}$ to these domains.

4.3.3 Translation Rules

The analogues of the maps $\mathcal{U}$ and $\mathcal{T}$ in the necklace theory [Sti93c] are $E = UQ$ and $F = VQ$, rather than U and V themselves. Since it is convenient to continue to work with U and V, the necklace theory must be translated. In fact, the maps QU and QV will be used as they are computationally more convenient to implement (the alternative would involve using the "square roots" of the domains). Since the operation of squaring is well-defined everywhere, the maps QU and QV will be defined wherever U and V are (notice also that a fixed point equation for (QU, QV) follows by squaring both sides of the 'uncle'). This means that it will be necessary later to "take the square root" (in a suitable sense) of the necklace curve that results, in order to deduce the existence of a corresponding curve for the "original" maps $E := UQ$ and $F := VQ$.

Fortunately, the translation of the necklace construction to use the maps QU and QV can be done almost mechanically:

- Replace the rescaling map $\mathcal{B}$ by

  $$B(z) = C\alpha^{-2}(z).$$

  (Thus $B^{-1} = \alpha^2 C$.)

- Replace the symbol $\mathcal{U}$ in the necklace theory by QU.

- Replace the symbol $\mathcal{T}$ in the necklace theory by QV.

- Replace $\mathcal{D}_{\mathcal{U}}$ in the necklace theory by Δ_U.

- Replace $\mathcal{D}_{\mathcal{T}}$ in the necklace theory by Δ_V.

The next subsections demonstrate how each part of the construction is translated. In particular, the proofs given in [Sti93c] stay valid, with the exception of the proof of Lemma 7.15, which requires the maps $\mathcal{U}$ and $\mathcal{T}$ to be global homeomorphisms in order that the inverse of certain Fibonacci tails may be taken. The necessary modifications are discussed in the relevant subsection. Essentially, it is required that the map V is a (local) homeomorphism on a certain domain, which will be called the *fundamental domain*.

4.3.4 The Accretive Fixed Point

The necklace construction requires that the following (accretive) fixed point equation holds:

$$\mathcal{U} = \mathcal{B}\mathcal{T}\mathcal{B}^{-1} \quad \text{on } \mathcal{D}_{\mathcal{U}} \tag{4.32a}$$
$$\mathcal{T} = \mathcal{B}\mathcal{U}\mathcal{T}\mathcal{B}^{-1} \quad \text{on } \mathcal{D}_{\mathcal{T}}. \tag{4.32b}$$

Under the above translation rules, the corresponding requirement is that QU and QV obey the following relation:

$$\mathsf{QU} = \mathsf{BQVB}^{-1} \quad \text{on } \Delta_{\mathsf{U}} \tag{4.33a}$$
$$\mathsf{QV} = \mathsf{BQUQVB}^{-1} \quad \text{on } \Delta_{\mathsf{V}}, \tag{4.33b}$$

i.e.

$$\mathsf{U} = \mathsf{C}\alpha^{-2}\mathsf{QV}\alpha^{2}\mathsf{C} \quad \text{on } \Delta_{\mathsf{U}} \tag{4.34a}$$
$$\mathsf{QV} = \mathsf{C}\alpha^{-2}\mathsf{QUQV}\alpha^{2}\mathsf{C} \quad \text{on } \Delta_{\mathsf{V}}. \tag{4.34b}$$

Notice that this fixed point equation is actually a consequence of the earlier (stronger) accretive 'uncle' equation for U and V, namely equation 4.31.

$$\mathsf{U} = \mathsf{C}\alpha^{-1}\mathsf{V}\alpha^{2}\mathsf{C} \quad \text{on } \Delta_{\mathsf{U}} \tag{4.35a}$$
$$\mathsf{V} = \mathsf{C}\alpha^{-1}\mathsf{UQV}\alpha^{2}\mathsf{C} \quad \text{on } \Delta_{\mathsf{V}}. \tag{4.35b}$$

Thus, if the validity of the uncle is established, then the above fixed point equation holds. Recall that the analytic continuations of U and V to the new domains Δ_{U} and Δ_{V} obey the uncle equation. Thus, the accretive fixed point equation needed for the necklace construction will hold provided that the analytic continuations can actually be carried out.

The uncle equation (4.31) will be needed here only once, namely in the modification of [Sti93c, Lemma 7.15] to take into account the fact that the maps we must deal with here are not global homeomorphisms (recall Remark 71). For future convenience the uncle equation is rewritten as follows:

$$\mathsf{U} = \mathsf{B}_1\mathsf{V}\mathsf{B}^{-1} \quad \text{on } \Delta_{\mathsf{U}} \tag{4.36a}$$
$$\mathsf{V} = \mathsf{B}_1\mathsf{UQV}\mathsf{B}^{-1} \quad \text{on } \Delta_{\mathsf{V}}, \tag{4.36b}$$

where $\mathsf{B}_1(z) = \mathsf{C}\alpha^{-1}(z)$.

4.3.5 The Translated Necklace Hypotheses

Translated into the new framework, the necklace hypotheses [Sti93c] look as follows:

- Domain Extension (c.f. equation 4.6):

$$B^{-1}\overline{\Delta_U} \subset \Delta_V, \tag{4.37a}$$
$$B^{-1}\overline{\Delta_V} \subset \Delta_V, \tag{4.37b}$$
$$QVB^{-1}\overline{\Delta_V} \subset \Delta_U. \tag{4.37c}$$

- Connectedness (c.f. equation 4.7):

$$\Delta_U \cap \Delta_V \neq \emptyset, \tag{4.38a}$$
$$QU\Delta_V \cap \Delta_V \neq \emptyset, \tag{4.38b}$$
$$QUB^{-1}\Delta_V \cap B^{-1}\Delta_U \neq \emptyset. \tag{4.38c}$$

- Contractivity (c.f. equation 4.8): The map

$$QVB^{-2}, \tag{4.39}$$

  is required to be a uniform contraction on $\Delta_U \cup \Delta_V$.

- Disjointness (c.f. equation 4.9):

$$QU\overline{\Delta_V} \cap \overline{\Delta_U} = \emptyset \tag{4.40a}$$
$$B^{-1}\overline{\Delta_V} \cap QUB^{-1}\overline{\Delta_V} = \emptyset \tag{4.40b}$$
$$B^{-1}\overline{\Delta_V} \cap QUB^{-1}\overline{\Delta_U} = \emptyset \tag{4.40c}$$
$$B^{-1}\overline{\Delta_U} \cap QUB^{-1}\overline{\Delta_U} = \emptyset. \tag{4.40d}$$

The proofs in [Sti93c, §1] all remain valid under the translation. In particular, the domain extension conditions (equation 4.37, in particular equation 4.37a) imply that the map U (and hence QU) extends to an analytic map on the union of the domains $\Delta_U \cup \Delta_V$.

For the Siegel disc case, the definition of the *domain pairs* [Sti93c, §1] then looks as follows:

$$(M_0, N_0) = (\Delta_U,\ \Delta_V), \tag{4.41a}$$
$$(M_{j+1}, N_{j+1}) = (QV\alpha^2 CN_j,\ \alpha^2 C(M_j \cup N_j)). \tag{4.41b}$$

Similarly, all of the (combinatorial) results obtained for Fibonacci strings remain unchanged by the translation. In particular, [Sti93c, Theorem 3.5] yields a non-empty compact connected invariant pair for the fixed point, that contains the orbit of the origin.

Further, the proof of [Sti93c, Theorem 4.13] is also unchanged, yielding a non-empty compact connected invariant pair that contains the orbit of the origin for any *asymptotically self-similar* pair.

A modification of the proof of [Sti93c, Lemma 7.15], given in section 4.5.9, finally yields a continuous parameterisation for the invariant curve. However, as indicated earlier, we have been working with the maps QU and QV, and the resulting curve is invariant for these maps rather than for $E = UQ$ and $F = VQ$. This problem is dealt with below. First it is necessary to correct an error in [Sti93c].

4.3.6 Global Hölder Exponent

With the translation rules spelled out above, the necklace construction [Sti93c, 7.30] yields a continuous function $\phi = \phi(t)$ with

$$\phi(t-1) = Q\mathcal{U}\phi(t), \quad \text{for } 1-\omega \le t \le 1 \tag{4.42}$$
$$\phi(t+\omega) = Q\mathcal{V}\phi(t), \quad \text{for } -\omega \le 1 \le 1-\omega \tag{4.43}$$

We now need to show that the function ϕ is Hölder continuous. (The corresponding proof for this in [Sti93c, 7.31] is erroneous.) This can be done by looking at the decay rate for the size of the patches as the generation index j increases.

Consider the recursion formula for the sets L_j (recall section 4.2.2):

$$\mathcal{L}_{j+2} = \mathcal{B}^{-1}\mathcal{L}_{j+1} \cup \mathcal{T}\mathcal{B}^{-2}\mathcal{L}_j,$$

and suppose that λ and μ are bounds for the contractivities of the maps $\mathcal{B}^{-1}$ and $\mathcal{T}\mathcal{B}^{-2}$ respectively, i.e.

$$|\mathcal{B}^{-1}(x) - \mathcal{B}^{-1}(y)| \le \lambda|x-y|,$$
$$|\mathcal{T}\mathcal{B}^{-2}(x) - \mathcal{T}\mathcal{B}^{-2}(y)| \le \mu|x-y|.$$

(Where the inequalities hold for all x, y in suitable initial domains.) In fact, the contractivity constant of the inverse rescaling $\mathcal{B}^{-1} : z \mapsto \alpha^2 z$ is just $|\alpha|^2$ (strictly smaller than one).

Now, we put

$$\kappa = \max(\lambda, \mu^{\frac{1}{2}}).$$

This obviously implies that $\lambda \le \kappa$ and $\mu \le \kappa^2$. (Notice that the above square root is erroneously missing in the proof of [Sti93c, 7.31]. It is necessary since the index on the set $\mathcal{T}\mathcal{B}^{-2}\mathcal{L}_j$ lags behind that of $\mathcal{L}_{j+2}$ by one more than does the index on $\mathcal{B}^{-1}\mathcal{L}_{j+1}$.)

Lemma 76 (Decay Rate of Patch Size)
Let d_j be the maximal diameter of a patch in generation j. We claim that there exists a constant c such that

$$d_j \le c\kappa^j.$$

Proof We can make the claim true for $j = 0$ and $j = 1$ by choosing c large enough. Assume now that $j \ge 0$ and that

$$d_j \le c\kappa^j$$
$$d_{j+1} \le c\kappa^{j+1}.$$

Then, by the mean value theorem of differential calculus, we obtain

$$d_{j+2} \le \max(\lambda c\kappa^{j+1}, \mu c\kappa^j) = \max(c\kappa^{j+2}, c\kappa^{j+2}) = c\kappa^{j+2},$$

and the induction is complete. □

So the maximal diameter of the patches of the j-th generation is bounded by $c\kappa^j$ with a positive constant c. Finally, the number

$$\mu := \frac{\log \kappa}{\log \omega},$$

(so that $\kappa = \omega^\mu$) is a global Hölder exponent of ϕ. To see why, fix the generation index $j \geq 1$ and consider two distinct points s and t in the interval $\phi^{-1}(L)$ with $\omega^{j+1} \leq |s-t| < \omega^j$. It follows [Sti93c, Lemma 7.21] that the points $\phi(s)$ and $\phi(t)$ cannot be separated by more than seven patches of generation j, giving (from above) the estimate

$$|\phi(s) - \phi(t)| \leq 9c\kappa^j.$$

Thus

$$\frac{|\phi(s) - \phi(t)|}{|s-t|^\mu} \leq \frac{9c}{\omega^\mu} \left(\frac{\kappa}{\omega^\mu}\right)^j = \frac{9c}{\omega^\mu},$$

i.e. μ is a Hölder exponent for ϕ.

4.3.7 Taking the Square Root

As summarised above, the necklace construction yields a Hölder continuous curve $\phi(t)$ which satisfies

$$QU\phi(t) \equiv \phi(t-1), \quad 1-\omega \leq t \leq 1 \tag{4.44a}$$
$$QV\phi(t) \equiv \phi(t+\omega), \quad -\omega \leq t \leq 1-\omega. \tag{4.44b}$$

The invariant curve ϕ thus defined, however, is not the one we want: It is invariant for the pair (QU, QV). What we would like to have is a curve $\chi(t)$ invariant for $(UQ, VQ) =: (E, F)$, i.e.,

$$UQ\chi(t) \equiv \chi(t-1), \quad 1-\omega \leq t \leq 1 \tag{4.45a}$$
$$VQ\chi(t) \equiv \chi(t+\omega), \quad -\omega \leq t \leq 1-\omega. \tag{4.45b}$$

It turns out that the cure is easy: we "take the square root" of ϕ. In brief, the argument goes as follows: Applying Q to both sides of equation 4.45 yields

$$QUQ\chi(t) \equiv Q\chi(t-1), \quad 1-\omega \leq t \leq 1$$
$$QVQ\chi(t) \equiv Q\chi(t+\omega), \quad -\omega \leq t \leq 1-\omega.$$

and this implies that the curve $Q\chi(t)$ solves equation 4.44. Since the necklace curve is uniquely determined (it is the closure of the orbit of the origin), it follows that, if χ exists at all, then we must have

$$\phi(t) \equiv Q\chi(t) = \chi^2(t),$$

i.e., χ is the "square root" of ϕ.

Lemma 77 (Square Root of ϕ)
Define χ by

$$\chi(t) = \begin{cases} U\phi(t+1) & -\omega \le t \le 0 \\ V\phi(t-\omega) & 0 < t \le 1 \end{cases}$$

We claim that χ is a continuous "square root" of ϕ and, further, that it is invariant for the pair $(E, F) := (UQ, VQ)$.

Proof: (1) The function χ is well-defined and continuous with the possible exception of the origin. At the origin, however, we get

$$\begin{aligned} \chi(0^-) &= U\phi(1) = 0 \\ \chi(0^+) &= V\phi(-\omega) = 0, \end{aligned}$$

and it follows that χ is continuous everywhere.

(2) Now we need to prove that χ is a square root of ϕ. For $-\omega \le t \le 0$, we have $-\omega + 1 \le t + 1 \le 1$, giving

$$Q\chi(t) = QU\phi(t+1) = \phi(t).$$

For $0 \le t \le 1$, on the other hand, we have $-\omega \le t - \omega \le 1 - \omega$, giving

$$Q\chi(t) = QV\phi(t-\omega) = \phi(t).$$

Thus χ is a continuous square root of ϕ.

(3) It remains to show that χ is invariant for (UQ, VQ). For $1 - \omega \le t \le 1$, we have $-\omega \le t - 1 \le 0$, so that

$$UQ\chi(t) = U\phi(t) = U\phi((t-1)+1) = \chi(t-1),$$

For $-\omega \le t \le 1 - \omega$, on the other hand, we have $0 \le t + \omega \le 1$,

$$VQ\chi(t) = V\phi(t) = V\phi((t+\omega) - \omega) = \chi(t+\omega),$$

completing the proof. □

(The relationship between the invariant set for the maps $(E, F) := (UQ, VQ)$, given by the "square root" above, and the set actually obtained by using (QU, QV) will be clarified a little more in section 5.2.1.)

4.3.8 Hölder Exponent of χ at the Origin

In order to establish that the curve χ thus defined is not differentiable on a dense set of points, we use the Hölder exponent obtained above. We note that the contractivity of B^{-1} is $|\alpha|^2$. By the proof of [Sti93c, Lemma 7.33], the number

$$\frac{\log |\alpha|^2}{\log \omega} = 2\frac{\log |\alpha|}{\log \omega}$$

is the optimal local Hölder exponent of the curve ϕ at the origin. Accordingly, half of that number is the optimal Hölder exponent of the square root χ (of ϕ) at the origin. In particular,

$$\frac{\log|\alpha|}{\log\omega}$$

is an upper bound on the global Hölder exponent. This number is smaller than one; the curve χ, therefore, is not differentiable. In fact, it is not differentiable on the whole backward orbit of the origin, which is dense.

4.4 Proof of the Main Results

Taking into account the above modifications, we now prove that the consequences of the necklace construction imply our main results. (The computer verification of the relevant hypotheses will then be given.)

In this section, we are working with the original maps E and F, rather than with U and V. (Recall that $E := UQ$ and $F := VQ$.) We put

$$\Delta_F := \{z : z^2 \in \Delta_V\}.$$

Moreover, we put Δ_E equal to one of the two connected components of the set $\{z : z^2 \in \Delta_U\}$.

Proof of Proposition 43: The existence proof of chapter 2 yielded an inverse fixed point. The verification of the analytic extensions for the maps of the fixed point will be given the computer part of the proof (section 4.5). Together, these imply that a fixed point for the accretive order exists (proposition 43). □

4.4.1 Invariant Curve of the Fixed Point

We now prove Theorem 44. Namely, that for the critical fixed point (E, F) there exists the Hölder continuous injective curve defined by equation 4.2 and that this curve is not differentiable on a dense set of points.

Proof of Theorem 44: By the necklace construction, the pair (QU, QV) has a self-similar transitive invariant curve ϕ which is Hölder continuous and passes through the origin. It satisfies equation 4.23:

$$\begin{aligned} QU\phi(t) &= \phi(t-1) \quad \text{for } 1-\omega \le t \le 1, \\ QV\phi(t) &= \phi(t+\omega) \quad \text{for } -\omega \le 1 \le 1-\omega, \\ \phi(0) &= 0. \end{aligned}$$

The argument presented in section 4.3.7 yields a "square root" χ which is invariant for $(E, F) := (UQ, VQ)$ and which satisfies

$$E\chi(t) = \chi(t-1) \quad \text{for } 1-\omega \le t \le 1, \tag{4.46a}$$

$$F\chi(t) = \chi(t+\omega) \quad \text{for } -\omega \le 1 \le 1-\omega, \tag{4.46b}$$

$$\chi(0) = 0. \tag{4.46c}$$

It now remains to prove equation 4.2c, namely

$$\alpha C\chi(t) = \chi(-\omega t).$$

(This equation is an "extra" one, not present in the original necklace theory. It expresses the exact self-similarity of the universal curve itself. As such, it holds only at the critical renormalization fixed point, not for pairs that are attracted.)

It is enough to prove it for $t = q\omega - p$, where the pair (q, p) characterises an arbitrary Fibonacci tail $\mathcal{F}$ (q and p counting the number of occurrences of the symbols E and F in the string $\mathcal{F}$, respectively).

Let ρ give the value of $t = q\omega - p$ for a given string. The fixed point equation establishes the following identity:

$$\mathcal{B}^{-1}\mathcal{F}\mathcal{B} = \Sigma(\mathcal{F}), \tag{4.47}$$

i.e. the substitution operator Σ (introduced in section 4.2) acts formally by conjugation by $\mathcal{B}^{-1}$.

A direct calculation [Sti93c, Lemma 7.6] verifies that

$$\rho(\Sigma\mathcal{F}) = -\omega\rho(\mathcal{F}). \tag{4.48}$$

Further, since q and p count the number of occurences of the symbols E and F respectively, then applying equation 4.46 q times and equation 4.46 p times yields:

$$\chi(\rho(\mathcal{F})) = \mathcal{F}(0). \tag{4.49}$$

Using these identities we obtain:

$$\begin{aligned}
B^{-1}\chi(t) &= B^{-1}\chi(\omega q - p) \\
&= B^{-1}\mathcal{F}(0), && \text{by equation 4.49} \\
&= B^{-1}\mathcal{F}B(0) \\
&= (\Sigma\mathcal{F})(0), && \text{by equation 4.47} \\
&= \chi(\rho(\Sigma\mathcal{F})), && \text{by equation 4.49} \\
&= \chi(-\omega t), && \text{by equation 4.48,}
\end{aligned}$$

i.e. $\alpha C\chi(t) = \chi(-\omega t)$ (note that we are dealing with the maps E and F here, so that the rescaling is given by $\mathcal{B}^{-1} = \alpha C$ rather than $\alpha^2 C$). □

4.4.2 Invariant curve of pairs attracted

Proof of Theorem 47: This follows as a direct consequence of the results in [Sti93c, §4] (see section 4.2, in particular theorem 57). Therefore the above result (with the exception of the extra equation 4.2c) holds for asymptotically self-similar pairs as well. □

4.4.3 Invariant curve of functions attracted

Proof of Theorem 49: This now follows in the same way as Theorem 2.1 of [Sti93a], which dealt with the simple fixed point. Compare, in particular, the second part of [Sti93a, A.1, p.955–956]. (A brief description of the argument is sketched below.) □

The invariant curve obtained in the previous proof was for pairs attracted to the critical fixed point. [Sti93c] employs an "extension procedure" that yields an invariant curve for individual functions that are attracted.

As indicated above, there is no problem here for the critical fixed point: the parameterisation for the universal curve (which is unbounded) was explicitly constructed for a *piece* of the curve. It was then extended to the whole curve by making use of the exact self-similarity. (This entailed proving that the parameterisation can be extended to satisfy the "extra" equation 4.2c.)

A problem arises, however, when performing the extension procedure for maps attracted to the fixed point. The procedure, as presented in the appendix of [Sti93c], makes use of inverse maps. Fortunately, it may be rewritten to avoid this, using exactly the same method that is given in section 4.5.9 for a related problem. The problem is that during the extension procedure the injectivity of the map parameterizing the invariant curve is lost. In the case of twist maps [Sti93c], it was easily restored. Here, however, we obtain a parameterisation of the invariant curve that covers it m times (where m is some positive integer), rather than just once, as the parameter varies over $[-\omega, 1]$. For more details, see [Sti93a], which deals with the invariant curve for functions attracted to the simple fixed point.

Proof of Corollary 50: According to a classical theorem by Dini and Lipschitz [Zyg59, Theorem 10-3, p.63], the Fourier series of a Hölder continuous function converges uniformly. It follows from Abel's theorem [Rud82, Theorem 8.2, p.160] that the Schröder series converges on the boundary of its domain of definition and represents the critical invariant curve. This establishes the first part of corollary 50.

It remains to prove the last statement of the corollary, i.e. that the critical invariant curve passes through a stationary point of the map. Suppose this were not true, i.e., we assume that we are given a function f which is attracted by Widom's fixed point and that the critical invariant curve of f does not pass through a stationary point of f. Then, since the boundary curve is the closure of the orbit of the origin, f is not stationary along the orbit of the origin. Denoting by (E_j, F_j) the iterates under renormalization, it follows that $F_j'(0) \neq 0$ for all j. (Observe that the F_j are rescaled iterates of the original map f, and apply the chain rule.) Letting α_j denote the relevant rescaling factors we have, by definition, that

$$(F_{j+1})_* = \frac{1}{\alpha_{j+1}} E_j F_j \alpha_{j+1},$$

(where we have used the notation $f_*(z) = \bar{f}(\bar{z}) = CfC(z)$) and we obtain

$$(F_{j+1})_*'(0) = E_j'(\alpha_{j+1}) \cdot F_j'(0). \tag{4.50}$$

Now, at the fixed point (E, F), we have

$$E'(\alpha) = 2\alpha U'(\alpha^2).$$

Using the fact that $U_*(0) = 1 = U(0)$, we may evaluate the second component of the fixed point equation at 0 to give $V(1) = |\alpha|^2 = V_*(1)$. Then, differentiating the first component and evaluating at $\bar{\alpha}^2$ we have

$$U'_*(\bar{\alpha}^2) = \alpha V'(|\alpha|^4).$$

Combining these relations yields

$$U'(\alpha^2) = \bar{\alpha} V'_*(V_*^2(1)).$$

A similar (lengthy) argument [Sti93b, Equations 1–14] yields

$$\frac{1}{2|\alpha|^3} = |V'(V^2(1))|.$$

Thus

$$|E'(\alpha)| = |2\alpha U'(\alpha^2)| = \frac{1}{|\alpha|},$$

which is strictly greater than 1. Along with equation 4.50, this implies that the numbers $F_j'(0)$ grow geometrically, which obviously contradicts convergence to the fixed point. Thus f must have a stationary point on the invariant curve. □

Assuming that the relevant hypotheses are satisfied, this completes the proof of the main results. It now remains to show that the computer-assisted parts of the proof (namely, verifying the analytic extensions, the necklace hypotheses, and the existence of the fundamental domain) can actually be carried out. This is done in the next section.

4.5 Verifying the Hypotheses

The computer-assisted proof relies heavily on the manipulation of objects called "coarse curves".

4.5.1 Coarse curves

In order to verify the necklace hypotheses, it is necessary to represent the domains D_U and D_V in a rigorous way, and to apply rigorous representations of maps to them. In the existence proof, complex numbers were represented by the Cartesian product of a pair of intervals, namely the intervals bounding the real and imaginary parts. An extension of the techniques of interval analysis to these *rectangle* objects was implemented for the existence proof (recall section 3.4). This provided a means of applying maps rigorously to complex numbers.

Here, however, a further extension is required to provide some means of applying maps to whole domains (in this case discs) in the complex plane, and then checking, for example,

whether certain images of these domains intersect each other, contain the images of other domains, or are disjoint from each other. For verifying the necklace hypotheses, it is enough to represent the domains (in this case discs and their images under certain maps) by the curves defining their *boundaries*.

The Open Mapping Theorem The justification for this comes from the *open mapping theorem* (see, for example, [Rud80]) which says that a non-constant analytic function f is an *open mapping*, i.e. that f maps open sets to open sets.

In fact, we can argue more strongly that domains (i.e. *connected* open sets) are mapped to domains. (The continuous image of any connected set is connected, by a standard result of topology. The invariance of domains under non-constant analytic maps then follows from this together with the open mapping theorem.)

The key observation is then that an interior point of a domain must be mapped to an interior point of the image of the domain; it cannot be mapped to a boundary point of the image of the domain.

It then follows that the *boundary of the image* of a domain is a subset of the *image of the boundary*:

Lemma 78 (Image of Boundary Contains Boundary of Image)
Let f be a non-constant analytic map defined on a domain D. In addition, let D be bounded and let f be continuous on the closure of D. Then

$$\partial f(D) \subseteq f(\partial D).$$

The two additional assumptions are required because, a priori, it is not clear that f is well defined on the boundary of D. The first assumption implies that the closure of D is compact. The second implies that so is the image of the closure, being the continuous image of a compact set.

Proof: The proof of the main claim now goes as follows. Clearly,

$$f(D) \subseteq f(\overline{D}).$$

From this, since $f(\overline{D})$ is compact (hence closed) we conclude that even

$$\overline{f(D)} \subseteq f(\overline{D}).$$

Now, take an arbitrary point w of the *boundary* of the image, $w \in \partial f(D)$. In particular, w is in the closure of the image, and it follows from the above that w is in the image of the closure, i.e., that $w \in f(\overline{D})$. It follows that there exists $z \in \overline{D}$ such that $w = f(z)$. If z were an interior point of D, then, by the open mapping theorem, w would be an interior point of the image, in contradiction to our assumption. Therefore, there is a pre-image z that is a boundary point of D, and we have proved that

$$\partial f(D) \subseteq f(\partial D),$$

i.e. the boundary of the image is contained in the image of the boundary. □

We do not necessarily have strict equality in the above because it is possible for a point on the image of the boundary, say $w \in f(\partial D)$, to have more than one pre-image, say $z_1 \in \partial D$ (on the boundary of the original domain) and $z_2 \in D$ (in the *interior* of the original domain), i.e. a point on the boundary of the original domain can be mapped to the interior of the image. (It may be helpful to think of squaring a general disc that contains the origin: the image of the boundary has a self-intersection, and an extra "loop" that is not part of the boundary of the image.)

Remark 79
The above results obviously also hold for non-constant anti-*analytic maps.*

Now we will see why this is important for verifying the necklace hypotheses. The first thing to note is that the domains of the maps we are concerned with *are* bounded, and the maps themselves *are* continuous on the closures of their domains. The strategy is to rigorously apply a map f to a *covering* $\text{cover}(\partial D) \supseteq \partial D$ of the boundary ∂D of a domain. We then have

$$f(\text{cover}(\partial D)) \supseteq f(\partial D) \supseteq \partial f(D).$$

i.e. The resulting set is guaranteed to cover (at least) the boundary of the true (i.e. exact) image of the domain.

For example, to verify the domain extension conditions we need to show that the images of certain domains lie within certain discs. We can do this by covering the boundary of the domain, applying the map rigorously to the covering, and finally testing that the image of the covering lies within the disc. By the above reasoning, it would follow that the true image of the *whole* domain must be contained in the disc. (This is fairly obvious: if there were a point of the image set *outside* of the disc, there would also be a *boundary* point of it outside of the disc.)

Coarse-curves Since it is relatively straightforward to apply maps rigorously to rectangles having their edges parallel to the axes (code to achieve this had already been implemented for the existence proof), it was decided that the covering for each domain boundary would be made from a union of these rectangles. Such a union will be called a *coarse-curve* (shorthand for *coarse-grained curve*) and the individual rectangles will be called *grains*.

The maps themselves are represented in the framework used in the existence proof, namely they are stored as `Function` objects. Recall that each `function` object consists of a `Circle` representing the boundary of the domain, and a `Ball` representing the corresponding function ball on the unit disc. As in the existence proof, the truncation degree used was 160.

To test the necklace hypotheses, it then suffices to provide the following facilities:

1. Construction of a coarse-curve to cover any given circle.

2. Construction of a coarse-curve to cover any given line-segment (this allows polygons to be used as domain boundaries).

3. Verification that the region in the complex plane whose boundary is represented by a given coarse curve must be a subset of a given disc.

4. Verification that the region in the complex plane whose boundary is represented by a given coarse curve has a non-empty intersection with the interior of a given disc.

5. Verification that two regions in the complex plane whose boundaries are represented by two given coarse curves are disjoint from each other.

Covering a curve with a coarse-curve

Methods are provided (in C++-jargon they are called *constructors*) for making a coarse-curve with a specified number of grains to cover a line-segment, a polygon, or a circle. The following discussion demonstrates how this is achieved for a general curve, the specific case of the circle is then dealt with.

The general covering process is based on the following observation:

Observation 80
Let Γ be a continuous curve in the complex plane. If the segment of the curve joining two points z_1 and z_2 is monotonic (or constant) in both its real and imaginary parts, then the portion of Γ joining z_1 and z_2 is contained within a rectangular region having its sides parallel to the coordinate axes, and having the two end-points z_1 and z_2 as diagonally opposite corners.

(For example, notice that the upper right-hand quadrant of the unit circle is contained within the unit square (having $z_1 = (1, 0)$, $z_2 = (0, 1)$.) Using the above observation, the strategy for covering the objects is as follows:

1. Subdivide the curve to be covered into segments, each monotonic in both the real and imaginary parts (a line-segment already has this property, so the whole segment may be used. A polygon is simply divided into its individual line-segments, each of which will certainly be monotonic (or constant). And a circle is dealt with one quadrant at a time, each quadrant being a monotonic curve.) Figure 4.5(a) shows a curve AB which may be divided into monotonic segments AC and CB. Figure 4.5(b) demonstrates that both segments are then contained within rectangles parallel to the axes having their end-points as corners.

2. For each of the above segments, some small 'sample-rectangles' are chosen that lie on the segment (this is done rigorously, so that these rectangles are each guaranteed to cover a point on the segment). In addition, we ensure that each end-point is covered by a sample-rectangle.

3. Observe that the segments of the curve between adjacent sample-rectangles will also be monotonic, because the whole segment that they are a part of is monotonic. Cover

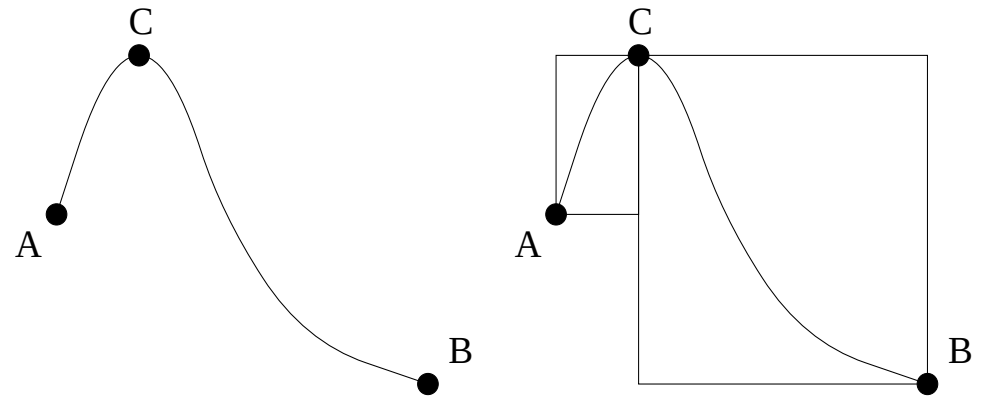

(a) A curve AB. (b) Monotonic segments AC and CB.

Figure 4.5: Division of a curve into monotonic segments.

each adjacent pair of sample-rectangles by the minimal rectangle that covers them both. Each of these final covering rectangles will be a grain making up the coarse-curve. Notice that (1) each grain completely covers the segment of curve between the two sample-points that it was placed over (this is clear from the monotonicity observation above) and (2) the union of the grains is connected because adjacent grains are guaranteed to overlap at the sample-rectangles, thus there can be no 'holes' in the covering.

Figure 4.6 indicates how the above procedure might be carried out for the segment CB. Some example sample-rectangles are shown in Figure 4.6(a). The resulting grains that cover adjacent pairs of them are shown in Figure 4.6(b). (It is worth noting that the

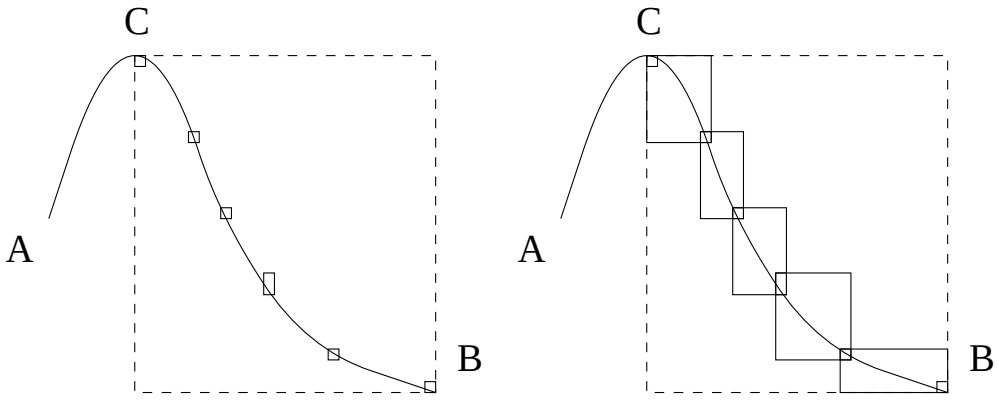

(a) Sample rectangles. (b) Grains covering adjacent pairs.

Figure 4.6: Covering each monotonic segment with rectangular grains.

sample rectangles will typically be degenerate in either the real or imaginary directions, meaning that adjacent grains will touch along a small horizontal or vertical line-segment, rather than overlapping on a rectangle.)

The entire segment of curve is now covered completely by the union of the grains. The

process is repeated for the other monotone sections of the curve, until the whole curve has been covered by the rectangular grains.

Covering a circle

Firstly, notice that in the case of a circle it is sufficient to cover a circle of the correct radius centred at the *origin*, since the resulting grains may then be mapped to cover the desired circle by translation, an operation that can be performed rigorously (using rectangle addition). Furthermore, it is actually enough to cover 1/8th of the circle, for example the piece between the x-axis and the line $y = x$ in the upper-right quadrant. The rest of the circle may then be covered using symmetry, by performing rigorous operations to take the real and imaginary parts of rectangles and exchange and/or negate them as required.

Firstly, the end-points of the region are covered by small sample-rectangles. The remaining sample-points are calculated by using a constant step-size in the angle θ between the x-axis and a point on the circle. As a result, the x coordinates of the sample points will be clustered more closely together near to the point where the circle intersects the x-axis. This is desirable, so that the rectangular grains covering the circle are all of approximately the same size. The x-values are typically degenerate intervals (i.e. single points).

Pythagoras' theorem is then applied rigorously to find the y-coordinate intervals that are guaranteed to contain the y-coordinates corresponding to the sample x-coordinates. Note that although the x-coordinates were degenerate intervals, the rigorous nature of the programs mean that the corresponding y-intervals are typically *not* degenerate (although they will tend to be very small) due to allowances for rounding error. In this way sample-rectangles are found that cover points lying on the 1/8th of the circle, including the end-points of the region.

Adjacent pairs of sample-rectangles are then covered by rectangular grains. Because the segments of the circle are monotonic in the half-quadrant that is to be covered, each of these grains is also guaranteed to contain the whole of the curve that lies between the adjacent sample-points.

Finally, symmetry is used to create negativised, conjugated and reflected copies of the grains produced, so that the whole circle at the origin is covered.

The rectangle representing the desired centre of the circle is then added to each of the grains, which has the effect of translating the coarse-curve into the correct position.

Figure 4.7 shows part of a coarse-curve that has been constructed on the unit circle. The Figure 4.7(a) shows a coarse-curve made up from 32 rectangular grains, whereas Figure 4.7(b) shows a coarse-curve made up from 256 grains. In both cases the fact that adjacent grains overlap is barely visible, but is guaranteed by construction.

The next sections detail the implementation of the rigorous tests needed to verify the necklace hypotheses.

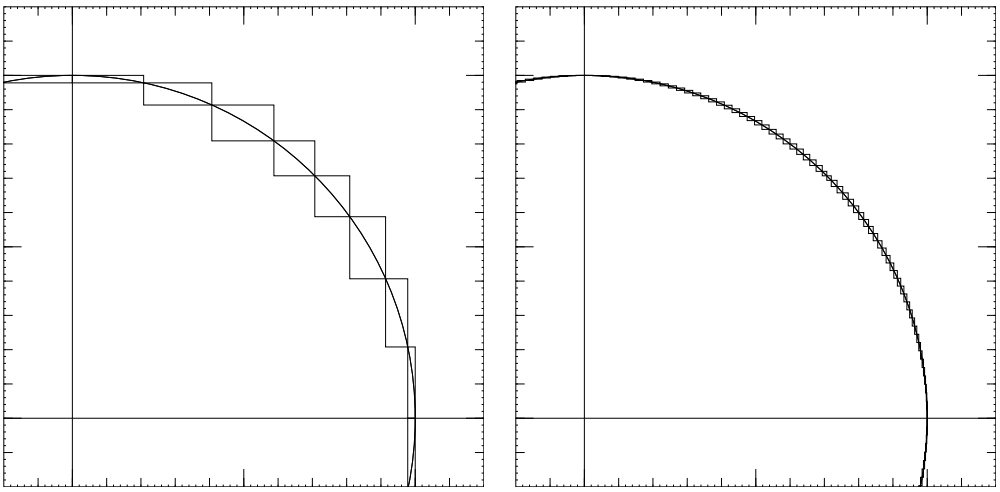

(a) Unit circle covered by 32 grains. (b) Unit circle covered by 256 grains.

Figure 4.7: Covering the Unit-circle

Verifying containment within a disc

Program Fragment 4.2 Checking containment within discs.

```
bool Circle::Contains(const Rectangle& z) const
{
  return (sup(abs(z - c)) < inf(abs(r)));
}

bool Circle::Contains(const CoarseCurve& C) const
{
  GrainIterator nextGrain(C);        //access the grains of C.
  const Rectangle* R;                //each is a rectangle.
  while( R = nextGrain() )           //check each grain.
    if( !Contains( *R ) )
      return false;                  //fail if any not contained.
  return true;                       //otherwise succeed.
}
```

In order to verify that the region whose boundary curve is covered by a given coarse-curve is contained inside a disc, it suffices to show that *all* of the grains making up the coarse-curve lie within the disc. This can be done using rectangle arithmetic by calculating the distance of each rectangle from the centre of the circle and ensuring that it is strictly less than the radius. Program fragment 4.2 demonstrates two functions: the first checks that a rectangle lies within a circle, and the second function uses the first to do the same for a coarse-curve, by checking each rectangular grain in turn to see if it lies in the circle. (The code uses an "iterator" object [GHJV94] to step through the grains one at a time, this is an elegant alternative to using structures like "for-loops"; it shields the programmer from having to know details of the actual storage method used for the grains.)

Verifying non-empty intersection with a disc

In order to verify that the region whose boundary curve is covered by a given coarse-curve has non-empty intersection with a disc, it is sufficient to test that *at least one* of the rectangles making up the coarse curve lies entirely inside the disc.

To see why this is so, recall that we begin by covering a domain boundary, each rectangle in the cover contains at least one point of the boundary. Then each of these boundary points is mapped either to a boundary point of the image, or to the interior of the image. In either case, if the covering rectangle is strictly contained within the disc, it follows that the (open) disc contains at least one boundary point or interior point, i.e. it has a non-empty intersection with the image of the original domain. This test is performed is a similar way to the one shown above.

Verifying disjointness of coarse-curves

For the purposes of verifying the necklace hypotheses, it turns out that in order to prove that the regions whose boundary curves are covered by two coarse-curves are disjoint, it is enough to identify some directed straight line that divides the plane into two regions, one entirely containing the first coarse curve, and the other entirely containing the second. Since the coarse-curves cover (at least) the whole boundary of the domain in each case, it follows that the domains must be disjoint. The vector dot-product (or, rather, a complex-number version of it) is used to verify that all the rectangles of one coarse-curve lie on the *right* of the line whereas all the rectangles of the other coarse-curve lie on the *left* of the line.

Remark 81 (Proving Disjointness)
Note that it would not *be sufficient, for establishing disjointness, to simply prove that none of the individual rectangles of the first coarse-curve intersect the second, because this does not rule out the possibility of one domain being a subset of the other (i.e. one boundary would "surround" the other).*

Remark 82 (Rigorous Tests)
It is important to note what is meant by a 'rigorous test'. If such a test succeeds, then the mathematical statement that it embodies is true. On the other hand, if the test fails, nothing can be said about the statement. For example, suppose that a test to verify that a coarse-curve is contained in a disc succeeds. It immediately follows that the coarse-curve (and any curve that it covers) is indeed contained in the disc. However, if the test fails, then it might still be that the true (mathematical) curve is contained in the disc, but that the error bounds were such that this could not be concluded with certainty — It may equally well be the case that the true curve was not *contained.*

Remark 83 (A Note About Pictures)
In the following subsections, pictures will be given demonstrating the verification of the various hypotheses by computer. It is important to note that the programs that actually perform the proofs do in fact check each condition using rigorous calculations. As such,

the pictures are designed for illustration purposes only and are not considered as a proof in themselves.

In some of the pictures, a grid will be superimposed over the images of some of the domains in order to distinguish them from the other domains involved. In addition, some coarse-curves will appear to have a thick boundary in places: this indicates that the covering rectangles have grown to accommodate the error bounds present.

4.5.2 Domains

In order to prove that it is possible to take the analytic continuation to the new domains, the main proof program must again use the domains D_U and D_V that were used in the existence proof (see Chapter 2). Both are discs, $D_U = \{z : |z - c_U| \leq r_U\}$ where

$$c_U = -0.2188497414079558 + i0.2328147240271490 \quad (4.51a)$$
$$r_U = 0.3640985354093064, \quad (4.51b)$$

and $D_V = \{z : |z - c_V| \leq r_V\}$ where

$$c_V = 0.5672961438978619 + i0.1229664702397770 \quad (4.52a)$$
$$r_V = 0.636. \quad (4.52b)$$

The analytic extensions of the maps U and V are taken to the new domains Δ_U and Δ_V respectively. Δ_U is the disc whose centre and radius are given by

$$c_{\Delta_U} = -0.365 + i0.635,$$
$$r_{\Delta_U} = 0.47.$$

Δ_V is a connected union of three discs, having centres and radii given by:

$$c_{\Delta_{V_1}} = 0.363691 + i0.174008, \quad r_{\Delta_{V_1}} = 0.271654$$
$$c_{\Delta_{V_2}} = -0.110599 + i0.192411, \quad r_{\Delta_{V_2}} = 0.157012$$
$$c_{\Delta_{V_3}} = 0.0575719 + i0.0272709, \quad r_{\Delta_{V_3}} = 0.203782$$

These domains are illustrated in Figure 4.8.

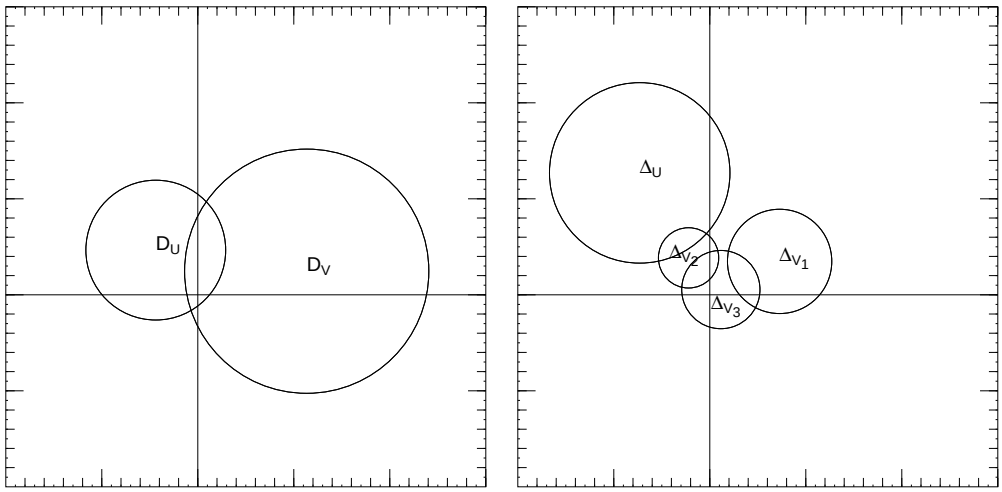

(a) Old domains. (b) New domains.

Figure 4.8: Domains.

In addition, the proof uses the approximate fixed point that was output by the existence proof, including the corresponding rigorous error bounds.

4.5.3 Verifying the Analytic Extensions

The first step is to verify that the fixed point has an analytic continuation to the new domains.

Let U and V be the components of the fixed point of the existence proof and let D_U and D_V be their domains (the discs defined above). Let Δ_U and Δ_V be the new domains as defined above. Note that $\Delta_{V_1} \subset D_V$. Note, moreover, that the new domains intersect the old ones: $\Delta_U \cap D_U \neq \emptyset$ and $\Delta_V \cap D_V \neq \emptyset$, as shown in Figure 4.9.

A proof is needed that U and V can be extended analytically to the unions of their old and new domains. Since the unions are connected, these extensions are uniquely determined if they exist (see, for example, [Rud80]). The unions of the old and new domains will be denoted $\Omega_U = D_U \cup \Delta_U$ and $\Omega_V = D_V \cup \Delta_V$.

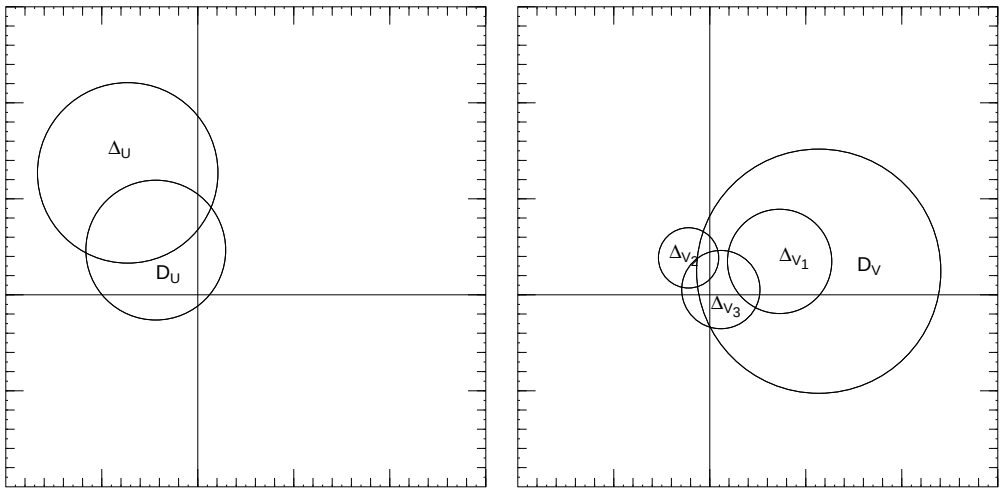

(a) Ω_U. (b) Ω_V.

Figure 4.9: Composite Domains Ω_U and Ω_V.

Extension of U:

In order to prove that U can be analytically extended to Ω_U, note that, by the fixed point equation (4.30a), the following identity holds:

$$U = C\alpha^{-1}V\alpha^2 C =: \hat{U} \quad \text{on } D_U.$$

We see that the right hand side of this equation ($\hat{U}$) is well defined on Δ_U if $\alpha^2 C\Delta_U \subset \Delta_{V_1}$ because, as mentioned, $\Delta_{V_1} \subset D_V$. This is verified by the program. (See Figure 4.10.)

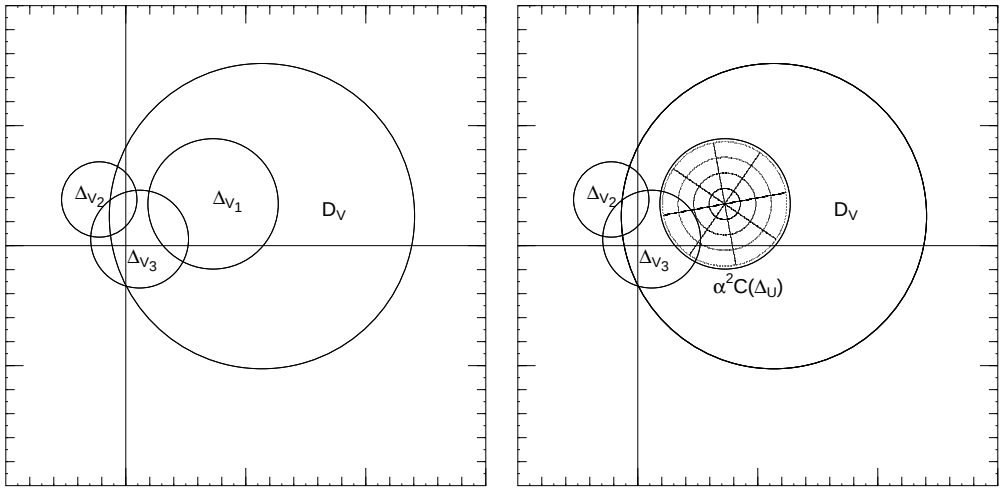

(a) $\Delta_{V_1} \subset D_V$. (b) $\alpha^2 C\Delta_U \subset \Delta_{V_1}$.

Figure 4.10: Verifying that U extends to $D_U \cup \Delta_U$.

It follows that $C\frac{1}{\alpha}V\alpha^2 C$ is the (unique) analytic extension of U to $D_U \cup \Delta_U = \Omega_U$. From now on, U may be regarded as an analytic function on Ω_U.

Extension of V to $D_V \cup \Delta_{V_1}$:

Notice that $\Delta_{V_1} \subset D_V$. Therefore, analytic extension to $D_V \cup \Delta_{V_1}$ is trivial.

Extension of V to $D_V \cup \Delta_{V_2}$:

In order to prove that V can be extended to $D_V \cup \Delta_{V_2}$, recall that, by the fixed point equation (4.33) and by commutativity of the fixed point, the identity

$$V = \frac{1}{|\alpha|^2} V\alpha^2 CQV\alpha^2 C =: \hat{W}, \tag{4.53}$$

holds on a non-empty open neighbourhood of the origin. In addition, consider the disc

$$\nu\Delta_{V_2} = \{z : |z - c_{\Delta_{V_2}}| \leq 1.5 r_{\Delta_{V_2}}\},$$

which is concentric with Δ_{V_2} and which contains both Δ_{V_2} and the origin, as shown in Figure 4.11(a). The program verifies that the right hand side of the above identity (4.53), denoted by $\hat{W}$, is well defined on $\nu\Delta_{V_2}$ as well, i.e. that $\alpha^2 CQV\alpha^2 C(\nu\Delta_{V_2}) \subset D_V$. See Figure 4.11(b).

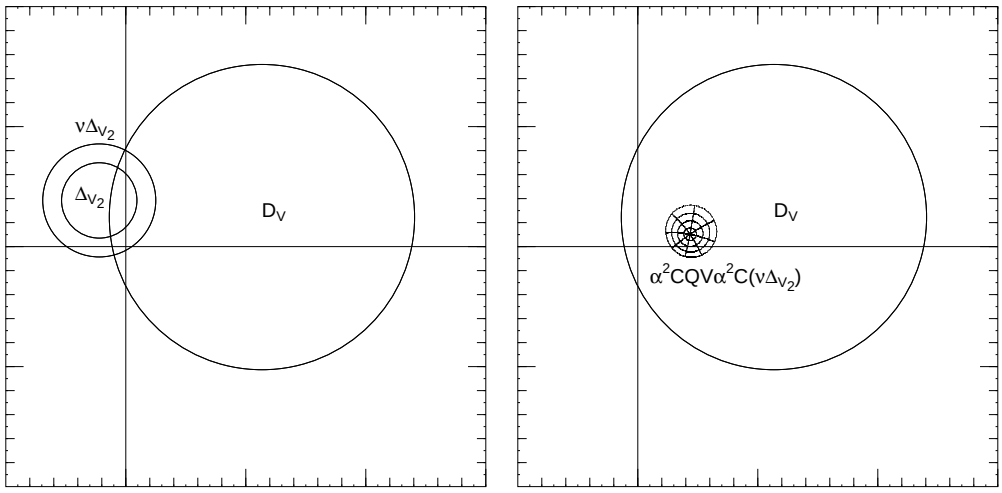

(a) $\nu\Delta_{V_2}$. (b) $\alpha^2 CQV\alpha^2 C(\nu\Delta_{V_2}) \subset D_V$.

Figure 4.11: Verifying that V extends to $D_V \cup \Delta_{V_2}$.

It follows that $\hat{W}$ is an analytic extension V to $D_V \cup \Delta_{V_2}$. Since the intersection of the open discs D_V and Δ_{V_2} is not empty, this extension is uniquely determined.

Extension of V to $D_V \cup \Delta_{V_3}$:

In order to prove that V can be extended to $D_V \cup \Delta_{V_3}$, we observe once again that the fixed point equation gives us an identity:

$$V = C\frac{1}{\alpha} VQC\frac{1}{\alpha} V|\alpha|^4 =: \hat{V} \quad \text{on } D_V. \tag{4.54}$$

The program verifies that the right hand side of the above equation (4.54), denoted by $\hat{V}$, is well defined on Δ_{V_3} as well. To do this, it checks that both $|\alpha|^4(\Delta_{V_3})$ and $QC\frac{1}{\alpha}V|\alpha|^4(\Delta_{V_3})$ are contained in D_V. (See Figure 4.12.)

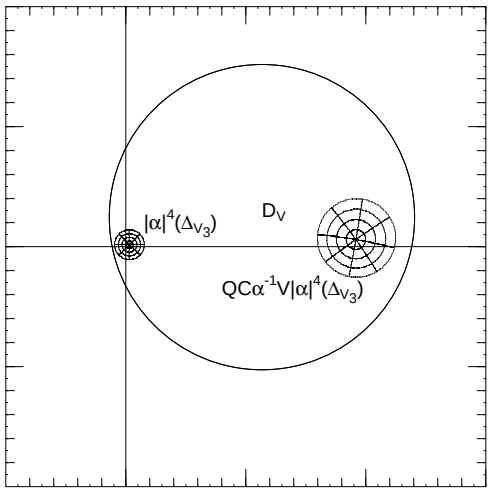

Figure 4.12: Verifying that V extends to $D_V \cup \Delta_{V_3}$.

It follows that $\hat{V}$ is an analytic extension of V to $D_V \cup \Delta_{V_3}$. Since the intersection of the open discs D_V and Δ_{V_3} is not empty, this extension is uniquely determined.

Extension of V to Ω_V:

Finally, we notice that the intersection of the open discs D_V, Δ_{V_2}, and Δ_{V_3} is not empty. It follows that V can be extended to

$$D_V \cup \Delta_{V_1} \cup \Delta_{V_2} \cup \Delta_{V_3} = \Omega_V.$$

From now on, we may regard V as an analytic function on the domain Ω_V.

4.5.4 The Fixed Point Equation

We are going to prove that the extended functions U and V satisfy the accretive fixed point equation (4.31) on the new domains. The first component of this fixed point equation,

$$U = C\frac{1}{\alpha}V\alpha^2C,$$

holds by the definition of the extended function U on Ω_U and, a fortiori, on Δ_U (since $\Delta_U \subset \Omega_U$).

To prove the second one, we note that, by commutativity, the identity

$$V = C\frac{1}{\alpha}UQV\alpha^2C$$

holds close to the origin. This equation will continue to hold on the connected component of the domain of definition of the right hand side. The domain extension conditions to be

verified below imply that the right hand side is well defined on Δ_V. Therefore, and since Δ_V contains the origin, the second component of the fixed point equation holds on Δ_V.

4.5.5 Domain Extension

The program verifies the three domain extension conditions (equations 4.37a–4.37c). Recall that the domains Δ_U and Δ_V consist of a disc and a connected union of three discs, respectively.

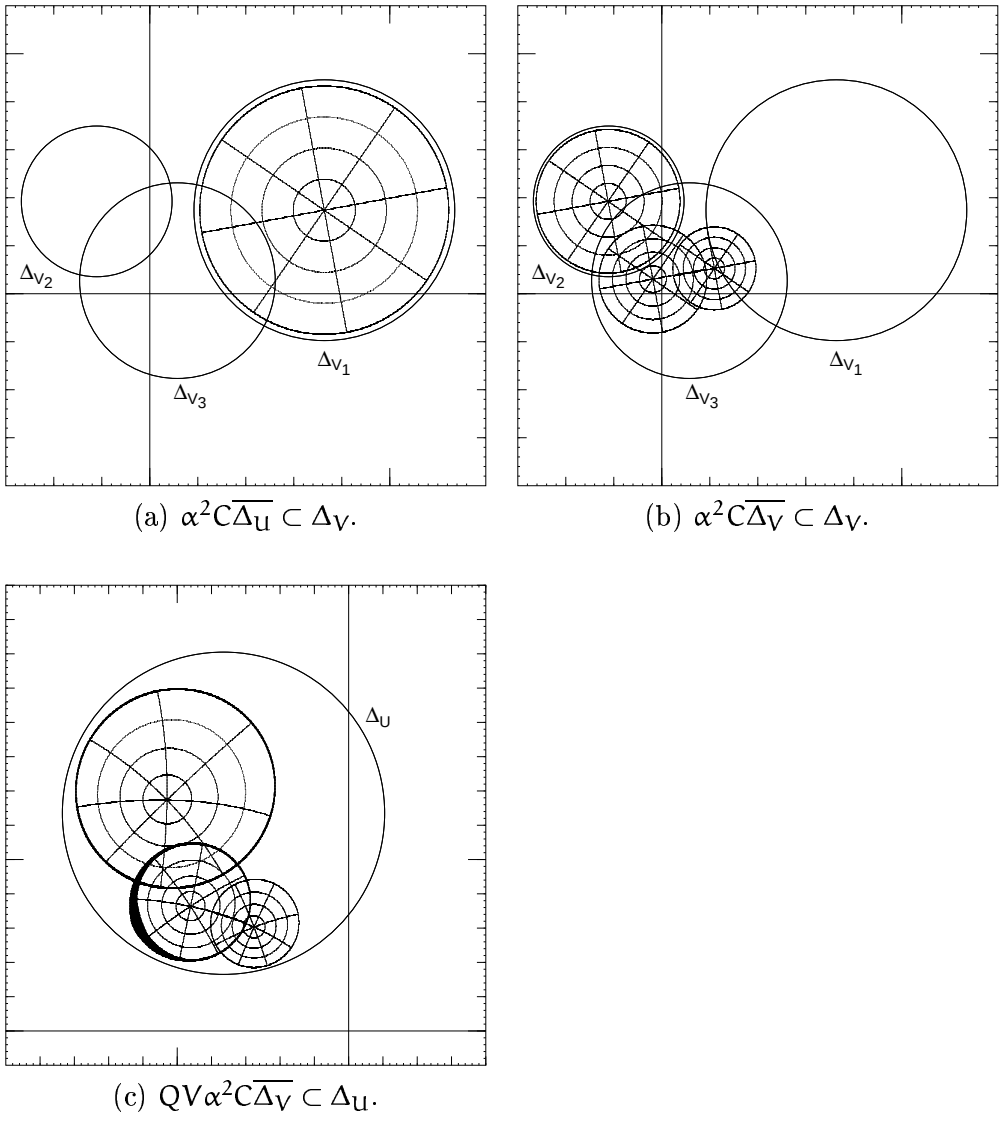

(a) $\alpha^2 C\overline{\Delta_U} \subset \Delta_V$. (b) $\alpha^2 C\overline{\Delta_V} \subset \Delta_V$.

(c) $QV\alpha^2 C\overline{\Delta_V} \subset \Delta_U$.

Figure 4.13: Verifying the Domain Extension Conditions

To verify each of the conditions, it is required to check that certain images of these domains are contained within the discs that represent the domains themselves. The coarse curve construction is used to apply the relevant maps rigorously to the original domains. It only remains for the program to use the facilities provided for checking containment of coarse curves within discs to complete this part of the proof. The computer verification is illustrated in Figure 4.13. Some of the program code is given in Program Fragment 4.3. (In particular, notice the use of the analytic continuations $\widehat{W}$ and $\widehat{V}$ in the program. Note also that whereas only 4096 grains are needed in the coarse-curve to verify the first condition,

32768 have been used for the later conditions in order to get good enough bounds. The badly-behaved numerics reveal themselves in Figure 4.13(c), where the coarse-curve covering the boundary of the image of Δ_{V_3} contains some large grains, giving the appearance of a thicker boundary.)

Program Fragment 4.3 Computer Verification of Domain Extension.

```
const int numberOfGrains = 32768;
CoarseCurve cc4( Q(W_hat(a2(C))), deltaV1, numberOfGrains );
CoarseCurve cc5( Q(V(a2(C))),     deltaV2, numberOfGrains );
CoarseCurve cc6( Q(V_hat(a2(C))), deltaV3, numberOfGrains );
if( deltaU.Contains(cc4)
     && deltaU.Contains(cc5)
      && deltaU.Contains(cc6) )
    cerr<<"Domain Extension (3) verified"<<endl;
else
    cerr<<"Domain Extension (3) failed"<<endl;
```

4.5.6 Connectedness

The connectedness conditions (equations 4.38a–4.38c) are also verified by performing rigorous comparisons of coarse curves (representing the images of the domains) with the discs of the original domains.

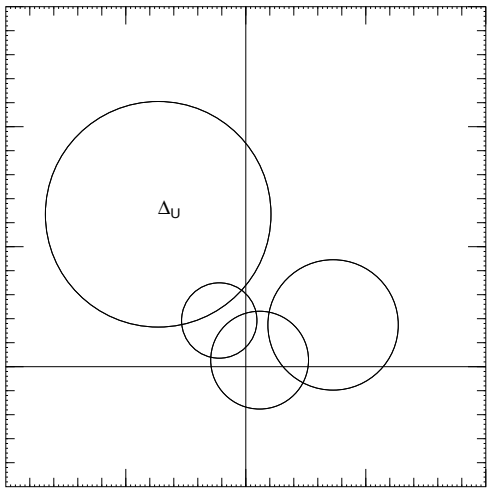

(a) $\Delta_U \cap \Delta_V \neq \emptyset$.

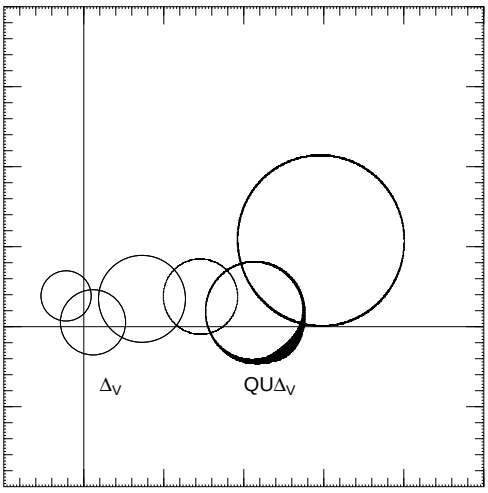

(b) $Q U \Delta_V \cap \Delta_V \neq \emptyset$.

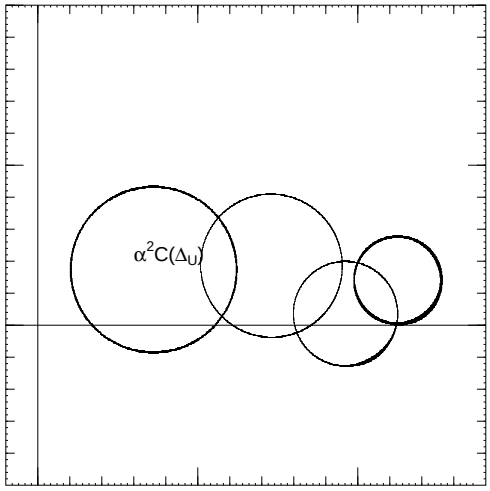

(c) $C\frac{1}{\alpha^2} Q U \alpha^2 C \Delta_V \cap \Delta_U \neq \emptyset$.

Figure 4.14: Verifying the Connectedness Conditions.

In this case, the test for intersection of coarse curves with discs is used. This requires that we apply the map $B = C\frac{1}{\alpha^2}$ to the equation for the last connectedness condition (equation 4.38c) to give

$$C\frac{1}{\alpha^2} Q U \alpha^2 C \Delta_V \cap \Delta_U \neq \emptyset,$$

so that the second region making up the intersection on the left hand side is just a disc (the comparison of coarse curves with discs is computationally much simpler than that

of coarse curves with other coarse curves.) The computer verification is illustrated in Figure 4.14, along with a portion of the computer code in Program Fragment 4.4.

Program Fragment 4.4 Computer Verification of Connectedness Conditions

```
if(  deltaU.Intersects( CoarseCurve(id,deltaV2,4096) )  )
    cerr<<"Connectedness (1) verified"<<endl;
else
    cerr<<"Connectedness (1) failed"<<endl;
```

4.5.7 Disjointness

To verify the disjointness conditions, we again use the extension $\hat{\mathsf{U}} = \mathsf{C}\frac{1}{\alpha}\mathsf{V}\alpha^2\mathsf{C}$ of U. To establish the disjointness of the relevant images of the domains, it turns out to be sufficient to identify a line that separates them. The program then uses the rigorous test described earlier to verify that the regions lie on opposite sides of the line. (In fact, in this case it suffices to use vertical lines, so that we may simply test the real parts of the rectangular grains making up each coarse curve.)

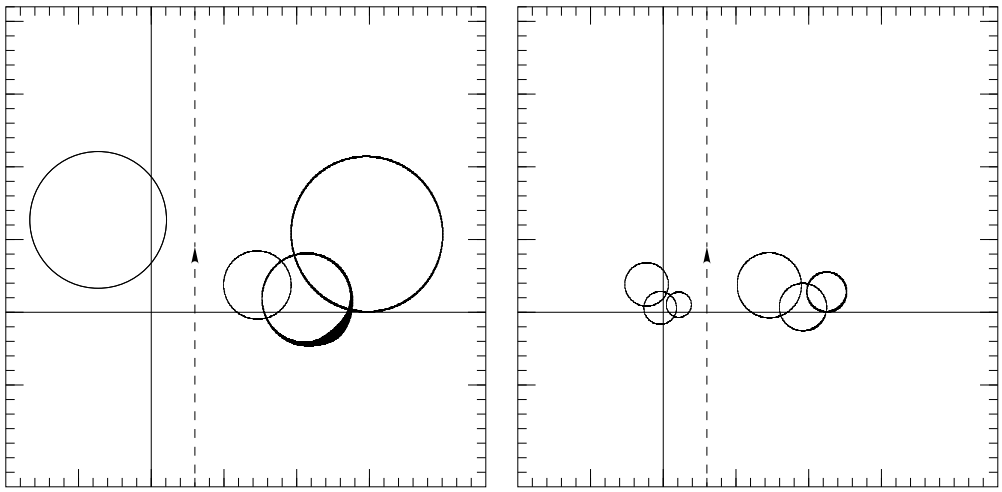

(a) $\overline{\Delta_U} \cap Q \cup \overline{\Delta_V} = \emptyset$. (b) $\alpha^2 C\overline{\Delta_V} \cap Q \cup \alpha^2 C\overline{\Delta_V} = \emptyset$.

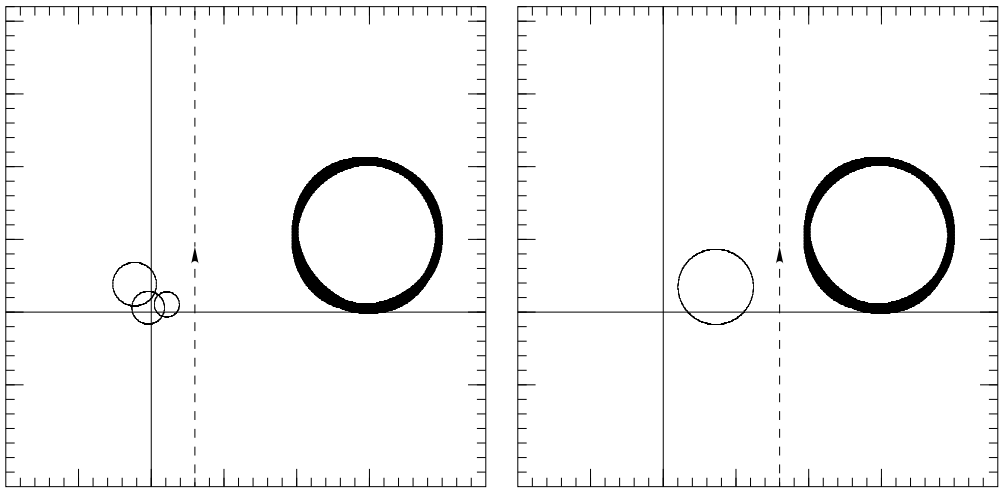

(c) $\alpha^2 C\overline{\Delta_V} \cap Q \cup \alpha^2 C\overline{\Delta_U} = \emptyset$. (d) $\alpha^2 C\overline{\Delta_U} \cap Q \cup \alpha^2 C\overline{\Delta_U} = \emptyset$.

Figure 4.15: Verifying the Disjointness Conditions.

For the first three conditions, the line $\Re z = 0.3$ was used. For the final condition, the line $\Re z = 0.8$ was used. The computer verification is illustrated in Figure 4.15 and Program Fragment 4.5. As the program code indicates, when using the analytic extensions the map $Q \cup B^{-1}$ (for example) becomes

$$Q \cup \alpha^2 C = Q\hat{U}\alpha^2 C = QC\alpha^{-1}V\alpha^2 C\alpha^2 C = QC\alpha^{-1}\hat{W}|\alpha|^4,$$

on the new domains.

Program Fragment 4.5 Computer Verification of the Disjointness Conditions.

```
...
CoarseCurve QUBinvDeltaU( Q(C(b(W_hat(a2(C(a2(C)))))), deltaU, 4096);
CoarseCurve BinvDeltaU(   a2(C),                        deltaU, 4096);

const LineSegment line2(rectangle(interval(0.8), interval(-1.0)),
                        rectangle(interval(0.8), interval(1.0)));

if( OnRight(QUBinvDeltaU, line2) && OnLeft(BinvDeltaU, line2) )
    cerr<<"Disjointness Verified"<<endl;
else
    cerr<<"Disjointness Failed"<<endl;
...
```

Alternatively, one could identify a pair of disjoint circles, and verify that each contained one set of the relevant domains.

4.5.8 Contractivity

According to the translation rules, the contractivity condition requires the map QVB^{-2} to be a uniform contraction on $\Delta_U \cup \Delta_V$. The program simply verifies that the range of the derivative of this map is contained in the unit disc. This derivative is given by

$$(QVB^{-2}(z))' = (V^2(|\alpha|^4 z))' = 2V(|\alpha|^4 z) \cdot |\alpha|^4 V'(|\alpha|^4 z).$$

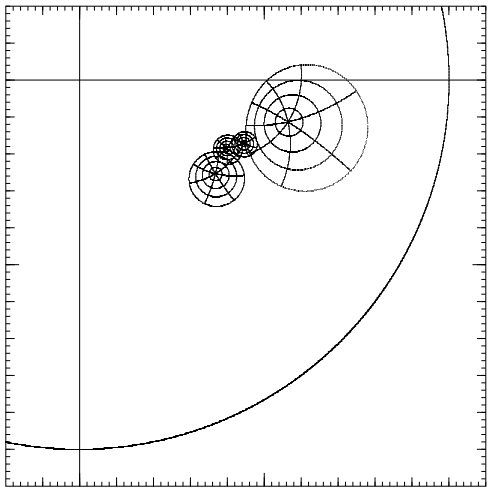

Figure 4.16: Verifying the Contractivity Condition.

In fact, the proof establishes more than this; it turns out that the range of the derivative lies within an annulus centred on the origin, bounded above and below by circles of radius 0.795694 and 0.380371 respectively (these numbers will be used later, in chapter 5, to obtain some crude bounds on the dimension of the Siegel disc boundary). The computer

verification is illustrated in Figure 4.16, where the values of the derivative are shown in comparison to the unit disc.

4.5.9 Univaluedness

As mentioned in the translation of the necklace construction, a problem arises: in order to carry out the proof of [Sti93c, Lemma 7.15] (which proves that patches overlap if and only if they are neighbours in the succession) we need to invert our maps. It turns out to be enough to require that the map V is locally homeomorphic on some domain, i.e. V must be univalued (schlicht) on a certain domain, which will be called a *fundamental domain*. If V maps the boundary of a domain Ω into a Jordan curve, it follows that it is univalued on the interior of Ω ([Mar67, Theorem 4.5, p. 118].)

The Fundamental Domain:

In order to understand the following, look at the proof of [Sti93c, Lemma 7.15], which establishes that patches overlap if and only if they are neighbours in the succession. The idea is to take three successive patches in the succession and look at the combinations of legal parents that are possible for them. The proof then entails showing that disjointness of the first and third child always reduces to one of the disjointness conditions. It turns out that there are three relevant cases.

In the first case, for instance, the inverse of the map F is taken, where F is a Fibonacci tail in $\{F_{j-1}\}$. This works fine in the twist maps case, but in our case the map F is not a global homeomorphism. The two patches concerned in that part of the proof are of the form $FF_jB^{-(j+1)}\overline{\Delta_V}$ and $FB^{-(j+1)}\overline{\Delta_U}$. The condition for disjointness, namely

$$FF_jB^{-(j+1)}\overline{\Delta_V} \cap FB^{-(j+1)}\overline{\Delta_U} = \emptyset,$$

may be rewritten as follows:

$$FB^{-(j-1)}B^{-1}B^jF_jB^{-j}B^{-1}\overline{\Delta_V} \cap FB^{-(j-1)}B^{-2}\overline{\Delta_U} = \emptyset.$$

Suppose, for the moment, that we can invert $FB^{-(j-1)}$. Noticing that, by [Sti93c, Lemma 2.2], $B^jF_jB^{-j} = QV$, we obtain

$$B^{-1}QVB^{-1}\overline{\Delta_V} \cap B^{-2}\overline{\Delta_U} = \emptyset,$$

and after applying the homeomorphism B^2 to that equation, we obtain, using the fixed point equation, one of the disjointness conditions:

$$QU\overline{\Delta_V} \cap \overline{\Delta_U} = \emptyset.$$

(c.f. equation 4.40a.)

Now, we observe that for any Fibonacci tail F in $\{F_j\}$, the map FB^{-j} is a homeomorphism on any region where V is a homeomorphism. This follows simply from the equation

$$B_1^j\hat{F}_jB^{-j} = V,$$

where F_j is the j-th Fibonacci string with the symbol Q dropped. (This equation, in turn, is obtained by iterating the strong form (equation 4.36), i.e. the 'uncle', of the fixed point equation.)

Therefore, the argument works if we can find a region Ω on which V is a homeomorphism, and which contains the two regions $B^{-1}QVB^{-1}\overline{\Delta_V}$ and $B^{-2}\overline{\Delta_U}$.

In the same way, the whole proof of [Sti93c, Lemma 7.15] can be made to work, provided that Ω contains all the relevant patches that are used.

The following patches must be considered (in what follows we refer to the cases in the proof of [Sti93c, Lemma 7.15]):

- **First Case:** This case was treated above in detail. It turned out that the fundamental domain Ω must contain the patches $B^{-1}QVB^{-1}\overline{\Delta_V}$ and $B^{-2}\overline{\Delta_U}$.
- **Second Case:** In the second case, the crucial line in the proof of [Sti93c], Lemma 7.15, is

  $$FF_{j-2}B^{-(j+1)}\overline{\Delta_V} \cap FF_jB^{-(j+1)}\overline{\Delta_V} = \emptyset$$

  where $F \in \{F_{j-1}\}$ and $FF_{j-2} \in \{F_j\}$. We rewrite this as

  $$FF_{j-2}B^{-j}B^{-1}\overline{\Delta_V} \cap FF_{j-2}B^{-j}B^jF_{j-1}B^{-j}B^{-1}\overline{\Delta_V} = \emptyset.$$

  (We have used $F_j = F_{j-2}F_{j-1}$.) Since $FF_{j-2} \in \{F_j\}$, we can take the inverse of $FF_{j-2}B^{-j}$, provided that the patches $B^{-1}\overline{\Delta_V}$ and $QUB^{-1}\overline{\Delta_V}$ are contained in Ω. (Observe that, by [Sti93c, Lemma 2.2], $B^jF_{j-1}B^{-j} = QU$.)
- **Third Case:** In the first sub-case of the third case, the crucial line is

  $$FB^{-(j+1)}\overline{\Delta_V} \cap FF_{j-1}B^{-(j+1)}\overline{\Delta_V} = \emptyset,$$

  with $F \in \{F_{j-2}\}$. We rewrite this as

  $$FB^{-(j-2)}B^{-3}\overline{\Delta_V} \cap FB^{-(j-2)}B^{-1}B^{j-1}F_{j-1}B^{-(j-1)}B^{-2}\overline{\Delta_V} = \emptyset.$$

  We can take the inverse of $FB^{-(j-2)}$, provided that the patches $B^{-3}\overline{\Delta_V}$ and $B^{-1}QVB^{-2}\overline{\Delta_V}$ are contained in Ω. (Observe that $B^{(j-1)}F_{j-1}B^{-(j-1)} = QV$.) The remaining two subcases follow precisely the above pattern; only the domains change.

A region Ω which does the job is the ("coffin-shaped") polygon given by the following vertices:

$$\begin{aligned} P_1 &= -0.1 - i0.15, & P_2 &= 1.1 - i0.15, \\ P_3 &= 1.3 + i0.00, & P_4 &= 1.3 + i0.25, \\ P_5 &= 1.1 + i0.44, & P_6 &= 0.0 + i0.44, \\ P_7 &= -0.3 + i0.30, & P_8 &= -0.3 + i0.10. \end{aligned}$$

The program verifies that the region whose boundary is represented by this polygon indeed contains the necessary patches. To see how this is achieved, we note that the polygon is *convex*. It suffices to show that by traversing the edges of the polygon in anticlockwise fashion, the necessary patches always lie to the left of the lines that are the extensions of each edge, as illustrated in figure 4.17.

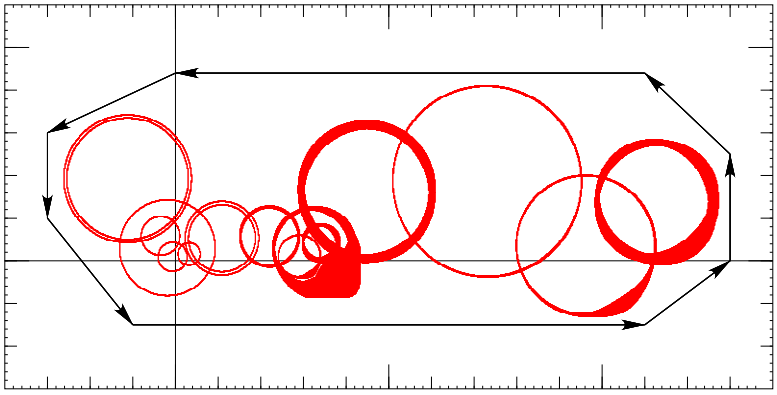

Figure 4.17: Containment of the patches within the Fundamental Domain.

The Variation of the Argument:

To prove that the image of this polygon has no self-intersections (and, therefore, that is really *is* a fundamental domain), it suffices to show that the argument with respect to some fixed interior point varies monotonically (see, for example, [Rud80]).

The total change of the argument with respect to a point w of a curve $t \mapsto f(z(t))$ is the imaginary part of

$$\frac{1}{2\pi}\oint_{\Omega}\frac{f'(z(t))}{f(z(t))-w}z'(t)dt,$$

(where here Ω indicates the boundary of the domain.) For our purposes, it suffices to show that the integrand stays within the upper half-plane.

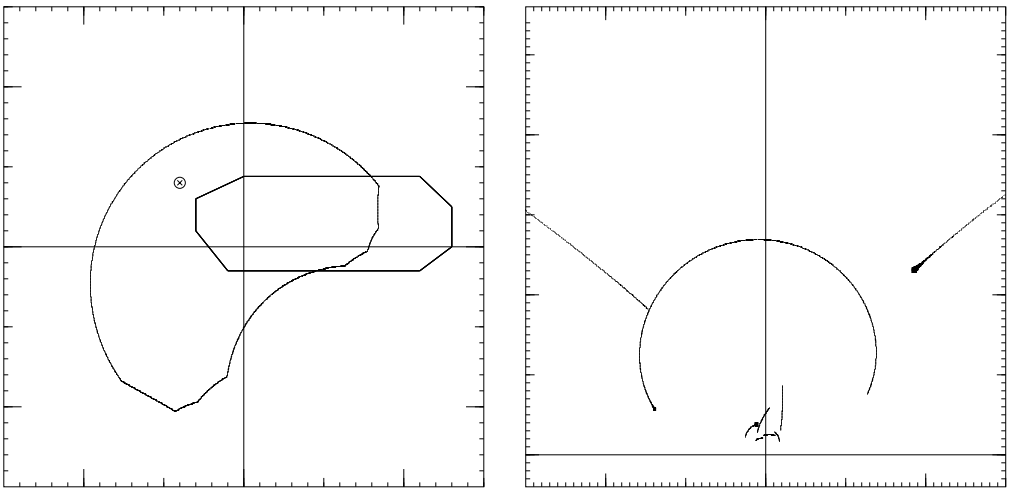

(a) Image of the Fundamental Domain. (b) Values of the Integrand.

Figure 4.18: The Variation of the Argument.

We parameterise each edge of the polygon by $z(t) = P_n + t(P_{n+1} - P_n)$, where $1 \leq n \leq 8$ and we identify P_9 with the vertex P_1. Then we take $z'(t) = (P_{n+1} - P_n)$ along the edge connecting P_n to P_{n+1}.

The point about which we will evaluate the total change of argument lies within the image of the domain under the map V. In fact, we take $w = -0.4 + i0.4$. This point is indicated by a cross in Figure 4.18(a).

The program determines the value of $z'(t)$ along each of the edges of the polygon and then performs a lengthy calculation to evaluate the integrand on each of the rectangles making up a coarse curve that covers the boundary of the domain, confirming that the integrand always lies within the upper half-plane.

(In fact, matters are somewhat more complicated than indicated above: it turned out that at certain positions on the boundary of the fundamental domain, the numerics were extremely badly behaved. This made it difficult to obtain useful bounds. The solution was to sub-divide some of the line segments into smaller pieces and to use a much greater number of grains, i.e. smaller-sized grains, in regions where there was a problem. This increases the time taken for the computation considerably, but eventually a suitable sub-division was found which enabled the proof to be completed. The computation time for this last condition was of the order of a week.)

4.6 Conclusions

In order to apply the results obtained in this chapter to complex dynamics, we need information on the set of analytic functions attracted by Widom's fixed point. (In particular, it would be nice if this set was not empty!) To that end, we would need some control over the stable manifold at Widom's fixed point: we would need to prove that the standard quadratic function

$$f(z) = \frac{e^{2\pi i\omega}}{2}z^2 + 1 - \frac{e^{2\pi i\omega}}{2},$$

is contained in it, and that the complex one parameter family $t \mapsto f_t$ with

$$f_t(z) = \frac{t}{2}z^2 + 1 - \frac{t}{2}$$

intersects it transversally.

Unfortunately, we have not been able to achieve anything in this direction. We could not even prove that the fixed point is hyperbolic. (The reason is that the derivative of the operator is too far away from being diagonal, with respect to the basis we were using.) To prove that the standard quadratic function is attracted would probably be even more difficult.

Experimentally, however, the situation is clear: The fixed point is hyperbolic, the stable manifold has real codimension 2 in the subspace of even commuting maps (see [Sti93b, Section 1]), the standard quadratic function is within the domain of attraction, and the above intersection is transversal. Taking this for granted, it is possible to prove, using the method in [Sti93a], that theorem 49 and corollary 50 hold for functions

- which are even,

- which have an indifferent fixed point with multiplier $e^{2\pi\omega i}$, and
- which are close enough to the quadratic map (with respect to the l_1 norm).

(Here, the assumption of evenness is a bit too restrictive. It would be enough to require a non-degenerate critical point at the origin.)

We have not given up the hope that we might finally be able to prove this result. Such a proof would involve finding a suitable basis such that the derivative of the renormalization operator may be "contracted" to a matrix that is sufficiently close to being diagonal that bounds on the spectrum may be directly read off. The idea is that a "contracted matrix" is a set of matrices in the same way that an interval is a set of real numbers. In particular, it turns out that if a contracted matrix C contains a linear operator L and λ is a complex eigenvalue of L, then the rectangle given by $\det(C - \lambda)$ must contain zero. This is the basic idea behind rigorous computation of the spectrum.

Chapter 5

Hausdorff Dimension of Boundary Curves

5.1 Introduction

Chapter 2 described a proof which yielded a (critical) renormalization fixed point for the Siegel disc operator (using the rigorous computer framework described in Chapter 3). Chapter 4 then established the existence of a Siegel disc, having a boundary curve with a Hölder continuous parameterization ϕ, for any map attracted to the fixed point. The aim of this chapter is to find rigorous bounds for the Hausdorff dimension of these boundary curves. (An extension to this project would be to find rigorous bounds for the *multi-fractal dimensions* [Fal90] of the invariant measure on the boundary.)

Recall that the renormalization fixed point has a (universal) invariant curve through the origin. The first step is to examine the dimension of this invariant curve, it should then be possible to extend the results to the Siegel disc boundaries of maps *attracted* to the fixed point.

In chapter 4, the existence of the universal invariant curve was deduced by applying the *necklace construction* of [Sti93c], in which a sequence of sets (the domain pairs) is constructed iteratively by applying the maps of the fixed point to their domains. Consequently, the invariant curve (or, rather, a certain piece of it) may be written as the invariant set of an Iterated Function System (IFS). To see this, recall the definition of the domain pairs given in section 4.2.2, in particular equation 4.10b:

$$(\mathcal{M}, \mathcal{N}) = (\mathcal{T}\mathcal{B}^{-1}\mathcal{N},\ \mathcal{B}^{-1}(\mathcal{M} \cup \mathcal{N})). \tag{5.1}$$

This may be regarded as defining the fixed point of an IFS made from the maps on the right-hand side (acting on some initial pair of domains). This form of IFS will be called a *Fibonacci system*, by analogy with the definition of the Fibonacci numbers and strings.

Remark 84
It is worth noting that the 'Fibonacci system' given above is not actually an IFS in the usual sense. Ordinarily an IFS consists of a single initial set, along with a set of maps. At each step, the next set is formed from the union of the images of the previous set under

all *of the maps. In contrast, the Fibonacci system has* a pair *of initial domains and a particular combination of the maps is applied at each stage to produce a new pair of sets.*

Alternatively, we may substitute for M in the equation for N (see equation 4.10b) to obtain

$$\mathcal{N} = \mathcal{B}^{-1}(\mathcal{T}\mathcal{B}^{-1}\mathcal{N} \cup \mathcal{N})) = (\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\mathcal{N}) \cup (\mathcal{B}^{-1}\mathcal{N}). \tag{5.2}$$

This second form describes the fixed point of an IFS consisting of the two maps $\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$ and $\mathcal{B}^{-1}$ and some initial (single) domain. It will be called a *Binary IFS*, as each generation of the IFS consists of 2^j *binary patches* (each being an image of the initial domain under a suitable 'binary' composition of the two maps).

Remark 85
The 'binary' formulation is an IFS in the conventional sense. There is a single initial domain, namely $\mathcal{N}$, and both maps are applied at each stage. The union of the images of the current set under both maps produces the new set.

The advantage of expressing (a piece of) the universal curve in this way is that the dimension of the invariant set of an IFS can often be estimated in a relatively straightforward manner by examining the maps of the IFS themselves. The constitutive maps of the IFS used here (consisting of compositions of the maps of the renormalization fixed point, along with suitable rescalings) are analytic, respectively anti-analytic; in particular, they are local similarities. This has a great advantage in that we can use the machinery described in Falconer [Fal90], which mainly deals with similarities, almost unchanged. (The pertinent theory will be described briefly below, after which it will be adapted and applied to the problem at hand.)

This approach will yield both lower and upper bounds for the Hausdorff dimension. The only substantial condition to verify (for the case of lower bounds) is the so-called "open set condition" of [Fal90, §9.2]. The full details will be given in the rest of this chapter.

5.1.1 Dimensions of Self-similar Sets

Many fractals (and, in particular, the critical invariant curve) exhibit some kind of self-similarity. The self-similarity may often be used to *define* the fractal, by means of a formulation known as an *Iterated Function System* (IFS).

The idea is that the fractal (say N) may be assembled from a "collage" of smaller (perhaps distorted) copies of itself. Each "copy" is obtained from the whole object by applying some transformation, say ψ_i, with the fractal itself being invariant under the set of transformations, i.e.

$$\mathsf{N} = \bigcup_{i=1}^{m} \psi_i(\mathsf{N}). \tag{5.3}$$

The idea is that, given some initial set N_0, we perform the iteration

$$\mathsf{N}_{k+1} = \bigcup_{i=1}^{m} \psi_i(\mathsf{N}_k), \tag{5.4}$$

with the hope that, in some sense, $N_k \to N$ as $k \to \infty$. It follows [Fal90, Theorem 9.1] that such a family of transformations will define a unique (non-empty) compact invariant set provided that they are *contractions*.

This approach has the advantage that bounds for the Hausdorff dimension of the invariant set may often be obtained by examining the defining transformations (in this case, the ψ_i).

Similarities In the case where the transformations are all *similarities* (i.e. the invariant set is exactly self-similar) and where the union $\cup_{i=1}^{m}\psi_i(N)$ is "almost disjoint" it follows [Fal90, 9.2] that the Hausdorff dimension may be obtained by solving the *partition equation*:

$$\sum_{i=1}^{m} c_i^s = 1,$$

for s, where the c_i are the contractivity rates of the contractions, i.e.

$$|\psi_i(x) - \psi_i(y)| = c_i|x - y|$$

for all points x, y in some domain containing N.

The Open Set Condition The requirement that the union be "almost disjoint" in the above needs to be clarified. What is needed is that the components $\psi_i(N)$ of the "collage" do not overlap too much. This is the *open set condition*. More formally, we need to be able to find some non-empty bounded open set that contains its images under the defining transformations, and for which the union of the images is disjoint.

Definition 86 (Open Set Condition)
If there exists an open set Ω such that

$$\Omega = \bigcup_{i=1}^{m} \psi_i(\Omega), \tag{5.5}$$

where the union on the right is disjoint (i.e. $\psi_i(\Omega) \cap \psi_j(\Omega) = \emptyset$ for $i \neq j$), then the IFS is said to satisfy the open set condition.

The above gives the dimension for an invariant set defined by *similarities*. The maps defining the necklace curve are not similarities. However (see [Fal90, 9.3]), the above results may be extended to give an *upper bound* for the Hausdorff dimension in the more general case where the maps ψ_i are contractions. In this case, the solution of the partition equation gives an upper bound on the dimension and the open set condition is *not* required. The key idea is that the invariant set is also invariant for the m^k possible compositions of the form $\psi_{i1} \circ \cdots \circ \psi_{ik}$, and that for high values of k these compositions are, in some sense, close to being similarities.

The above technique will be used to obtain an upper bound for the Hausdorff dimension of the universal curve. In fact, we will also use a further extension of the technique discussed

above to yield a *lower* bound for the dimension (a similar result was presented in [Fal90], but it relied on the limit set N being totally disconnected, which is not the case here). For lower bounds, the open set condition is required. These considerations are presented in section 5.5. Firstly, we discuss the IFS formulation for the necklace.

5.1.2 The Iterated Function Systems

The situation is somewhat more complicated than indicated above: although there is essentially only *one* universal curve, there are four different ways in which it may be conveniently expressed as the limit set of an IFS. To see this, recall that the existence proof itself does not *directly* yield a fixed point equation, but rather what we decided to call an *uncle* of it, namely equation 2.16:

$$\begin{aligned} \mathsf{U}(z) &= \mathsf{C}\alpha^{-1}\mathsf{V}\alpha^{2}\mathsf{C} \quad \text{on } \mathrm{D}_{\mathsf{U}}, \\ \mathsf{V}(z) &= \mathsf{C}\alpha^{-1}\mathsf{VQU}\alpha^{2}\mathsf{C} \quad \text{on } \mathrm{D}_{\mathsf{V}}, \end{aligned}$$

From this 'uncle', *two* different fixed point equations may be deduced. They will be called the *regular* one, and the *singular* one: the regular fixed point equation is the one containing the composition VQUQ ($=\mathsf{FE}$), whereas the singular one contains QVQU (both contain the same composition as the uncle, namely VQU).

An additional 'degree of freedom' is added to the problem by the presence of the two different *composition orders*. Recall that the existence proof used the *inverse* composition order (establishing the validity of a VQU uncle equation) because the bounds produced by the alternative order were not sufficiently tight to give a rigorous proof. In contrast, the natural composition order to use in the necklace construction was the *accretive* order (UQV uncle) because this enabled us to use the combinatorial properties of the Fibonacci strings. Recall that in order to use the former (inverse) fixed point in the latter (accretive) necklace construction, it was therefore necessary to take *analytic continuations* of the maps involved (see 4.3.2), which established the validity of the UQV uncle from its VQU counterpart.

These two 'degrees of freedom' are *orthogonal*: i.e. the choice of an *inverse or accretive* uncle may be made independently of the choice of a *regular or singular* fixed point equation. This means that the alternatives may be combined to yield *four* different fixed point equations. Correspondingly, there are *four* different IFSs that may be used. These will be denoted as follows:

1. Accretive Regular
2. Accretive Singular
3. Inverse Regular
4. Inverse Singular

The following subsection demonstrates how an IFS arises for the inverse composition order, after which the full details of the different IFSs will be given.

5.1.3 The Inverse ($\mathcal{TU}$) Necklace

Section 5.1 demonstrated how the accretive necklace construction gives rise to the accretive binary IFS. In order to see how an IFS arises for the inverse composition order, $\mathcal{TU}$, we look at the corresponding "inverse" form of the necklace construction.

Heuristics

Let us first do it in a heuristic way. We start with the usual $\mathcal{UT}$ necklace construction:

$$\begin{aligned} \mathcal{M} &= \mathcal{TB}^{-1}\mathcal{N} \\ \mathcal{N} &= \mathcal{B}^{-1}(\mathcal{M} \cup \mathcal{N}). \end{aligned}$$

Rewrite this as

$$\begin{aligned} \mathcal{M} &= \mathcal{TB}^{-1}\mathcal{N} \\ \mathcal{N} &= \mathcal{B}^{-1}(\mathcal{TN} \cup \mathcal{UM}). \end{aligned}$$

(The equation $\mathcal{M} \cup \mathcal{N} = \mathcal{TN} \cup \mathcal{UM}$ follows by looking at the parametrisation $\phi(t)$.) It follows that

$$\begin{aligned} (\mathcal{TUM}) &= \mathcal{TUTB}^{-1}\mathcal{N} = \mathcal{B}^{-1}(\mathcal{UTN}) \\ (\mathcal{UTN}) &= \mathcal{UTB}^{-1}\mathcal{TN} \cup \mathcal{UTB}^{-1}\mathcal{UM} = \mathcal{UB}^{-1}(\mathcal{UTN}) \cup \mathcal{B}^{-1}(\mathcal{TUM}). \end{aligned}$$

This, in turn, implies that the sets $\tilde{\mathcal{M}} := \mathcal{TUM}$, $\tilde{\mathcal{N}} := \mathcal{UTN}$ satisfy the fixed point equation

$$\tilde{\mathcal{M}} = \mathcal{B}^{-1}\tilde{\mathcal{N}} \tag{5.6a}$$

$$\tilde{\mathcal{N}} = \mathcal{B}^{-1}\tilde{\mathcal{M}} \cup \mathcal{UB}^{-1}\tilde{\mathcal{N}}. \tag{5.6b}$$

This is the basic form of the necklace construction corresponding to the $\mathcal{TU}$ fixed point equation

$$\begin{aligned} \mathcal{U} &= \mathcal{BTB}^{-1} \\ \mathcal{T} &= \mathcal{BTUB}^{-1}. \end{aligned}$$

Continuous Extension of the Necklace Curve

The above discussion suggests that we define

$$\begin{aligned} \tilde{\mathcal{M}} &= \phi([0, \omega]) \\ \tilde{\mathcal{N}} &= \phi([-1, 0]). \end{aligned}$$

However, there might be a problem with the second equation since ϕ is not defined, a priori, for $t < -\omega$ (recall that, for the usual necklace construction, $\mathcal{L} = \phi([-\omega, 1])$). We

therefore have to extend ϕ in some way. Recall that

$$\begin{aligned} \mathcal{U}\phi(t) &= \phi(t-1), \quad 1-\omega \le t \le 1 \\ \mathcal{T}\phi(t) &= \phi(t+\omega), \quad -\omega \le t \le 1-\omega \\ \mathcal{B}^{-1}\phi(t) &= \phi(-\omega t), \quad -\omega \le t \le 1, \end{aligned}$$

and, in particular,

$$\phi(t) = B\phi(-\omega t), \quad -\omega \le t \le 1.$$

The right hand side is defined for $-\frac{1}{\omega} \le t \le 1$, and as a result we can use it to continuously extend ϕ to the interval $[-1, 1]$.

With this extension, the above definitions of $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{N}}$ are justified and can be written as

$$\begin{aligned} \tilde{\mathcal{M}} &= \phi([0, \omega]) \\ \tilde{\mathcal{N}} &= B\phi([0, \omega]). \end{aligned}$$

It is now necessary to verify that $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{N}}$ actually satisfy the inverse ($\mathcal{TU}$) necklace equations. Firstly,

$$\mathcal{B}^{-1}\tilde{\mathcal{N}} = \phi([0, \omega]) = \tilde{\mathcal{M}}.$$

For the second equation, observe that

$$\mathcal{B}^{-1}\tilde{\mathcal{M}} = \mathcal{B}^{-1}\phi([0, \omega]) = \phi([-\omega^2, 0]),$$

and that

$$\begin{aligned} \mathcal{U}\mathcal{B}^{-1}\tilde{\mathcal{N}} &= \mathcal{U}\phi([0, \omega]) \\ &= \mathcal{B}\mathcal{T}\mathcal{B}^{-1}\phi([0, \omega]) \\ &= \mathcal{B}\mathcal{T}\phi([-\omega^2, 0]) \\ &= \mathcal{B}\phi([-\omega^2 + \omega, \omega]) \\ &= \phi([-1, \omega - 1]) \\ &= \phi([-1, -\omega^2]). \end{aligned}$$

Finally, from equation 5.6, we obtain the fixed point equation for the inverse binary IFS:

$$\tilde{\mathcal{N}} = \mathcal{B}^{-2}\tilde{\mathcal{N}} \cup \mathcal{U}\mathcal{B}^{-1}\tilde{\mathcal{N}}.$$

The next section will clarify the distinction between the 'regular' and 'singular' flavours of IFSs, after which the full details of each of the four possibilities for an IFS will be given.

5.2 Regular and Singular IFS

The aim of this section is to clarify the relationship between the regular and singular varieties. Consider, for example, the translation of the necklace construction given in

Chapter 4. The result was a pair of functions U and V, defined on domains Δ_U and Δ_V, which satisfy the 'accretive uncle' equation:

$$U = C\frac{1}{\alpha}V\alpha^2C, \tag{5.7a}$$

$$V = C\frac{1}{\alpha}UQV\alpha^2C, \tag{5.7b}$$

with $\alpha = V(0)$.

From this equation, the following two different fixed point equations may be deduced by squaring:

$$QU = C\frac{1}{\alpha^2}QV\alpha^2C \tag{5.8a}$$

$$QV = C\frac{1}{\alpha^2}QUQV\alpha^2C \tag{5.8b}$$

and

$$UQ = C\frac{1}{\alpha}VQ\alpha C \tag{5.9a}$$

$$VQ = C\frac{1}{\alpha}UQVQ\alpha C. \tag{5.9b}$$

The former holds on the domains Δ_U and Δ_V, the latter on the "square root" of these, namely on domains $\tilde{\Delta}_U$ and $\tilde{\Delta}_V$ such that $Q\tilde{\Delta}_U = \Delta_U$ and $Q\tilde{\Delta}_V = \Delta_V$.

The latter is the one we are actually interested in, since it corresponds to $(E, F) = (UQ, VQ)$. It describes the *scaling limit about the critical point* of the Siegel disc boundary, i.e., it is the fixed point equation for the universal pair obtained when renormalising the standard quadratic map with respect to its critical point. This is called the *singular* scaling limit.

It is plausible that the former one describes the *scaling limit about a point of the forward orbit* of the critical point, i.e., the universal pair obtained when renormalising the standard quadratic map with respect to a point of the forward orbit of its critical point. (This has not been proved, however.) This is the *regular* scaling limit.

Notice that both (5.8) and (5.9) are of the same shape:

$$\mathcal{U} = \mathcal{B}\mathcal{T}\mathcal{B}^{-1} \tag{5.10a}$$

$$\mathcal{T} = \mathcal{B}\mathcal{U}\mathcal{T}\mathcal{B}^{-1}. \tag{5.10b}$$

The "generic" notation (i.e., using $\mathcal{U}$, $\mathcal{T}$, $\mathcal{B}$, $\mathcal{M}$, $\mathcal{N}$, and $\mathcal{L}$) will be used where possible, because it is simpler.

5.2.1 Square Root of the Necklace

Recall that $(E, F) := (UQ, VQ)$. If we put $(\hat{E}, \hat{F}) := (QU, QV)$, then it may be helpful to bear in mind the following commutative diagrams:

$$\begin{array}{ccc} z & \xrightarrow{Q} & w \\ {\scriptstyle E}\downarrow & & \downarrow{\scriptstyle \hat{E}} \\ z' & \xrightarrow{Q} & w' \end{array} \qquad \begin{array}{ccc} z & \xrightarrow{Q} & w \\ {\scriptstyle F}\downarrow & & \downarrow{\scriptstyle \hat{F}} \\ z' & \xrightarrow{Q} & w' \end{array} \tag{5.11}$$

(Notice that U and V both map w to z'.)

If $t \mapsto \chi(t)$ is the parameterization of the necklace curve for equation 5.8 (regular) and $t \mapsto \phi(t)$ is the parameterization for equation 5.9 (singular), then we have

$$\phi(t) \equiv (\chi(t))^2,$$

as shown in section 4.3.7 (which demonstrated that $\chi(t)$ is a continuous square root of $\phi(t)$).

Since, apart from at the origin, squaring is a local similarity, it follows (by Corollary 2.4 of Falconer [Fal90]) that the Hausdorff dimensions of these curves are equal. It is computationally much more convenient to work with the regular IFS, primarily because the singular IFS uses the "square roots" of the domains. Thus, for the calculation of the dimensions, we shall be using the regular IFS and, according to the above remarks, the results will hold for the singular IFS as well.

Recall that (in section 4.5.9) we proved that the map V is a homeomorphism on an open set containing $N = \phi([-\omega, 1-\omega])$. (It was not actually verified there that the "coffin" domain really contains the whole set N. This is done before the corresponding results are used.) Because U is just a rescaled version of V, it follows that U is a homeomorphism on a neighbourhood of $M = \phi([1-\omega, 1])$. Since these maps are also analytic, they are *bi-Lipschitz*.

Notice, then, that the curve χ has been pieced together from two bi-Lipschitz images of pieces of the curve ϕ. By [Fal90, Corollary 2.4], the Hausdorff dimensions of these curves are equal. (This process of piecing together also preserves the natural measure. Therefore, all the dynamical dimensions of both curves will agree, which would be useful if this project were extended to examine them.)

5.2.2 Necklace Curves as Limit Sets of IFSs

Letting $j \to \infty$ in equation 4.11, we obtain (using the generic notation)

$$\mathcal{L} = \mathcal{B}^{-1}\mathcal{L} \cup \mathcal{T}\mathcal{B}^{-2}\mathcal{L},$$

i.e., the necklace set $\mathcal{L}$ itself (rather than just the piece $\mathcal{N}$) is the limit set of an IFS. Multiplying this equation by $\mathcal{B}^{-1}$ and using $\mathcal{N} = \mathcal{B}^{-1}\mathcal{L}$, we get the binary IFS,

$$\mathcal{N} = \mathcal{B}^{-1}\mathcal{N} \cup \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\mathcal{N},$$

introduced earlier. The maps $\mathcal{TB}^{-2}$ and $\mathcal{B}^{-1}\mathcal{TB}^{-1}$ are conjugate by $\mathcal{B}$; it is a matter of taste with which IFS to work. In particular, the contractivity of $\mathcal{TB}^{-2}$ on $\mathcal{L}$ is equal to the contractivity of $\mathcal{B}^{-1}\mathcal{TB}^{-1}$ on $\mathcal{N}$.

5.3 The four IFSs in detail

In what follows, it is useful to keep in mind that the accretive and the inverse invariant curve are different pieces of the same universal curve. The accretive curve is the piece parameterised between $-\omega$ and 1. The argument given in the previous section showed that the inverse one is the piece between -1 and ω. The details of the various IFSs are now given.

5.3.1 Inverse Order (VQU)

Domains The basic domains are D_U and D_V, as used in the existence proof (c.f. Chapter 2).

Uncle The inverse uncle resulted from the existence proof (c.f. equation 2.16):

$$U = C\alpha^{-1}V\alpha^2C, \quad \text{on } D_U \tag{5.12a}$$
$$V = C\alpha^{-1}VQU\alpha^2C, \quad \text{on } D_V \tag{5.12b}$$

Singular Scaling Limit (VQUQ):

Domains The singular scaling limit takes place on the 'square root domains' $\sqrt{D_U}$ and $\sqrt{D_V}$ (note the position of squaring, Q, in the corresponding fixed point equation, below):

$$\sqrt{D_U} := \{z : z^2 \in D_U\} \tag{5.13a}$$
$$\sqrt{D_V} := \{z : z^2 \in D_V\} \tag{5.13b}$$

Fixed Point Equation

$$UQ = C\alpha^{-1}VQ\alpha C, \quad \text{on } \sqrt{\Delta_U} \tag{5.14a}$$
$$VQ = C\alpha^{-1}VQUQ\alpha C, \quad \text{on } \sqrt{\Delta_V} \tag{5.14b}$$

Invariant Curve We denote the invariant curve by $\tilde{\chi}(t)$. It is characterised by

$$UQ\tilde{\chi}(t) = \tilde{\chi}(t-1), \quad 0 \le t \le \omega \tag{5.15a}$$
$$VQ\tilde{\chi}(t) = \tilde{\chi}(t+\omega), \quad -1 \le t \le 0 \tag{5.15b}$$
$$\alpha C\tilde{\chi}(t) = \tilde{\chi}(-\omega t), \quad -1 \le \omega. \tag{5.15c}$$

Invariant Pair We denote the invariant pair of sets by $(\tilde{M}_S, \tilde{N}_S)$:

$$\tilde{M}_S = \tilde{\chi}([0,\omega]) \tag{5.16a}$$
$$\tilde{N}_S = \tilde{\chi}([-1,0]). \tag{5.16b}$$

Fibonacci System

$$\tilde{M}_S = \alpha C\tilde{N}_S \tag{5.17a}$$
$$\tilde{N}_S = \alpha C\tilde{M}_S \cup UQ\alpha C\tilde{N}_S. \tag{5.17b}$$

Binary IFS

$$N_S = |\alpha|^2 N_S \cup UQ\alpha CN_S. \tag{5.18}$$

Regular Scaling Limit ($QVQU$):

Domains The regular scaling limit takes place on the base domains D_U and D_V.

Fixed Point Equation

$$QU = C\alpha^{-2}QV\alpha^2 C, \quad \text{on } \Delta_U \tag{5.19a}$$
$$QV = C\alpha^{-2}QVQU\alpha^2 C, \quad \text{on } \Delta_V \tag{5.19b}$$

Invariant Curve We denote the invariant curve by $\tilde{\phi}(t)$. It is characterised by

$$QU\tilde{\phi}(t) = \tilde{\phi}(t-1), \quad 0 \le t \le \omega \tag{5.20a}$$
$$QQ\tilde{\phi}(t) = \tilde{\phi}(t+\omega), \quad -1 \le t \le 1 \tag{5.20b}$$
$$\alpha^2 C\tilde{\phi}(t) = \tilde{\phi}(-\omega t), \quad -1 \le t \le \omega. \tag{5.20c}$$

Invariant Pair We denote the invariant pair of sets by $(\tilde{M}_R, \tilde{N}_R)$:

$$\tilde{M}_R = \tilde{\phi}([0,\omega]) \tag{5.21a}$$
$$\tilde{N}_R = \tilde{\phi}([-1,0]). \tag{5.21b}$$

Fibonacci System

$$\tilde{M}_R = \alpha^2 C\tilde{N}_R \tag{5.22a}$$
$$\tilde{N}_R = \alpha^2 C\tilde{M}_R \cup QU\alpha^2 C\tilde{N}_R. \tag{5.22b}$$

Binary IFS

$$N_R = |\alpha|^4 N_R \cup QU\alpha^2 CN_R. \tag{5.23}$$

5.3.2 Accretive Order (UQV)

Domains The base domains are Δ_{U} and Δ_{V} as defined in Chapter 4 and used to verify the necklace hypotheses.

Uncle The accretive uncle was obtained by analytically extending the maps U and V to the domains Δ_{U} and Δ_{V} (see chapter 4):

$$\mathsf{U} = \mathsf{C}\alpha^{-1}\mathsf{V}\alpha^2\mathsf{C}, \quad \text{on } \Delta_{\mathsf{U}} \tag{5.24a}$$
$$\mathsf{V} = \mathsf{C}\alpha^{-1}\mathsf{UQV}\alpha^2\mathsf{C}, \quad \text{on } \Delta_{\mathsf{V}} \tag{5.24b}$$

Singular Scaling Limit (UQVQ):

Domains The singular scaling limit takes place on the "square root domains" $\sqrt{\Delta_{\mathsf{U}}}$ and $\sqrt{\Delta_{\mathsf{V}}}$:

$$\sqrt{\Delta_{\mathsf{U}}} := \{z : z^2 \in \Delta_{\mathsf{U}}\} \tag{5.25a}$$
$$\sqrt{\Delta_{\mathsf{V}}} := \{z : z^2 \in \Delta_{\mathsf{V}}\} \tag{5.25b}$$

Fixed Point Equation

$$\mathsf{UQ} = \mathsf{C}\alpha^{-1}\mathsf{VQ}\alpha\mathsf{C}, \quad \text{on } \sqrt{\Delta_{\mathsf{U}}} \tag{5.26a}$$
$$\mathsf{VQ} = \mathsf{C}\alpha^{-1}\mathsf{UQVQ}\alpha\mathsf{C}, \quad \text{on } \sqrt{\Delta_{\mathsf{V}}} \tag{5.26b}$$

Invariant Curve We denote the invariant curve by $\chi(t)$. It is characterised by

$$\mathsf{UQ}\chi(t) = \chi(t-1), \quad 1-\omega \le t \le 1 \tag{5.27a}$$
$$\mathsf{VQ}\chi(t) = \chi(t+\omega), \quad -\omega \le t \le 1-\omega \tag{5.27b}$$
$$\alpha\mathsf{C}\chi(t) = \chi(-\omega t), \quad -\omega \le t \le 1. \tag{5.27c}$$

Invariant Pair We denote the invariant pair of sets by $(\mathsf{M}_{\mathsf{S}}, \mathsf{N}_{\mathsf{S}})$:

$$\mathsf{M}_{\mathsf{S}} = \chi([1-\omega, 1]) \tag{5.28a}$$
$$\mathsf{N}_{\mathsf{S}} = \chi([-\omega, 1-\omega]). \tag{5.28b}$$

Fibonacci System

$$\mathsf{M}_{\mathsf{S}} = \mathsf{VQ}\alpha\mathsf{C}\mathsf{N}_{\mathsf{S}} \tag{5.29a}$$
$$\mathsf{N}_{\mathsf{S}} = \alpha\mathsf{C}(\mathsf{M}_{\mathsf{S}} \cup \mathsf{N}_{\mathsf{S}}) \tag{5.29b}$$

Binary IFS

$$\mathsf{N}_{\mathsf{S}} = \alpha\mathsf{C}\mathsf{N}_{\mathsf{S}} \cup \alpha\mathsf{C}\mathsf{V}\mathsf{Q}\alpha\mathsf{C}\mathsf{N}_{\mathsf{S}}. \tag{5.30}$$

Regular Scaling Limit ($QUQV$):

Domains The regular scaling limit takes place on the domains Δ_U and Δ_V.

Fixed Point Equation

$$QU = C\alpha^{-2}QV\alpha^2C, \quad \text{on } \Delta_U \tag{5.31a}$$
$$QV = C\alpha^{-2}QUQV\alpha^2C, \quad \text{on } \Delta_V \tag{5.31b}$$

Invariant Curve We denote the invariant curve by $\phi(t)$. It is characterised by

$$QU\phi(t) = \phi(t-1), \quad 1-\omega \le t \le 1 \tag{5.32a}$$
$$QV\phi(t) = \phi(t+\omega), \quad -\omega \le t \le 1-\omega \tag{5.32b}$$
$$\alpha^2C\phi(t) = \phi(-\omega t), \quad -\omega \le t \le 1. \tag{5.32c}$$

Invariant Pair We denote the invariant pair of sets by (M_R, N_R):

$$M_R = \phi([1-\omega, 1]) \tag{5.33a}$$
$$N_R = \phi([-\omega, 1-\omega]). \tag{5.33b}$$

Fibonacci System

$$M_R = QV\alpha^2CN_R \tag{5.34a}$$
$$N_R = \alpha^2C(M_R \cup N_R) \tag{5.34b}$$

Binary IFS

$$N_R = \alpha^2CN_R \cup \alpha^2CQV\alpha^2CN_R. \tag{5.35}$$

5.3.3 Discussion

In principle, the presence of four different (binary) IFS formulations may actually confer an advantage: the most suitable IFS formulation can be chosen for a given task. For example, one of the IFS's may be better suited to getting upper bounds on the dimension, another may be better suited to getting lower bounds. In addition, the inverse flavour is computationally more convenient to implement (since it is possible to work directly with the maps U and V on simple disc domains, rather than having to use their analytic continuations on a union of discs).

There are some problems, however: it is necessary to prove that the Hausdorff dimensions for the invariant sets of the IFSs are the same. (If the project were extended, it would also be necessary to prove that their dynamical dimensions are the same.) Fortunately, this is not too difficult (it may be done using [Fal90, Corollary 2.4], which says that Lipschitz maps do not increase the Hausdorff dimension. This will be done in section 5.2). In addition, for the computation of *lower* bounds, the *open set condition* must be verified for any IFS that is used. This last condition proves to be rather more difficult.

Fibonacci Systems and Binary IFSs: It is worth clarifying the distinction between a 'binary IFS' and a 'Fibonacci system'. The Fibonacci system formulation is closest in spirit to the original necklace construction (in fact, for the accretive case they are identical). In particular (for lower bounds) it is possible to use the details of the necklace construction to prove that the *open set condition* holds for the Fibonacci systems (this is done in section 5.6). However, the Fibonacci systems are not genuine IFSs (c.f. remarks 84 and 85). In order to apply the relevant results on IFSs to obtain bounds on the dimension, we must use the binary IFSs or look at dimensions for more generalised IFS constructions. It turns out that the open set condition can be carried across from the Fibonacci system to the Binary IFS in the case where the accretive composition order is used. In the inverse case, however, no proof was found.

Since the inverse composition order is much easier to use (and in the hope that the open set condition might eventually be verified to yield lower bounds), the first results that will be presented in this chapter pertain to that order. Using the inverse order first also provided a way of gaining some experience in programming the method. Later results attempt to use the accretive IFSs. In the latter case the programming and calculations are correspondingly more difficult and time-consuming.

5.4 Elasticity of the Necklace

In order to estimate the Hausdorff dimension, we use essentially the same approach as Falconer [Fal90], in that we obtain upper bounds for the *contractivities* and lower bounds for the *coercivities* of the constituent maps of the IFS evaluated on a partition of the necklace curve. By solving the so-called "partition equations" we then obtain upper and lower bounds for the dimension. The aim of this section is to establish some definitions and to develop the techniques used for estimating the contractivities and coercivities. It will also show that, in addition to being *Hölder continuous*, the necklace is also *Hölder coercive* (a term which will be defined later in this section). These two facts will enable us to get some preliminary bounds for the dimension. (We say that the properties of contractivity and coercivity determine the *elasticity* of the necklace.)

5.4.1 Contractivity and Coercivity

First, we recall the following definitions. In the following let X, Y be normed spaces and let f be a map from X to Y.

Definition 87 (Lipschitz)
The map f is called Lipschitz *if there exists a constant $\ell > 0$ such that*

$$|f(x) - f(y)| \leq \ell |x - y| \tag{5.36}$$

for all $x, y \in X$, and ℓ is called a Lipschitz *constant.*

(Thus a Lipschitz function is a Hölder continuous function with Hölder exponent equal to one.)

Definition 88 (Contractive)
If $\ell < 1$ for a Lipschitz function f, then f is called contractive, *and ℓ is called a* contractivity *constant.*

Definition 89 (Coercive)
The map f is called coercive *if there exists a constant $r > 0$ such that*

$$|f(x) - f(y)| \geq r|x - y| \tag{5.37}$$

for all $x, y \in X$. The number r is called a coercivity *constant.*

It follows readily that the inverse of a coercivity constant of f is a Lipschitz constant of its inverse map. On the other hand, the inverse of a Lipschitz constant of an invertible map is a coercivity constant of its inverse map.

Provided that the map f is *differentiable*, then contractivity and coercivity constants can be calculated from bounds on its derivative, as the following lemmas illustrate.

Lemma 90 (Contractivity)
Let X be a Banach space, and let f be a continuously differentiable map from an open set Ω of X to X. Then, if A is a compact convex *subset of Ω, the quantity*

$$\max_{x \in A} |Df(x)| \tag{5.38}$$

is a Lipschitz constant of f on A.

Proof This follows directly from the mean value theorem. □

Lemma 91 (Coercivity)
(Let f be as above.) Let A be a compact subset of Ω, and assume that f is injective on A and that $Df(x)$ is regular (i.e., that $Df(x)^{-1}$ exists) at each point x of A. If $f(A)$ is convex, *then the quantity*

$$(\max_{x \in A} |Df(x)^{-1}|)^{-1} \tag{5.39}$$

is a coercivity constant of f.

Proof Put $y = f(x)$, and observe that

$$Df^{-1}(y) = (Df(x))^{-1}. \tag{5.40}$$

The lemma then follows by the proof of the previous one. □

Remark 92 (Bounds on Convex Sets)
The need for bounds to be established over convex *sets in the above is somewhat awkward, since the necklace curve is not convex. Also, recall that the domains of the maps are not convex for every IFS; they are sometimes unions of convex sets (discs). It is enough, however, that the bounds be obtained on some convex set that contains the relevant invariant set (for example, a region bounded by a suitable convex polygon), rather than over the whole domains of definition of the maps.*

Note that the map $\mathcal{B}^{-1}$ is either of the form $\alpha^2 C$ (regular) or of the form αC (singular), and is therefore always an anti-similarity (therefore having contractivity and coercivity constants equal).

The other map to consider is either of the form $\mathcal{T}\mathcal{B}^{-2}$ on the whole necklace set $\mathcal{L}$ (recall section 5.2.2) or $\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$ on the piece $\mathcal{N}$.

Consider, for example, the map $\mathcal{T}\mathcal{B}^{-2}$ on $\mathcal{L}$ in the case of the *regular* necklace. (In the case of the singular necklace, this map has a critical point and is, therefore, not coercive.) This map is analytic. Its derivative, therefore, is given by a single complex number, which means that

$$|Df(x)^{-1}|^{-1} = |Df(x)|, \tag{5.41}$$

for all x in its domain of definition. It follows that a positive *lower* bound for $|Df(x)|$ on the *pre-image* under f of any *convex* set gives a coercivity constant of f on that set. (The need for bounds to be taken on the *pre-image* of a convex set is also somewhat awkward.)

Note that the map $f = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$ (on $\mathcal{N}$) is also analytic, so the above argument applies.

Regular IFS In the case of the curve ϕ, $\mathcal{B}^{-1}$ is the map $\alpha^2 C$, and

$$\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1} = \alpha^2 C Q V \alpha^2 C. \tag{5.42}$$

Both maps are *contractive* on a neighbourhood of $\mathcal{N}$. (For the first map, this is obvious. For the second, this was proved when the contractivity condition was verified for the necklace construction in section 4.5.8.)

Moreover, both maps are *coercive* (see 5.4.1) on a neighbourhood of $\mathcal{N}$: their derivatives are bounded away from zero. Again, this is obvious for the first map, which is an exact (anti)-similarity. In the case of the second, we know that V is an analytic homeomorphism (on the fundamental domain of section 4.5.9). Of course, the derivative of QV will vanish at the "left" endpoint $\phi(-\omega)$ of the necklace, because $V\phi(-\omega) = 0$. But this point does not occur in the set $\alpha^2 C N_R = \mathcal{B}^{-1}\mathcal{N}$ (recall the descriptions of the IFSs in section 5.3). Alternatively, one could argue that the bounds obtained for the contractivity condition of the necklace (section 4.5.8) show directly that the derivative of $\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$ on $\mathcal{N}$ is bounded away from zero.

Remark 93
It is desirable for the constituent maps of an IFS to be coercive in addition to contractive: without coercivity, there would be no hope of obtaining lower bounds for the dimensions.

Singular IFS In the case of the curve χ, $\mathcal{B}^{-1}$ is the map αC, and

$$\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1} = C\alpha V Q \alpha C. \tag{5.43}$$

Numerically, both maps are observed to be contractivities (this has not been proved, but will not be used anyway). However, notice that the map $\mathcal{T}\mathcal{B}^{-1}$ is *not* coercive for the singular IFS; it has a critical point at the origin.

5.4.2 Decay Rate of Patch Size

Recall the proof for the upper bound on the decay rate of the patch size with increasing generation index (Lemma 76 in section 4.3.6): $d_j \leq c\kappa^j$ for some constant c. We now use a similar argument to get a lower bound for the *minimal* patch size. We will use these two facts to establish an "exponential disjointness" condition for the distance between two patches in a generation that are separated by exactly one patch. This will enable us to conclude that the necklace curve is "Hölder coercive" (in addition to being Hölder continuous) and will ultimately give lower bounds for the dimension.

Lemma 94 (Minimal Patch Size)
The following lower bound holds for the minimal *diameter δ_j of a patch of generation j:*

$$\delta_j \geq d\rho^j, \tag{5.44}$$

where d is a positive constant and where ρ is defined as follows: Let σ and τ be positive lower bounds of the infinitesimal contractivities of the maps $\mathcal{B}^{-1}$ and $\mathcal{T}\mathcal{B}^{-2}$:

$$|D\mathcal{B}^{-1}| \geq \sigma, \quad |D\mathcal{T}\mathcal{B}^{-2}| \geq \tau. \tag{5.45}$$

Then

$$\rho = \min(\sigma, \tau^{\frac{1}{2}}). \tag{5.46}$$

The above inequality for ρ obviously implies that $\sigma \geq \rho$ and $\tau \geq \rho^2$. Let δ_j be the minimal diameter of a patch in generation j. We claim that there exists a constant c such that

$$\delta_j \geq d\rho^j. \tag{5.47}$$

Proof We can make the claim true for $j = 0$ and $j = 1$ simply by choosing d small enough. Assume now that $j \geq 0$ and that

$$\begin{aligned} \delta_j &\geq d\rho^j \\ \delta_{j+1} &\geq d\rho^{j+1}. \end{aligned}$$

Recall that $\mathcal{L}_{j+2} = \mathcal{B}^{-1}\mathcal{L}_{j+1} \cup \mathcal{T}\mathcal{B}^{-2}\mathcal{L}_j$ (equation 4.11). Then, by the mean value theorem of differential calculus, we obtain

$$\delta_{j+2} \geq \min(\sigma d\rho^{j+1}, \tau d\rho^j) = \min(d\rho^{j+2}, d\rho^{j+2}) = d\rho^{j+2}, \tag{5.48}$$

and the induction is complete. □

5.4.3 Diameter of Patch Clusters

The purpose of this section is to use the above results to look at the diameter of *clusters* of patches, with the aim of eventually finding a lower bound for the dimension of the necklace.

First, some preliminary definitions:

Definition 95 (Diameter)
For any set A*, let* $|A|$ *denote its* diameter*:*

$$|A| = \sup\{\, |x - y| : x, y \in A \,\}. \tag{5.49}$$

Definition 96 (Distance)
For any two sets A *and* B*, we define their* distance *by*

$$dist(A, B) = \inf\{\, |x - y| : x \in A, y \in B \,\}. \tag{5.50}$$

Let P and Q be patches of the usual necklace construction, of the same generation j. Observe that

$$|P \cup Q| \geq \operatorname{dist}(P, Q), \tag{5.51}$$

i.e. the diameter of the union of two patches is bounded below by their distance. (If P and Q overlap, their distance is zero, and the estimate is vacuous.)

This fact will prove useful. Recall that, using the disjointness conditions, it follows that two patches which are not neighbours do not overlap (section 4.2.6). The proof of this (given in [Sti93c, 7.15]) may be modified to yield a lower bound for the minimal distance of two non-neighbouring patches P and Q, of the form

$$\operatorname{dist}(P, Q) \geq d\rho^j. \tag{5.52}$$

Where ρ is the global coercivity defined above, and d is a positive constant. (This will be proved below.)

In other words, clusters of three patches of generation j have diameter $\geq d\rho^j$.

Consider a cluster of $n + 1$ patches of generation j. Take the boundary patches of this cluster, and go back to their earliest ancestors which are separated by at least one patch. If we must go back k generations, then the diameter of the patch cluster is $\geq d\rho^{j-k}$.

The number k may be estimated by

$$c_1 \left(\frac{1}{\omega}\right)^k \leq n \leq c_2 \left(\frac{1}{\omega}\right)^k, \tag{5.53}$$

which follows from the fact that, for any solution (a_k, b_k) of the Fibonacci recursion $a_{k+1} = b_k$, $b_{k+1} = a_k + b_k$, the sum $a_k + b_k$ is bounded below and above by a multiple of ω^{-k}.

Exponential Disjointness

The details of the modification of the proof of [Sti93c, Lemma 7.15], which says that *patches of the same generation only overlap if and only if they are neighbours*, will now be given. This yields an exponential bound for the distance between two patches which are separated by precisely one patch.

Lemma 97 (Exponential Disjointness)
Let P *and* Q *be patches in the same generation that are separated by exactly* one *patch. There exists a positive constant* d *and a positive constant* $\rho < 1$ *such that*

$$dist(P, Q) \geq d\rho^j. \tag{5.54}$$

Moreover, if P *and* Q *are separated by* n *patches of generation* j*, then*

$$\text{dist}(P, Q) \geq d\rho^{j-k}, \tag{5.55}$$

where k *is controlled by equation 5.53.*

Proof: (In order to understand the following, recall the details of the necklace construction in section 4.2. In particular, the succession and descendence relations given in definitions 69 and 70) There are three possible cases for neighbouring parent patches A and B (this follows from the definition of succession). Here the full proof will be given for the first case, the others are completely analogous.

In the following, $\mathcal{F}$ will indicate a Fibonacci tail, and $\mathcal{F}_j$ the j-th Fibonacci string.

Case 1: $A = \mathcal{F}\mathcal{B}^{-j}\mathcal{D}_\mathcal{T}$ and $B = \mathcal{F}\mathcal{B}^{-j}\mathcal{D}_\mathcal{U}$. In this case, A has a single child

$$A_1 = \mathcal{F}\mathcal{F}_j\mathcal{B}^{-(j+1)}\mathcal{D}_\mathcal{T}. \tag{5.56}$$

The progeny of B consists of the two neighbouring patches

$$B_1 = \mathcal{F}\mathcal{B}^{-(j+1)}\mathcal{D}_\mathcal{U}. \quad \text{and} \quad B_2 = \mathcal{F}\mathcal{B}^{-(j+1)}\mathcal{D}_\mathcal{T} \tag{5.57}$$

Further, these patches occur in the order A_1, B_1, B_2 so that $A_1 \cap B_2 = \emptyset$, i.e.

$$\mathcal{F}\mathcal{F}_j\mathcal{B}^{-(j+1)}\overline{\mathcal{D}_\mathcal{T}} \cap \mathcal{F}\mathcal{B}^{-(j+1)}\overline{\mathcal{D}_\mathcal{U}} = \emptyset, \tag{5.58}$$

which may be rewritten as

$$\mathcal{F}\mathcal{B}^{-(j-1)}\mathcal{B}^{-1}\mathcal{B}^j\mathcal{F}_j\mathcal{B}^{-j}\mathcal{B}^{-1}\mathcal{D}_\mathcal{T} \cap \mathcal{F}\mathcal{B}^{-(j-1)}\mathcal{B}^{-2}\mathcal{D}_\mathcal{U} = \emptyset. \tag{5.59}$$

We know that the map $\mathcal{F}\mathcal{B}^{-(j-1)}$ is invertible (section 4.5.9 demonstrated that it is a homeomorphism on any region where V is). Noticing that, by [Sti93c, Lemma 2.2], $\mathcal{B}^j\mathcal{F}_j\mathcal{B}^{-j} = \mathcal{T}$ (which follows by inverting the action of the symbol substitution, Σ, j-times until we reach $\mathcal{F}_0 = \mathcal{T}$), we obtain

$$\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\overline{\mathcal{D}_\mathcal{T}} \cap \mathcal{B}^{-2}\overline{\mathcal{D}_\mathcal{U}} = \emptyset, \tag{5.60}$$

which, by applying $\mathcal{B}^2$, is reduced to one of the disjointness conditions that was verified in chapter 4. It follows that

$$\text{dist}(\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\overline{\mathcal{D}_\mathcal{T}}, \mathcal{B}^{-2}\overline{\mathcal{D}_\mathcal{U}}) \geq \delta > 0. \tag{5.61}$$

Now, we are going to make use of coercivity. The map $\mathcal{F}\mathcal{B}^{-(j-1)}$ is obtained by composing the coercive maps $\mathcal{B}^{-1}$ and $\mathcal{T}\mathcal{B}^{-1}$ (restricted to $\mathcal{D}_\mathcal{T}$), since each tail of a patch of generation

$j-1$ corresponds to a "lineage string" of length $j-1$ in which each symbol represents one of these two maps (recall section 4.2.7). Let ρ be a common coercivity constant. Then ρ^{j-1} is a coercivity constant of $\mathcal{FB}^{-(j-1)}$, giving

$$\operatorname{dist}(\mathcal{FB}^{-(j-1)}\mathcal{B}^{-1}\mathcal{TB}^{-1}\overline{\mathcal{D}_{\mathcal{T}}}, \mathcal{FB}^{-(j-1)}\mathcal{B}^{-2}\overline{\mathcal{D}_{\mathcal{U}}}) \geq \delta\rho^{j-1} > 0. \tag{5.62}$$

The same strategy can be used for the other cases for A and B given in [Sti93c, Lemma 7.15] and 4.5.9. □

5.4.4 Hölder Continuity and Hölder Coercivity

We know (section 4.3.8) that the function $\phi(t)$ is Hölder continuous, i.e.

$$|\phi(s) - \phi(t)| \leq c_0|s-t|^{\mu}, \tag{5.63}$$

where the positive constant μ depends on the contractivity constants of the maps $\mathcal{B}^{-1}$ and $\mathcal{TB}^{-2}$, and c_0 is a positive constant.

Lemma 98 (Hölder Coercivity)
Using the above exponential bound on the distance between patches of a fixed generation, we can also prove that there exist positive constants d_0 *and* ν *such that*

$$|\phi(s) - \phi(t)| \geq d_0|s-t|^{\nu}. \tag{5.64}$$

This property of the curve $\phi(t)$ *will be called* Hölder coercivity.

Proof: Fix s and t in $[-\omega, 1-\omega]$ with $s \neq t$, and fix the generation index j so large that $\phi(s)$ and $\phi(t)$ are separated by one or more patches of generation j. Denote the number of patches separating them by n. We consider the effect of the necklace construction on the domains of parameterization $[-\omega, 1-\omega] = \phi^{-1}(\mathcal{N})$ and $[1-\omega, 1] = \phi^{-1}(\mathcal{M})$, which we might call the "flat necklace". Notice that all the "flat" patches of generation j have length either ω^j or ω^{j+1}. Furthermore, s and t are separated by n of these "flat" patches (see equation 5.53), so that there exist positive constants c_3, c_4 with

$$c_3\frac{|s-t|}{\omega^j} \leq n \leq c_4\frac{|s-t|}{\omega^j}. \tag{5.65}$$

Along with equation 5.53, this gives

$$\begin{aligned} \log c_1 + k\log\frac{1}{\omega} &\leq \log n \leq \log c_2 + k\log\frac{1}{\omega} \\ \log c_3 + j\log\frac{1}{\omega} + \log|s-t| &\leq \log n \leq \log c_4 + j\log\frac{1}{\omega} + \log|s-t|. \end{aligned}$$

Since $\phi(s)$ and $\phi(t)$ are separated by n actual patches of generation j, we have the estimate (see equation 5.55)

$$|\phi(s) - \phi(t)| \geq d\rho^{j-k}. \tag{5.66}$$

An upper bound for $j-k$ is then given by

$$\begin{aligned} j-k &\leq j - \frac{-\log c_2 + \log n}{\log \frac{1}{\omega}} \\ &\leq j - \frac{-\log c_2 + \log c_3 + j\log\frac{1}{\omega} + \log|s-t|}{\log\frac{1}{\omega}} \\ &= \frac{-\log c_2 + \log c_3}{\log\frac{1}{\omega}} - \frac{\log|s-t|}{\log\frac{1}{\omega}} \\ &=: c_5 - \frac{\log|s-t|}{\log\frac{1}{\omega}}, \end{aligned}$$

(with the obvious definition of c_5). Combining these gives

$$|\phi(s)-\phi(t)| \geq d\rho^{j-k} \geq d\rho^{c_5 - \frac{\log|s-t|}{\log\frac{1}{\omega}}} \geq d_0|s-t|^{\nu}, \tag{5.67}$$

where d_0 is positive and $\nu = \frac{\log\rho}{\log\omega}$. Since s and t were arbitrary, this establishes Hölder coercivity for the curve ϕ. □

5.4.5 Lower Bound for Hausdorff Dimension

It turns out that *Hölder coercivity* immediately yields a lower bound for the Hausdorff dimension, using the so-called *mass distribution principle* of [Fal90, §4.2]:

Theorem 99 (Mass Distribution Principle)
Let m be a mass distribution on a set A, and suppose that for some σ there exist $c > 0$ and $\delta > 0$ such that

$$m(U) \leq c|U|^{\sigma}, \tag{5.68}$$

for all sets U with $|U| \leq \delta$. Then $\mathcal{H}^{\sigma}(A) \geq m(A)/c$ and

$$\sigma \leq \dim_H A, \tag{5.69}$$

where $\mathcal{H}^{\sigma}$ is the σ-dimensional Hausdorff measure, and $\dim_H$ is the Hausdorff dimension.

In order to apply this result here, we look at the *invariant measure* on the invariant curve. Let a and b be elements of $[-\omega, 1]$ (the full interval of parameterization in the necklace construction) and assume that $a < b$. For any points s and t of $[a,b]$, Hölder coercivity gives

$$|\phi(s)-\phi(t)| \geq d_0|s-t|^{\nu}. \tag{5.70}$$

In other words, the diameter of the image of $[a,b]$ under ϕ admits the lower bound

$$|\phi([a,b])| \geq d_0|s-t|^{\nu}. \tag{5.71}$$

Since this is valid for all s and t in $[a,b]$, we may choose $s=a$ and $t=b$ to give

$$|\phi([a,b])| \geq d_0|a-b|^{\nu}. \tag{5.72}$$

Dynamically Invariant Measure: The distribution of points on the necklace set under the dynamics of the maps of the IFS defines an invariant measure. Note that the necklace set is identified with the interval $[a, b]$ by means of the bijection ϕ (in fact, $\mathcal{L} = \phi([a, b])$). It follows that the measure of a subset of the necklace set is equal (modulo normalization) to the one-dimensional Lebesgue measure of its pre-image with respect to ϕ, inducing a measure on the interval.

The resulting measure is dynamically invariant, because the one-dimensional Lebesgue measure is invariant with respect to the translations $t \mapsto t - 1$ and $t \mapsto t + \omega$. (Of course, any constant multiple of the measure will also be invariant.)

Let m denote the dynamically invariant measure on the necklace set $\mathcal{L}$ ($\mathcal{L}$ is the *support* of m). The interval $[a, b]$ has Lebesgue measure $|b - a|$. Modulo normalization, it follows that it is equal to $m(\phi([a, b]))$, giving

$$|\phi([a, b])| \geq d_0 m(\phi([a, b]))^\nu. \tag{5.73}$$

Let U be any set which intersects the necklace curve. Let $[a, b]$ be the smallest interval such that the intersection of U with the necklace curve is contained in $\phi([a, b])$. (Put $a = \inf\{ t : \phi(t) \in U \}$ and $b = \sup\{ t : \phi(t) \in U \}$.) Notice that

$$m(U) = m(U \cap \mathcal{L}) \leq m(\phi([a, b])), \tag{5.74}$$

Moreover, by the minimality of the interval $[a, b]$,

$$|U| \geq |\phi([a, b])|. \tag{5.75}$$

This gives the estimate

$$|U| \geq |\phi([a, b])| \geq d_0 m(\phi([a, b]))^\nu \geq d_0 m(U)^\nu, \tag{5.76}$$

which implies that

$$m(U) \leq \frac{1}{d_0^{\frac{1}{\nu}}} |U|^{\frac{1}{\nu}}. \tag{5.77}$$

Finally, by the mass distribution principle (above), this gives that the number $\frac{1}{\nu}$ is a lower bound on the Hausdorff dimension of $\mathcal{L}$.

Unfortunately, the numerical values for ν that we find are much greater than 1 (recall that $\nu = \log \rho / \log \omega$ where ν is defined in lemma 94, and compare with the bounds found for the derivative of $\mathcal{TB}^{-2} \equiv \mathrm{QVB}^{-2}$ when verifying the contractivity hypothesis in section 4.5.8). Therefore, the above lower bound will be smaller than 1, which makes it useless. (The necklace set is a curve, and this immediately implies that 1 is a lower bound anyway.) However, the above reasoning is included here for completeness. The later sections of this chapter attempt to get useful bounds by rigorous numerics.

5.4.6 Upper Bound for Hausdorff Dimension

It is simpler to show that *Hölder continuity* yields an upper bound for the Hausdorff dimension of the necklace curve. (According to [Fal90, §2.3] this is true in general.)

Let a, b, s, and t be as above. Hölder continuity implies

$$|\phi(s) - \phi(t)| \leq c|s - t|^{\mu}. \tag{5.78}$$

This gives

$$|\phi(s) - \phi(t)| \leq c|a - b|^{\mu} = cm(\phi([a, b]))^{\mu}. \tag{5.79}$$

By the same argument as above, this implies that

$$|\phi([a, b])| \leq cm(\phi([a, b]))^{\mu}. \tag{5.80}$$

Consider any *partition* of the necklace curve by sets U of the form $\phi([a, b))$ (the images of half-open intervals). It follows that

$$\sum |U|^{\frac{1}{\mu}} \leq c^{\frac{1}{\mu}} \sum m(U) = \frac{1}{c^{\frac{1}{\mu}}} m(\mathcal{L}) < \infty. \tag{5.81}$$

Since this holds for arbitrary partitions, $\frac{1}{\mu}$ is an upper bound for the Hausdorff dimension of $\mathcal{L}$.

Numerically, the upper bound thus found is about 1.6. The following sections develop the theory to enable the calculation of rigorous bounds by computer, which will be used to produce sharper bounds.

5.5 Partition Function Approach to Hausdorff Dimension

In this section, the necessary theory for bounding the Hausdorff dimension will be developed. (For the following, compare Falconer [Fal90, Chapter 3]. In particular 9.2, Theorem 9.3, and Example 9.8.)

5.5.1 Binary Patches

Consider the $\mathcal{N}$ part of the necklace set (recall that $\mathcal{L} = \mathcal{M} \cup \mathcal{N}$, where $\mathcal{L}$ is the necklace set), together with the maps $\mathcal{B}^{-1}$ and $\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$. For simplicity, we denote $\mathcal{B}^{-1}$ by ψ_L and $\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$ by ψ_R. Notice that ψ_R is an *analytic* map and ψ_L is an exact *anti-similarity*. Both maps are globally *contractive*; let $\kappa < 1$ denote a common contractivity constant. (They are also also *homeomorphisms*, at least on a suitable domain.)

A *binary patch* is the set obtained by applying an arbitrary composition of the maps ψ_L and ψ_R to the set $\mathcal{N}$:

$$P = \psi_j \psi_{j-1} \cdots \psi_1(\mathcal{N}),$$

where each of the ψ_i is either ψ_L or ψ_R. The number j is called the *generation index* of the binary patch P. We denote the mapping defining P by Ψ_P:

$$\Psi_P = \psi_j \cdots \psi_1.$$

Since a binary patch is entirely determined by this mapping, we may identify P with Ψ_P.

Definition 100 (Product of Binary Patches)
The product $P*Q$ *of two binary patches is the patch identified by the composition* $\Psi_P \circ \Psi_Q$.

(Which means that $j(P*Q) = j(P) + j(Q)$.)

The union of all the binary patches with a fixed generation index j will be called the j-*th generation* and will be denoted by G_j.

Definition 101 (Product of Generations)
We defined the product of two generations by

$$G_j * G_k = \{ P*Q : P \in G_j, Q \in G_k \}.$$

Obviously, we have

$$G_{j+1} = G_j * G_1, \tag{5.82}$$

and

$$G_{2j} = G_j * G_j. \tag{5.83}$$

5.5.2 Difference Quotients

We are interested in the coercivity and contractivity of the maps Ψ_P. To that end, we look at *difference quotients*

$$\Delta_P(x,y) = \frac{\Psi_P(x) - \Psi_P(y)}{x-y},$$

where x and y range over $\mathcal{N}$ and $x \neq y$. Obviously, an upper bound of this difference quotient is a *Lipschitz constant* for Ψ_P (on $\mathcal{N}$), whereas a lower bound is a *coercivity constant* (see section 5.4.1).

Because the maps ψ_L and ψ_R are globally contractive, all these Lipschitz constants are strictly smaller than 1 and are therefore *contractivity* constants. Denote the infimum of $|\Delta_P(x,y)|$ by $\rho(P)$ and the supremum by $\kappa(P)$:

$$\rho(P) := \inf_{x,y \in \mathcal{N}, x \neq y} |\Delta_P(x,y)|, \tag{5.84a}$$

$$\kappa(P) := \sup_{x,y \in \mathcal{N}, x \neq y} |\Delta_P(x,y)|. \tag{5.84b}$$

(These numbers constitute the *optimal* contractivity and coercivity constants.) Notice that the *diameter* $|P|$ of P admits the estimates

$$\rho(P) \cdot |\mathcal{N}| \leq |P| \leq \kappa(P) \cdot |\mathcal{N}|, \tag{5.85}$$

(recall that ρ and κ are lower, respectively upper, bounds on difference quotients). Further, it follows that

$$\begin{aligned} \Delta_{P*Q}(x,y) &= \frac{\Psi_P(\Psi_Q(x)) - \Psi_P(\Psi_Q(y))}{x-y} \\ &= \frac{\Psi_P(\Psi_Q(x)) - \Psi_P(\Psi_Q(y))}{\Psi_Q(x) - \Psi_Q(y)} \cdot \frac{\Psi_Q(x)) - \Psi_Q(y)}{x-y}. \end{aligned} \tag{5.86}$$

(Here, the right hand side is defined whenever x is different from y.) Therefore

$$\rho(P * Q) \geq \rho(P)\rho(Q), \tag{5.87a}$$
$$\kappa(P * Q) \leq \kappa(P)\kappa(Q). \tag{5.87b}$$

In view of this, it might be helpful to call ρ "super-multiplicative" and κ "sub-multiplicative".

5.5.3 Partition Functions

Consider the lower and upper *partition functions* $Z_l^{(j)}(s)$ and $Z_u^{(j)}(s)$ defined by

$$Z_l^{(j)}(s) := \sum_{P \in G_j} \rho(P)^s, \tag{5.88a}$$
$$Z_u^{(j)}(s) := \sum_{P \in G_j} \kappa(P)^s. \tag{5.88b}$$

(Notice that both functions are strictly decreasing in s.)

Equation 5.87a ("super-multiplicativity") implies that

$$\begin{aligned} Z_l^{(2j)}(s) &= \sum_{P \in G_{2j}} \rho(P)^s = \sum_{Q,R \in G_j} \rho(Q * R)^s \\ &\geq \sum_{Q,R \in G_j} \rho(Q)^s \rho(R)^s = (Z_l^{(j)}(s))^2. \end{aligned} \tag{5.89}$$

In the same way, it follows ("sub-multiplicativity", equation 5.87b) that

$$Z_u^{(2j)}(s) \leq (Z_u^{(j)}(s))^2. \tag{5.90}$$

By analogy with Falconer [Fal90, §9], we now define real numbers r_j and t_j which solve equations where the lower, respectively upper, partition functions are equated to 1:

$$Z_l^{(j)}(r_j) = \sum_{P \in G_j} \rho(P)^{r_j} = 1, \tag{5.91a}$$
$$Z_u^{(j)}(t_j) = \sum_{P \in G_j} \kappa(P)^{t_j} = 1. \tag{5.91b}$$

The equations (5.89) and (5.90), together with the fact that the partition functions are strictly decreasing in s, implies that

$$r_j \leq r_{2j} \leq t_{2j} \leq t_j. \tag{5.92}$$

If, in addition, it turns out that $|r_j - t_j| \to 0$ as $j \to \infty$, it follows that the sequences r_j and t_j converge to a common limit. [Fal90, Theorem 9.3] would then suggest that this limit is the Hausdorff dimension of $\mathcal{N}$.

The next section proves that this is, in fact, the case. $r_j - t_j$ is shown to be a zero sequence by using the *bounded variation principle*. That the limit is equal to the Hausdorff dimension follows by modifying the proof of [Fal90, Theorem 9.3]. (In order to be able to apply this theorem, it is necessary to verify that the *open set condition* holds for any IFS that is used.) The details are given below, after which it remains to verify the open set condition for the IFSs.

5.5.4 The Bounded Variation Principle

Definition 102 (Bounded Variation)
The bounded variation principle, *to be established below, says that* the quotient of coercivity and contractivity remains uniformly bounded away from zero.

This fact will then be used to prove that the lower and upper bounds, $\mathfrak{r}$ and $\mathfrak{t}$ (obtained by solving the partition equations) must converge as the generation index tends to infinity, in the case where the contractivity and coercivity constants are *optimal*. (In general, the bounds we obtain using rigorous numerics and the mean value theorem to bound certain derivatives will not actually be optimal.) Then, by establishing the validity of the *open set condition*, it will follow that the common limit of the lower and upper bounds is the Hausdorff dimension of the necklace. (Finally, by proving that the dimension is invariant under renormalization, for maps that are close enough to the fixed point and attracted, the result will be extended to the dimension of the invariant curve for such maps.)

Estimating Difference Quotients

Consider the map $\psi_R = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$. This map is *analytic* on an open neighbourhood Ω of the set $\mathcal{N}$. Let Γ be a contour surrounding $\mathcal{N}$ in Ω, and let U be the interior of Γ. The following estimates then follow:

Lemma 103

1. *There exists a function $\gamma(z,w)$ analytic on $U \times U$ such that*

$$\frac{\psi_R(z) - \psi_R(w)}{z - w} - \psi_R{}'(z) = (z - w)\gamma(z,w), \tag{5.93}$$

   i.e. this function estimates the difference between the difference quotient and the derivative.

2. *There exists a function $\delta(z,w)$ analytic on $U \times U$ such that*

$$\psi_R{}'(z) - \psi_R{}'(w) = (z - w)\delta(z,w), \tag{5.94}$$

   i.e. this function estimates the difference between derivatives at different points.

Proof: By Cauchy's theorem [Rud80], we have

$$\begin{aligned} \psi_R(z) &= \frac{1}{2\pi i}\oint_\Gamma \frac{\psi_R(\zeta)\,d\zeta}{\zeta - z} \\ \psi_R(w) &= \frac{1}{2\pi i}\oint_\Gamma \frac{\psi_R(\zeta)\,d\zeta}{\zeta - w}, \end{aligned}$$

for any two points z and w of U. It follows that

$$\frac{\psi_R(z) - \psi_R(w)}{z - w} = \frac{1}{2\pi i}\oint_\Gamma \frac{\psi_R(\zeta)\,d\zeta}{(\zeta - z)(\zeta - w)}.$$

Moreover,

$$\psi_R'(z) = \frac{1}{2\pi i}\oint_\Gamma \frac{\psi_R(\zeta)\,d\zeta}{(\zeta - z)^2}.$$

Subtraction yields

$$\frac{\psi_R(z) - \psi_R(w)}{z - w} - \psi_R'(z) = \frac{1}{2\pi i}\oint_\Gamma \frac{\psi_R(\zeta)\,d\zeta}{(\zeta - z)^2(\zeta - w)} \cdot (w - z),$$

which verifies the first claim. Similarly, the equation

$$\psi_R'(z) - \psi_R'(w) = \frac{1}{2\pi i}\oint_\Gamma \frac{\psi_R'(\zeta)d\zeta}{(\zeta - z)(\zeta - w)} \cdot (z - w)$$

proves the second claim. □

Quotients of Difference Quotients In order to discuss quotients of coercivity and contractivity, it is necessary to estimate the quotient of two difference quotients. Assume that $x, y, x', y' \in \mathcal{N}$, with $x \neq y$, and $x' \neq y'$. Then

$$\begin{aligned} Q_{\psi_R}(x, y, x', y') &= \frac{\psi_R(x) - \psi_R(y)}{x - y} \Big/ \frac{\psi_R(x') - \psi_R(y')}{x' - y'} \\ &= \frac{\psi_R'(x) + (x - y)\gamma(x, y)}{\psi_R'(x') + (x' - y')\gamma(x', y')} \\ &= \frac{\psi_R'(x)}{\psi_R'(x')} \left(\frac{1 + (x - y)\frac{\gamma(x,y)}{\psi_R'(x)}}{1 + (x' - y')\frac{\gamma(x',y')}{\psi_R'(x')}} \right) \\ &= \left(1 + (x - x')\frac{\delta(x, x')}{\psi_R'(x')}\right) \left(\frac{1 + (x - y)\frac{\gamma(x,y)}{\psi_R'(x)}}{1 + (x' - y')\frac{\gamma(x',y')}{\psi_R'(x')}} \right). \end{aligned}$$

(ψ_R', i.e., $(\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1})'$, does not vanish on $\mathcal{N}$. So the right hand sides are well defined.) Since $\mathcal{N}$ is compact (by the necklace construction), there exists a constant $c_0 > 0$ such that

$$|Q_{\psi_R}(x, y, x', y')| \leq 1 + c_0 \max\{|x - y|, |x' - y'|, |x - x'|\}.$$

In particular, if the points x, y, x', and y' are confined to a subset of $\mathcal{N}$ of diameter d, the absolute value of Q_{ψ_R} is bounded by $1 + c_0 d$. Taking the logarithm yields the estimate

$$\log |Q_{\psi_R}| \leq c_0 d, \tag{5.95}$$

by the inequality $\log(1 + t) \leq t$ valid for all $t > -1$.

For the map $\psi_L = \mathcal{B}^{-1}$, which is an anti-similarity, the difference quotients have constant absolute value. Therefore $Q_{\psi_L} = 1$, and the bound $\log |Q_{\psi_L}| \leq c_0 d$ follows immediately.

Bounded Variation for Difference Quotients

Having done these preparations, we now come to the heart of the problem: the derivation of the bounded variation principle for difference quotients.

Note that, for any binary patch P and for any two different points x and y in $\mathcal{N}$, we have

$$\begin{aligned}\frac{\Psi_P(x)-\Psi_P(y)}{x-y} &= \frac{x_j-y_j}{x_{j-1}-y_{j-1}}\cdot\cdots\cdot\frac{x_1-y_1}{x_0-y_0}\\ &= \frac{\psi_j(x_{j-1})-\psi_j(y_{j-1})}{x_{j-1}-y_{j-1}}\cdot\cdots\cdot\frac{\psi_1(x_0)-\psi_1(y_0)}{x_0-y_0},\end{aligned}$$

where

$$\begin{aligned} x_0 &= x, & y_0 &= y,\\ x_i &= \psi_i(x_{i-1}), & y_i &= \psi_i(y_{i-1}), \quad i\geq 1.\end{aligned}$$

Here, each ψ_i is either $\psi_L=\mathcal{B}^{-1}$ or $\psi_R=\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$. It follows that

$$\begin{aligned}\frac{\Delta_P(x,y)}{\Delta_P(x',y')} &= \frac{\Psi_P(x)-\Psi_P(y)}{x-y}\Big/\frac{\Psi_P(x')-\Psi_P(y')}{x'-y'}\\ &= Q_{\psi_j}(x_{j-1},y_{j-1},x'_{j-1},y'_{j-1})\cdot\cdots\cdot Q_{\psi_1}(x_0,y_0,x'_0,y'_0).\end{aligned}$$

Observe that, for any i, the points x_i, y_i, x'_i, and y'_i are confined to one and the same binary patch of generation i (the sequence $\{\psi_i\}$ being the same for each of them). Recall that the diameter of such a patch is bounded by $c\kappa^i$, where κ is the global decay rate defined in lemma 76. Taking the logarithm of the absolute value of the last equation and using equation 5.95 yields

$$|\log\Delta_P(x,y)-\log\Delta_P(x',y')|\leq c_1\kappa^{j-1}+c_1\kappa^{j-2}+\cdots+c_1\kappa^0,$$

where $c_1=c_0c$. The geometric series on the right hand side is bounded uniformly in j. Taking the exponential of both sides yields the estimate

$$\frac{1}{c_3}\leq\left|\frac{\Delta_P(x,y)}{\Delta_P(x',y')}\right|\leq c_3,$$

where $c_3>1$ is a constant. Since this holds for all pairs (x,y) and (x',y') in $P\times P$, it continues to hold with the supremum of the numerator and the infimum of the denominator (*optimality*), giving

$$\frac{1}{c_3}\leq\frac{\kappa(P)}{\rho(P)}\leq c_3.$$

Since $\rho(P)\leq\kappa(P)$, we can restate this by saying that there exists a positive constant $c>1$ such that

$$\frac{1}{c}\kappa(P)\leq\rho(P)\leq\kappa(P). \tag{5.96}$$

This is the principle of bounded variation: *the quotient of coercivity and contractivity remains uniformly bounded away from zero.* (Notice that the bound depends neither on the patch P, nor on its generation.)

5.5.5 Convergence of Lower and Upper Bounds

Now, it is needed to prove that the differences $|r_j - t_j|$ converge to zero as $j = 2^i \to \infty$.

In order to simplify notation, fix $j = 2^i$ and put $r = r_j$ and $t = t_j$. Equation 5.92 gives $r \leq t$, i.e. $t = r + \delta$, where $\delta \geq 0$. By definition (equation 5.91), we have

$$\begin{aligned} \sum_P \rho(P)^r &= 1, \\ \sum_P \kappa(P)^t &= 1, \end{aligned}$$

where the sums are taken over all patches of generation j. From equation 5.96 (bounded variation) it follows that

$$\begin{aligned} 1 &= \sum_P \kappa(P)^t \leq c^t \sum_P \rho(P)^t \\ &= c^t \sum_P \rho(P)^{r+\delta} \\ &\leq c^t (\rho_{\max})^\delta \sum_P \rho(P)^r \\ &= c^t (\rho_{\max})^\delta. \end{aligned}$$

Observe that, since the sequence t_{2^i} is decreasing, the quantities $c^t = c^{t_{2^i}}$ are bounded above as $i \to \infty$. Moreover, $\rho_{\max} \to 0$ as $i \to \infty$ (see equation 5.85). Therefore, the product $c^t(\rho_{\max})^\delta$ can only stay bounded away from zero if $\delta \to 0$ as $i \to \infty$, i.e. if $r \to t$. □

5.6 Open Set Condition

The previous sections established that solutions to the upper partition-function equations give upper bounds on the dimension. In order that solutions to the *lower* partition-function equations give lower bounds on the dimension, it is necessary to verify the open set condition. This will be done first for the accretive composition order ($\mathcal{UT}$). It is simpler to construct a proof for the Fibonacci necklace construction, and then to adapt the results for the binary IFS. For the inverse composition order, we were only able to prove the open set condition for the Fibonacci system, we could not see how to carry it across to the corresponding binary IFS. The proof is included both for completeness, and in the hope that the open set condition might in future be verified for the inverse order.

5.6.1 Accretive Fibonacci Necklace

For the open set condition, recall that we need to find an open set that contains its images under the transformations of the IFS, and for which the union of the images is disjoint. For the Fibonacci system, we will construct this set from two pieces, giving:

Proposition 104
There exist non-empty disjoint *open sets* $\Omega_N \subset \mathcal{D}_\mathcal{T}$ *and* $\Omega_M \subset \mathcal{D}_\mathcal{U}$ *such that*

$$\begin{aligned} \mathcal{B}^{-1}\Omega_M &\subset \Omega_N, && (5.97a)\\ \mathcal{B}^{-1}\Omega_N &\subset \Omega_N, && (5.97b)\\ \mathcal{T}\mathcal{B}^{-1}\Omega_N &\subset \Omega_M. && (5.97c) \end{aligned}$$

Remark 105 (Forward Invariance and Full Invariance)
This means that the pair (Ω_M, Ω_N) *is* forward *invariant for the two maps* $\mathcal{B}^{-1}$ *and* $\mathcal{T}\mathcal{B}^{-1}$. *By contrast, the pair* $(\mathcal{M}, \mathcal{N})$ *is* fully *invariant:*

$$\begin{aligned} \mathcal{B}^{-1}(\mathcal{M} \cup \mathcal{N}) &= \mathcal{N}, \\ \mathcal{T}\mathcal{B}^{-1}\mathcal{N} &= \mathcal{M}. \end{aligned}$$

Outline of strategy. The idea is as follows: recall that the necklace is constructed from two transformed copies of itself, which intersect at the pivotal point. (The pivotal point itself is an image of the boundary points, and vice versa.) This violates the open set condition, which requires a disjoint union of the pieces. However, removing the boundary points and the pivotal point from the necklace curve, while destroying full invariance, leaves the remaining set *forward* invariant. In order to see this, notice that that

$$\begin{aligned} \mathcal{B}^{-1}\phi(1) &= \phi(-\omega), && (5.98a)\\ \mathcal{B}^{-1}\phi(-\omega) &= \phi(1-\omega), && (5.98b)\\ \mathcal{T}\mathcal{B}^{-1}\phi(1-\omega) &= \phi(1-\omega), && (5.98c)\\ \mathcal{T}\mathcal{B}^{-1}\phi(-\omega) &= \phi(1). && (5.98d) \end{aligned}$$

Let S denote the set $\{\phi(-\omega), \phi(1-\omega), \phi(1)\}$. Since the maps

$$\begin{aligned} \mathcal{B}^{-1} &: \mathcal{M} \cup \mathcal{N} \to \mathcal{N}, \\ \mathcal{T}\mathcal{B}^{-1} &: \mathcal{N} \to \mathcal{M}, \end{aligned}$$

are bijective, the equations (5.98a–5.98d) show that the *pre-image* of S with respect to both $\mathcal{B}^{-1}$ and $\mathcal{T}\mathcal{B}^{-1}$ is equal to S. In other words, no point lying outside S is mapped into S. Therefore, the complement of S in the pair $(\mathcal{M}, \mathcal{N})$ is forward invariant under $\mathcal{B}^{-1}$ and $\mathcal{T}\mathcal{B}^{-1}$. More precisely,

$$\begin{aligned} \mathcal{B}^{-1}(\mathcal{M} \setminus S) &\subset \mathcal{N} \setminus S, && (5.99a)\\ \mathcal{B}^{-1}(\mathcal{N} \setminus S) &\subset \mathcal{N} \setminus S, && (5.99b)\\ \mathcal{T}\mathcal{B}^{-1}(\mathcal{N} \setminus S) &\subset \mathcal{M} \setminus S. && (5.99c) \end{aligned}$$

Notice, incidentally, that S is *not* forward invariant, since

$$\mathcal{B}^{-1}\phi(1-\omega) = \phi(-\omega^3) \notin S.$$

(This implies that the pair $((\mathcal{M} \setminus S), (\mathcal{N} \setminus S))$ is not *backward* invariant.)

Modifying the necklace construction

The above observations suggest that suitable sets can be constructed (which satisfy the open set condition) by taking each generation of the necklace construction and removing the boundary patches and the two patches which contain the pivotal point. This turns each generation into a union of two disjoint open sets, one contained in $\mathcal{D}_{\mathcal{T}}$ and the other in $\mathcal{D}_{\mathcal{U}}$.

Definition 106 (Disjoint Sets Ω_M and Ω_N.)
For each generation, with the boundary and pivotal patches removed, let Ω_N be the union of the open sets contained in $\mathcal{D}_{\mathcal{T}}$, and Ω_M the union of the remaining ones.

The sets Ω_N and Ω_M thus defined satisfy the invariance conditions (5.97a–5.97c).

Before the proof of this is given, the notion of a *pivotal patch* needs to be re-defined.

Definition 107 (Pivotal Patches)
The pivotal patches are defined to be the patches containing the pivotal point $\phi(1-\omega)$.

It follows that there are *at most two* pivotal patches in any generation, because the intersection of three different patches of the same generation is always empty. In fact, since $\phi(1-\omega) \in \mathcal{M} \cap \mathcal{N}$, the pivotal point must be contained simultaneously in a patch of $\mathcal{N}_j$ and in a patch of $\mathcal{M}_j$, in each generation, which means that there are precisely *two* pivotal patches in each generation.

Remark 108
Notice that this definition is different from the one used in chapter 4, where only one *patch was called "pivotal", namely the one in the natural ordering after which the next symbol ν changes. This is the patch tagged by $\mathcal{T}$ whose Fibonacci tail $\mathcal{F}$ satisfies*

$$F_j = \begin{cases} \mathcal{T}\mathcal{U}\mathcal{F} & \text{for even } j \\ \mathcal{U}\mathcal{T}\mathcal{F} & \text{otherwise.} \end{cases}$$

According to the definition just given, both this patch and its successor are called pivotal. Notice, by the definition of succession (definition 69), that this successor is of the form $(\mathcal{F} - \mathcal{F}_{j-2})\mathcal{B}^{-j}\mathcal{D}_{\mathcal{U}}$ (if $j \geq 2$). So, for $j \geq 2$, one of the two pivotal patches is tagged by $\mathcal{T}$ and the other by $\mathcal{U}$.

Notation: In each generation j, let the left (resp. right) boundary patch be denoted by L^j (resp. R^j), and the left (resp. right) pivotal patch be denoted by P^j (resp. Q^j).

In order to prove forward invariance of (Ω_N, Ω_M), we shall first establish the following Lemma, which is designed with equations 5.98a–5.98d in mind.

Lemma 109
The following equations hold:

$$\mathcal{B}^{-1}\mathsf{R}^j = \mathsf{L}^{j+1}, \tag{5.100a}$$
$$\mathcal{B}^{-1}\mathsf{L}^j = \mathsf{P}^{j+1}, \tag{5.100b}$$
$$\mathcal{T}\mathcal{B}^{-1}\mathsf{P}^j = \mathsf{Q}^{j+1}, \tag{5.100c}$$
$$\mathcal{T}\mathcal{B}^{-1}\mathsf{L}^j = \mathsf{R}^{j+1}. \tag{5.100d}$$

Proof: To prove equation (5.100b), notice that the patch L^j contains the point $\phi(-\omega)$ (being the left-hand edge of the interval of parameterization for the necklace set). Therefore (by equation 5.98) the patch $\mathcal{B}^{-1}\mathsf{L}^j$ contains the point $\phi(\omega^2) = \phi(1-\omega)$, i.e., the pivotal point. It follows that it is one of the pivotal patches. Since it is in $\mathcal{N}_{j+1}$, it must be the left pivotal patch P^{j+1}.

To prove equation (5.100c), notice that the patch P^j contains the pivotal point $\phi(1-\omega)$. Therefore, the patch $\mathcal{T}\mathcal{B}^{-1}\mathsf{P}^j$ contains the point

$$\phi(\omega - \omega(1-\omega)) = \phi(\omega\cdot(1-(1-\omega))) = \phi(\omega^2) = \phi(1-\omega),$$

i.e., the pivotal point. Since it is in $\mathcal{M}_{j+1}$, it must be the right pivotal patch Q^{j+1}.

To prove the remaining equations, recall the definition of Σ (section 4.2.3) and of the *next symbol* ν, and let $\hat{\mathcal{F}}_j$ denote the maximal length Fibonacci tail in $\{\mathcal{F}_j\}$. It follows that

$$\begin{array}{llllllll} \text{If} & \nu(\hat{\mathcal{F}}_j) & = & \mathcal{U} & \text{then} & \Sigma(\hat{\mathcal{F}}_j) & = & \hat{\mathcal{F}}_{j+1}, \text{ and} \\ \text{if} & \nu(\hat{\mathsf{F}}_j) & = & \mathcal{T} & \text{then} & \mathcal{T}\Sigma(\hat{\mathcal{F}}_j) & = & \hat{\mathcal{F}}_{j+1}. \end{array}$$

To see this notice that, by definition, the maximal length Fibonacci tail of the string $\mathcal{F}_j$ is simply $\mathcal{F}_j$ with its leading symbol removed, i.e. the next symbol of $\hat{\mathcal{F}}_j$ is the leading symbol of $\mathcal{F}_j$. Notice, also, that the leading symbol of $\mathcal{F}_j$ is $\mathcal{U}$ if j is odd and $\mathcal{T}$ if j is even. Therefore $\nu(\hat{\mathcal{F}}_j) = \mathcal{U}$ means that $\mathcal{F}_j = \mathcal{U}\hat{\mathcal{F}}_j$, giving

$$\mathcal{F}_{j+1} = \Sigma(\mathcal{F}_j) = \Sigma(\mathcal{U}\hat{\mathcal{F}}_j) = \mathcal{T}\Sigma(\hat{\mathcal{F}}_j),$$

i.e. $\hat{\mathcal{F}}_{j+1} = \Sigma(\hat{\mathcal{F}}_j)$. The case of even j is completely analogous, proving the second implication stated above.

Now, recall that $\Sigma\mathcal{F} = \mathcal{B}^{-1}\mathcal{F}\mathcal{B}$ for any Fibonacci tail $\mathcal{F}$. Suppose that j is *odd* then, from the definition of succession (definition 69) the boundary patches are given by,

$$\begin{array}{llllll} \mathsf{L}^j & = & \hat{\mathcal{F}}_{j-1}\mathcal{B}^{-j}\mathcal{D}_{\mathcal{U}}, & \mathsf{R}^j & = & \hat{\mathcal{F}}_j\mathcal{B}^{-j}\mathcal{D}_{\mathcal{T}}, \\ \mathsf{L}^{j+1} & = & \hat{\mathcal{F}}_{j+1}\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{T}}, & \mathsf{R}^{j+1} & = & \hat{\mathcal{F}}_j\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{U}}. \end{array}$$

Noticing that $\nu(\hat{\mathcal{F}}_j) = \mathcal{U}$ (by the reasoning presented above), we obtain

$$\begin{aligned} \mathcal{B}^{-1}\mathsf{R}^j &= \mathcal{B}^{-1}\hat{\mathcal{F}}_j\mathcal{B}\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{T}} \\ &= \Sigma(\hat{\mathcal{F}}_j)\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{T}} \\ &= \hat{\mathcal{F}}_{j+1}\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{T}} \\ &= \mathsf{L}^{j+1}, \end{aligned}$$

proving equation 5.100a. Further, noticing that $\nu(\mathcal{F}_{j-1}) = \mathcal{T}$ (also above), we obtain

$$\begin{aligned}\mathcal{T}\mathcal{B}^{-1}\mathsf{L}^j &= \mathcal{T}\mathcal{B}^{-1}\hat{\mathcal{F}}_{j-1}\mathcal{B}\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{U}} \\ &= \mathcal{T}\Sigma(\hat{\mathcal{F}}_{j-1})\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{U}} \\ &= \hat{\mathcal{F}}_j\mathcal{B}^{-(j+1)}\mathcal{D}_{\mathcal{T}} \\ &= \mathsf{R}^{j+1},\end{aligned}$$

proving equation 5.100d.

(The case of even j is analogous.) □

Proof of Proposition 104: In view of the above argument, invariance is readily deduced: Let z be any point of Ω_M. There exists an index j such that z is contained in a patch of $\mathcal{M}_j$ *different* from the right boundary patch. It follows that $\mathcal{B}^{-1}(z)$ is contained in a patch of $\mathcal{N}_{j+1}$ which is different from both the left boundary patch and from the left pivotal patch. (The preceding lemma tell us that the left boundary patch of generation $j+1$ is the image under $\mathcal{B}^{-1}$ of precisely the right boundary patch of generation j, and that the left pivotal patch of generation $j+1$ is the image under $\mathcal{B}^{-1}$ of precisely the left boundary patch of generation j.) Therefore, $\mathcal{B}^{-1}(z) \in \Omega_N$, proving equation 5.97a. The remaining equations (5.97b and 5.97c) follow in the same way.

To complete the proof of the open set condition for the accretive Fibonacci necklace, it remains to prove that the limit sets Ω_M and Ω_N are disjoint. This is demonstrated below (unfortunately, it is not very illuminating, but is included for completeness).

To this end, consider a recursively defined *increasing* sequence of pairs of open sets (Ω_M^j, Ω_N^j) such that

$$\Omega_M = \bigcup_j \Omega_M^j, \qquad \Omega_N = \bigcup_j \Omega_N^j.$$

Start with $j = 2$. Let us write the second generation as an ordered collection of patches:

$$\mathcal{U}\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}} < \mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}} < \mathcal{B}^{-2}\mathcal{D}_{\mathcal{U}} < \mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}} < \mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{U}}.$$

Here, $\mathcal{U}\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}}$ is the left boundary patch, $\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{U}}$ is the right boundary patch, and $\mathcal{B}^{-2}\mathcal{D}_{\mathcal{U}}$ and $\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}}$ are the pivotal patches. Summing up:

$$\mathsf{L}^2 = \mathcal{U}\mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}}, \quad \mathsf{R}^2 = \mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{U}}, \quad \mathsf{P}^2 = \mathcal{B}^{-2}\mathcal{D}_{\mathcal{U}}, \quad \mathsf{Q}^2 = \mathcal{T}\mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}}.$$

Now we remove L^2, R^2, P^2, and Q^2 from this generation and put the remaining sets that lie in $\mathcal{D}_{\mathcal{T}}$ into Ω_N and those that lie in $\mathcal{D}_{\mathcal{U}}$ into Ω_M. To that end, define

$$\Omega_N^2 = \mathcal{B}^{-2}\mathcal{D}_{\mathcal{T}}, \qquad \Omega_M^2 = \emptyset.$$

In this case, it follows trivially that these two collections are disjoint from one another,

$$\overline{\Omega_N^2} \cap \overline{\Omega_M^2} = \emptyset. \tag{5.101}$$

Moreover, by the disjointness of non-neighbouring patches, and the fact that $\Omega^2_M = \emptyset$, we obtain

$$\overline{\Omega^2_N} \cap \overline{Q^2} = \emptyset, \tag{5.102a}$$
$$\overline{\Omega^2_N} \cap \overline{R^2} = \emptyset, \tag{5.102b}$$
$$\overline{\Omega^2_M} \cap \overline{P^2} = \emptyset, \tag{5.102c}$$
$$\overline{\Omega^2_M} \cap \overline{L^2} = \emptyset. \tag{5.102d}$$

The next generation looks as follows:

$$L'_1, \quad L'_2, \quad \Sigma_N, \quad P', \quad Q'_1, \quad Q'_2, \quad \Sigma_M, \quad R',$$

where L'_1 and L'_2 denote the children of the patch L^2. P' denotes the child of the patch P. Q'_1 and Q'_2 denote the children of the patch Q. Σ_N denotes the set of children of patches in Ω^2_N. Σ_M denotes the set of children of patches in Ω^2_M (here, this is just the empty set).

In order to understand the above scheme and the following definition of Ω^3_M and Ω^3_N, we make a few observations (it is useful here to refer back to sections 4.2 and 4.2.7):

- Patches with base set $\mathcal{D}_\mathcal{T}$ give rise to two children, whereas patches with base set $\mathcal{D}_\mathcal{U}$ have only one.
- The two boundary patches in each generation $j \geq 1$ are children of the corresponding boundary patches of the previous generation. (Recall that the descendence relation is compatible with the ordering of the patches, i.e. the ordering is *refined* in successive generations.)
- Observe that one boundary patch is tagged by $\mathcal{U}$ and the other by $\mathcal{T}$, so that they have three children together. So precisely one of the three children is not a boundary patch. This one has to be included in the next generation of either Ω_M or Ω_N.
- The two pivotal patches in each generation $j \geq 1$ are children of the pivotal patches of the preceding generation. (Suppose, on the contrary that a pivotal patch descends from a non-pivotal one. Since parents contain their children, it would follow that a non-pivotal patch contains the pivotal point, which is a contradiction.)
- Precisely one of the three children of the two pivotal patches (with $j \geq 2$) is not a pivotal patch. (By the same reasoning as above, we know that the two pivotal patches give rise to three children.) This patch must be included in the next generation of either Ω_M or Ω_N.
- Notice that one pivotal patch has to belong to $\mathcal{N}_j$ and the other one to $\mathcal{M}_j$ (recall that the next symbol changes on the "boundary" between these two sets). In the above example, P' is in $\mathcal{N}_3$, while both Q'_1 and Q'_2 are in $\mathcal{M}_3$. Therefore, P' is certainly pivotal. Since Q'_2 is not a neighbour of P', it is the non-pivotal child, and must be included in the next generation of Ω_M.
- The patches in Σ_N and in Σ_M all have parents that are already included in Ω_M, resp. Ω_N.

The above discussion motivates the following definition:

$$\begin{aligned}
\Omega_N^3 &= \Omega_N^2 \cup L_2', \\
\Omega_M^3 &= \Omega_M^2 \cup Q_2', \\
L^3 &= L_1', \\
P^3 &= P', \\
Q^3 &= Q_1', \\
R^3 &= R'.
\end{aligned}$$

By the disjointness of non-neighbouring patches, and by (5.101 – 5.102d), the conditions (5.101 – 5.102d) are satisfied for $j = 3$ as well. For instance, in order to prove that $\overline{\Omega_N^3} \cap \overline{\Omega_M^3} = \emptyset$, we observe that

- $\overline{\Omega_N^2} \cap \overline{\Omega_M^2} = \emptyset$,
- $\overline{L_2'}$ does not intersect Ω_M^2, because its parent L does not.
- $\overline{Q_2'}$ does not intersect Ω_N^2, because its parent Q does not.
- $\overline{L_2'} \cap \overline{Q_2'} = \emptyset$ by disjointness.

The remaining conditions (5.102a–5.102d) are proved in the same way.

Proceeding by induction, the sequences (Ω_N^j) and (Ω_M^j) thus defined are increasing and stay disjoint for all j. Therefore, their unions Ω_M and Ω_N are disjoint as well. □

5.6.2 Binary Necklace

Having established an open set condition for the accretive Fibonacci system, we now show how this carries over to the binary IFS,

$$\mathcal{N} = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\mathcal{N} \cup \mathcal{B}^{-1}\mathcal{N}. \tag{5.103}$$

Proposition 110
There exists a non-empty open set $\Omega \subset \mathcal{D}_{\mathcal{T}}$, such that

$$\begin{aligned}
\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\Omega &\subset \Omega, && (5.104a)\\
\mathcal{B}^{-1}\Omega &\subset \Omega, && (5.104b)\\
\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\Omega \cap \mathcal{B}^{-1}\Omega &= \emptyset. && (5.104c)
\end{aligned}$$

(This is just the open set condition for the binary IFS.)

Proof: Put $\Omega = \Omega_N$, with the set Ω_N as defined above. The claim now follows easily from the conclusions of the preceding lemma:

$$\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\Omega = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\Omega_N \subset \mathcal{B}^{-1}\Omega_M \subset \Omega_N = \Omega.$$

Further,

$$\mathcal{B}^{-2}\Omega = \mathcal{B}^{-2}\Omega_N \subset \mathcal{B}^{-1}\Omega_N \subset \Omega_N = \Omega,$$

and

$$\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\Omega \cap \mathcal{B}^{-2}\Omega = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}\Omega_N \cap \mathcal{B}^{-2}\Omega_N \subset \mathcal{B}^{-1}\Omega_M \cap \mathcal{B}^{-1}\Omega_N = \emptyset.$$

□

Now, all the machinery is in place to apply the proof of Falconer's Theorem 9.3. (This proof also requires the bounded variation principle, which was established previously.) It finally follows that the numbers r_{2^i} are indeed lower bounds of the Hausdorff dimension of the necklace set.

5.7 Renormalisation Invariance of Hausdorff Dimension

Now that it has been established that the solutions to the partition equations are bounds on the dimension, it remains to see how the dimension is affected by renormalization. In other words, how does the dimension of the critical invariant curve relate to the dimensions of the invariant curves for functions attracted to the (critical) renormalization fixed point.

By analogy with [Sti93c] and section 4.2, we look at "asymptotic self-similarity". In particular, consider the equations

$$\begin{aligned} \mathcal{M}_{j+1}^k &= \mathcal{T}_k\mathcal{B}_{k+1}^{-1}\mathcal{N}_j^{k+1} \\ \mathcal{N}_{j+1}^k &= \mathcal{B}_{k+1}^{-1}(\mathcal{M}_j^{k+1} \cup \mathcal{N}_j^{k+1}), \end{aligned}$$

which define the necklace construction in the asymptotically self-similar case. Letting $j \to \infty$, gives

$$\begin{aligned} \mathcal{M}^k &= \mathcal{T}_k\mathcal{B}_{k+1}^{-1}\mathcal{N}^{k+1} \\ \mathcal{N}^k &= \mathcal{B}_{k+1}^{-1}(\mathcal{M}^{k+1} \cup \mathcal{N}^{k+1}). \end{aligned}$$

Here, $\mathcal{M}^k$ and $\mathcal{N}^k$ are pieces of the necklace curve of the k-th iterate under renormalization. Denote by d_k, respectively e_k, the Hausdorff dimensions of $\mathcal{M}^k$, respectively $\mathcal{N}^k$:

$$d_k = \dim_H(\mathcal{M}^k), \qquad e_k = \dim_H(\mathcal{N}^k).$$

Since the maps $\mathcal{T}_k\mathcal{B}_{k+1}^{-1}$ and $\mathcal{B}_{k+1}^{-1}$ are bi-Lipschitz (*close enough to the fixed point*), we obtain

$$\begin{aligned} d_k &= e_{k+1}, \\ e_k &= \max\{d_{k+1}, e_{k+1}\}. \end{aligned}$$

Assuming, that $d_{k+1} \le e_{k+1}$, we conclude that $d_k = e_k$. If not, i.e., if $d_{k+1} > e_{k+1}$, we obtain $d_k = e_{k+1}$ and $e_k = d_{k+1}$, which implies that $d_k < e_k$. This, according to the first

remark, yields $d_{k-1} = e_{k-1}$. In other words, if the dimensions d and e are equal for any iterate, then they are equal for all of the previous iterates. By choosing k high enough, we see that indeed $e_0 = d_0$.

It follows that $e_k = d_k = d$ for all k, i.e., the Hausdorff dimension of the invariant sets of all the iterates are all the same; Hausdorff dimension is a renormalisation invariant for these maps.

At this stage, the value of d could still depend on the orbit. However, it turns out that the proof of [Fal90, Theorem 9.3] goes through for asymptotic self-similarity as well. (The proof of the open set condition, which above was verified for the fixed point itself, can be adapted to the asymptotically self-similar case too, in the same way that the results of [Sti93c, §4] are obtained.)

It follows that the Hausdorff dimension of the invariant sets of all the iterates is equal to the Hausdorff dimension of the invariant set of the renormalisation fixed point. In other words, Hausdorff dimension is invariant under renormalisation for maps that are both (1) attracted to the fixed point and (2) close enough to it, in the sense that the maps $\mathcal{T}_k\mathcal{B}_{k+1}^{-1}$ and $\mathcal{B}_{k+1}^{-1}$ are bi-Lipschitz for all k.

5.8 The Inverse (TU) Necklace

In order to get rigorous numerical values for the Hausdorff dimension, it is easier to work with the functions U and V of the existence proof directly, rather than with their analytic continuations on the new domains (Δ_U and Δ_V). This means looking at the inverse composition order $\mathcal{TU}$, and the corresponding necklace construction. Recall that section 5.1.3 demonstrated how the inverse IFSs arise.

The question now arises of whether the Hausdorff dimensions of the inverse and accretive necklace sets agree. In order to see that they do, consider the following points:

- $\dim_H \mathcal{M} = \dim_H \mathcal{N}$.
- $\dim_H \tilde{\mathcal{M}} = \dim_H \tilde{\mathcal{N}}$.
- $\tilde{\mathcal{M}} \subset \mathcal{M} \cup \mathcal{N}$. This implies that $\dim_H \tilde{\mathcal{M}} \leq \dim_H(\mathcal{M} \cup \mathcal{N})$.
- $\dim_H \tilde{\mathcal{M}} = \dim_H(\tilde{\mathcal{M}} \cup \tilde{\mathcal{N}} \cup \mathcal{B}\tilde{\mathcal{N}})$, and the latter set contains $\mathcal{M} \cup \mathcal{N}$. This implies that $\dim_H \tilde{\mathcal{M}} \geq \dim_H(\mathcal{M} \cup \mathcal{N})$.

Finally, it remains to verify the open set condition for the inverse necklace.

5.8.1 Open Set Condition for the $\mathcal{TU}$ Necklace

Although an analogue of the open set condition can be proved for the inverse necklace construction it does not, unfortunately, seem to imply the corresponding open set condition for the inverse binary IFS. The proof for the necklace construction is not very exciting, but will be sketched here for completeness.

Using the generic notation, the fixed point equation is

$$\mathcal{U} = \mathcal{B}\mathcal{T}\mathcal{B}^{-1} \tag{5.105a}$$
$$\mathcal{T} = \mathcal{B}\mathcal{T}\mathcal{U}\mathcal{B}^{-1}, \tag{5.105b}$$

where the first equation holds on $\mathcal{D}_{\mathcal{U}}$ and the second one on $\mathcal{D}_{\mathcal{T}}$. (Notice: $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{\mathcal{T}}$ stand for the domains of the existence proof.)

Recall the *domain extension conditions*

$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{U}}} \subset \mathcal{D}_{\mathcal{T}} \tag{5.106a}$$
$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \subset \mathcal{D}_{\mathcal{U}} \tag{5.106b}$$
$$\mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \subset \mathcal{D}_{\mathcal{T}}. \tag{5.106c}$$

In addition, an extra *disjointness condition* will be needed:

$$\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} \cap \mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_{\mathcal{T}}} = \emptyset. \tag{5.107}$$

(Note: this condition resembles one of the usual disjointness conditions, but involves the domains D_{U} and D_{V} of the existence proof given in chapter 2, rather than the domains Δ_{U} and Δ_{V} used in chapter 4.)

Proposition 111
There exist non-empty disjoint *open sets* Ω_{N} *and* Ω_{M} *such that*

$$\Omega_{\mathsf{N}} \subset \mathcal{D}_{\mathcal{T}}$$
$$\Omega_{\mathsf{M}} \subset \mathcal{D}_{\mathcal{U}}$$

and that

$$\mathcal{B}^{-1}\Omega_{\mathsf{M}} \subset \Omega_{\mathsf{N}}, \tag{5.108a}$$
$$\mathcal{B}^{-1}\Omega_{\mathsf{N}} \subset \Omega_{\mathsf{M}}, \tag{5.108b}$$
$$\mathcal{U}\mathcal{B}^{-1}\Omega_{\mathsf{N}} \subset \Omega_{\mathsf{N}}. \tag{5.108c}$$

Overview: Once again, the idea is to construct the sets Ω_{M} and Ω_{N} as the limits of increasing sequences of open sets. This time, however, we cannot use the usual necklace construction and simply omit some of the patches, and must therefore adopt a different approach. The idea is to start with the set $\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$ as the first building block of Ω_{N}. (Notice that, by (5.106c), it is contained in $\mathcal{D}_{\mathcal{T}}$.) Then we start iterating the map $\mathcal{B}^{-1}$. The set $\mathcal{B}^{-1}\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$ goes into Ω_{M}. The set $\mathcal{B}^{-2}\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$ goes back into Ω_{N}. The set $\mathcal{B}^{-3}\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_{\mathcal{T}}$ goes again into Ω_{M}, etc. (By (5.106a) and (5.106b), the increments are alternatingly in $\mathcal{D}_{\mathcal{U}}$ and in $\mathcal{D}_{\mathcal{T}}$. It follows that $\Omega_{\mathsf{N}} \subset \mathcal{D}_{\mathcal{T}}$ and that $\Omega_{\mathsf{M}} \subset \mathcal{D}_{\mathcal{U}}$.)

Intuitively, as the construction proceeds two "tentacles" come out of the sets Ω_{N} and Ω_{M} and grow towards each other, becoming thinner and thinner and coming closer and closer to the origin, one from the left, the other from the right. At each step, either the right

hand tentacle grows and the other remains stationary, or the other way around. (Their closures will intersect precisely at the origin.)

This construction has the advantage of making the *invariance condition* (5.108) explicit: Since $\Omega_N \in \mathcal{D}_\mathcal{T}$,

$$\mathcal{U}\mathcal{B}^{-1}\Omega_N \subset \mathcal{U}\mathcal{B}^{-1}\mathcal{D}_\mathcal{T},$$

and the latter set was the first building block of Ω_N. (The remaining conditions also follow directly from the construction.)

The resulting sets are non-empty and open. Also, it was already observed that the domain extension conditions imply that $\Omega_M \subset \mathcal{D}_\mathcal{U}$ and that $\Omega_N \subset \mathcal{D}_\mathcal{T}$.

It remains to show that the sets Ω_N and Ω_M are *disjoint*.

Proof: First, we formalise the construction outlined above. To this end, put

$$\Omega_N^0 = \emptyset = \Omega_M^0.$$

Then, for *even* $j \geq 0$, put

$$\Omega_N^{j+1} = \Omega_N^j \cup \mathcal{B}^{-j}\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_\mathcal{T}, \tag{5.109a}$$
$$\Omega_M^{j+1} = \Omega_M^j. \tag{5.109b}$$

For *odd* j, put

$$\Omega_N^{j+1} = \Omega_N^j, \tag{5.110a}$$
$$\Omega_M^{j+1} = \Omega_M^j \cup \mathcal{B}^{-j}\mathcal{U}\mathcal{B}^{-1}\mathcal{D}_\mathcal{T}. \tag{5.110b}$$

The sequence of pairs (Ω_M^j, Ω_N^j) is well defined and increasing. It is required to prove that the limit sets

$$\Omega_N = \bigcup_j \Omega_N^j, \quad \Omega_M = \bigcup_j \Omega_M^j,$$

are disjoint. It is enough to prove that the sets Ω_N^j and Ω_M^j are disjoint for all values of j. This can be done by verifying that a certain set of conditions holds for even values of j and that a second set of conditions holds to odd values of j. The conditions for even j are given by

$$\overline{\Omega_N^j} \cap \overline{\Omega_M^j} = \emptyset \tag{5.111a}$$
$$\mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{U}} \cap \overline{\Omega_N^j} = \emptyset \tag{5.111b}$$
$$\mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_M^j} = \emptyset \tag{5.111c}$$

The conditions for odd j are given by

$$\overline{\Omega_N^j} \cap \overline{\Omega_M^j} = \emptyset \tag{5.112a}$$
$$\mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{U}} \cap \overline{\Omega_M^j} = \emptyset \tag{5.112b}$$
$$\mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_N^j} = \emptyset \tag{5.112c}$$

Even Case For $j = 0$, the even conditions are trivially satisfied. Assume that j is even and greater than or equal to zero and that the even conditions hold for j. The aim is to verify the odd conditions for $j+1$.

(5.112a): Using the definition (5.109), the domain extension condition (5.106c), and the induction hypothesis (5.111c), we obtain

$$\begin{aligned}\overline{\Omega_N^{j+1}} \cap \overline{\Omega_M^{j+1}} &\subset (\overline{\Omega_N^{j}} \cap \overline{\Omega_M^{j}}) \cup (\mathcal{B}^{-j}\mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_M^{j}}) \\ &\subset \emptyset \cup (\mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_M^{j}}) \\ &= \emptyset.\end{aligned}$$

(5.112b): Using the definition (5.109), the domain extension condition (5.106a), and the induction hypothesis (5.111c), gives

$$\mathcal{B}^{-(j+1)}\overline{\mathcal{D}_\mathcal{U}} \cap \overline{\Omega_M^{j+1}} \subset \mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_M^{j}}, = \emptyset.$$

(5.112c): Using the domain extension condition (5.106b), the induction hypothesis (5.111b), and the disjointness condition (5.107):

$$\begin{aligned}\mathcal{B}^{-(j+1)}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_N^{j+1}} &\subset (\mathcal{B}^{-(j+1)}\overline{\mathcal{D}_\mathcal{T}} \cap \overline{\Omega_N^{j}}) \\ &\qquad \cup (\mathcal{B}^{-(j+1)}\overline{\mathcal{D}_\mathcal{T}} \cap \mathcal{B}^{-j}\mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_\mathcal{T}}) \\ &\subset (\mathcal{B}^{-j}\overline{\mathcal{D}_\mathcal{U}} \cap \overline{\Omega_N^{j}}) \\ &\qquad \cup \mathcal{B}^{-j}(\mathcal{B}^{-1}\overline{\mathcal{D}_\mathcal{T}} \cap \mathcal{U}\mathcal{B}^{-1}\overline{\mathcal{D}_\mathcal{T}}) \\ &= \emptyset.\end{aligned}$$

We now have verified all the odd conditions for $j+1$. The second case follows by analogous reasoning (we assume that j is odd and that the odd conditions hold and proceed to verify that the even conditions hold for $j+1$). □

Unfortunately, no way was seen to carry the above proof over to the corresponding binary IFS.

5.9 Results

In this section, we present the results that were obtained from rigorous computation of lower and upper bounds on the dimension. Firstly, the inverse composition order was used (in fact, the inverse regular binary IFS) as this enabled the maps U and V themselves to be used on simple (disc) domains of definition. Recall that the open set condition has not been proved for this IFS and, therefore, that the calculations for the *lower* bounds (although rigorous in themselves) cannot be said to give true lower bounds.

The rigorous framework was then extended to support the analytic extensions of the maps to the domains for the accretive order and the relevant (accretive regular binary) IFS was used. Because of the need to use the analytic extensions of the maps, the calculations

were found to be extremely time-consuming. Indeed, I did not manage to get rigorous lower bounds in the time available. However, rigorous upper bounds for the dimension were obtained and were found to be surprisingly good.

Numerically (i.e. non-rigorously) the dimension is found to be in the range

$$1.00119 \leq \dim_{\mathrm{H}} \leq 1.07967.$$

(This figure was calculated by Stirnemann [Sti].) Notice, in particular, that the dimension seems to be very close to 1. Finding a useful lower bound will therefore be extremely difficult. (Indeed, in order to obtain the above numerical estimate it was necessary to use the *eleventh* generation of the IFS.)

5.9.1 Summary of the Technique

Upper Bounds

Upper bounds are the easiest to obtain, since an upper bound on the magnitude of the derivative of a function over a convex domain directly gives a Lipchitz constant for the function on that domain. For the maps used here, these Lipschitz constants will be less than one and are, therefore, *contractivity* constants.

Here are the relevant bits of theory:

Lemma 112
Let f *be an analytic map, well-defined and differentiable on a convex set* c. *Then (by the mean value theorem), the quantity*

$$\max_{\mathsf{x}\in\mathsf{c}} |\mathsf{D}\mathsf{f}(\mathsf{x})|,$$

is a Lipschitz constant for f *on* c.

Lemma 113
Let f, g *be functions. Let* a *be a Lipschitz constant for* g *on a set* c. *Let* b *be a Lipschitz constant for* f *on a set* $\mathsf{d} \supseteq \mathsf{g}(\mathsf{c})$. *Then* ab *is a Lipschitz constant for* $\mathsf{f}(\mathsf{g})$ *on* c.

Recall that in order to get upper bounds for the dimension, we require contractivity constants on the binary patches of some generation of the binary IFS. By plugging these constants into the upper partition equation (5.88b) and solving for s, we get an upper bound for the dimension.

We take some initial (convex) domain P that contains the invariant set (to check this we ensure that it contains its images under both maps). Recall that the first generation of the binary IFS consists of the images of the initial domain under the two constitutive maps ψ_{L} and ψ_{R}, one of which is analytic, the other being anti-analytic. For the binary IFSs, the anti-analytic map is in fact an anti-similarity and a contractivity constant for this is thus trivial to obtain from the magnitude of the rescaling. For the analytic map, the above lemma tells us that an upper bound for the modulus of its derivative on a convex domain gives a contractivity constant over the domain.

So we can calculate a contractivity constant for each of the maps producing the two patches $\psi_L(P)$ and $\psi_R(P)$ in the first generation. These contractivity constants can be plugged into the upper partition function to give an upper bound on the Hausdorff dimension. The next step is to see how these results can be used to get contractivity constants for higher generations.

For example, to see how we obtain bounds for the contractivities of the maps for the second generation (namely $\psi_L\psi_L$, $\psi_L\psi_R$, $\psi_R\psi_L$, and $\psi_R\psi_R$), we apply the two maps of the IFS to the initial domain to produce two images $\psi_L(P)$ and $\psi_R(P)$. We then check that these images lie within the original domain. (The original domain will be taken to be either a disc or a convex polygon, and the images will be covered by coarse-curves. The corresponding containment test is the same as that used to verify the existence of a fundamental domain in section 4.5.9.)

The next step is to find contractivity constants for the maps ψ_L and ψ_R on each of the images $\psi_L(P)$ and $\psi_R(P)$. A problem arises because the images $\psi_L(P)$ and $\psi_R(P)$ might not be convex. The solution is to find some convex supersets of them and use the same procedure as above to get contractivity constants on these sets instead. Since we have taken supersets, the numbers thus obtained are also contractivity constants on the images themselves. The idea is that we can now use the above lemma concerning compositions of maps get contractivity constants for the four binary maps $\psi_L\psi_L$, $\psi_R\psi_L$, $\psi_L\psi_R$, $\psi_R\psi_R$ on the initial domain P, by multiplying the constants found for ψ_L and ψ_R on P in the previous step by those just found for ψ_L and ψ_R on each of $\psi_L(P)$ and $\psi_R(P)$. Now these new contractivities can be plugged into the partition equation yielding (hopefully) a better upper bound.

This process is then repeated for subsequent generations. At each stage we combine the contractivity bounds already obtained for the previous generation with the bounds for contractivities of the individual maps ψ_L and ψ_R on each of the binary patches.

Discussion

Some important considerations are raised by the description of the above method.

Finding an initial domain: The shape of the initial domain chosen is important to the success (or otherwise) of the method. As was the case for the existence proof presented in chapter 2, choosing a good initial domain is something of a black art. Firstly, the initial domain for the IFS must be within the domains of definition of the relevant maps and must contain the invariant set for the IFS. (It is not practical to use the boundaries of the domains of the maps themselves. In particular, for the accretive case, one of the domains is a union of three discs and is not convex.) In fact, it is desirable that the initial domain P is contained well inside the domains of the maps with some room to spare, particularly when using the analytic continuations, otherwise the numerics are not sufficiently good to enable us to prove that the images of the initial domain are contained within it.

Secondly, we need that both maps are homeomorphisms on the initial domain chosen (in order that the results concerning exponential disjointness, etc., are valid). We must

therefore choose a domain that is contained within the coffin-shaped fundamental domain presented in section 4.5.9. (Alternatively, we would need to find another fundamental domain which did contain the initial domain.)

The domain chosen, for example the interior of some polygon, can be refined by calculating high-generation approximations to the necklace and choosing a polygon that just contains them with a little room to spare. This is demonstrated schematically in Figure 5.1, which shows an initial polygonal domain P that is improved to give a domain P′ by looking at the images of P for the third generation of the binary IFS.

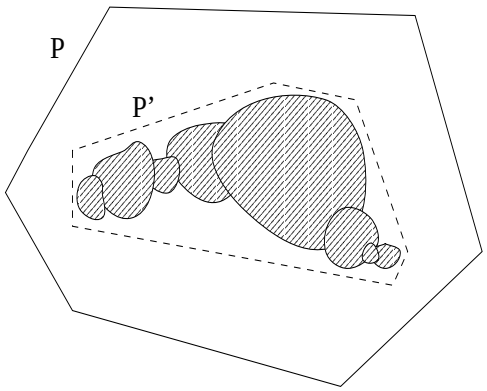

Figure 5.1: Improving an initial domain.

Convex Bounding Sets: A major difficultly with the above method is the need to obtain bounds for the derivatives on *convex* sets: some rigorous method has to be used for constructing a convex superset of a given coarse-curve. Whilst it is conceptually simple to construct a *convex hull* for a set of points in the plane, it turns out to be very difficult to make the construction rigorous for the kinds of objects that we deal with here. To appreciate one source of difficulty, consider three points which lie on a straight line: to machine accuracy, it may be impossible to determine whether the central point lies to the left or to the right of the line joining the other two or whether, indeed, it actually lies *on* the line. For this reason alone, it can be difficult to rigorously construct a convex bounding polygon directly from the grains of a given coarse curve. A second problem is the need to rigorously prove that the resulting convex set really contains the object that it was designed to. In particular, if the bounding set fits too closely then it may be impossible to prove containment using rigorous numerics.

It was realised that the difficulty of constructing a convex set could be bypassed if we were willing to settle for getting bounds on just the first few generations. For the zeroth generation, the initial domain may be chosen to be a convex polygon anyway, so there is no difficulty there. For successive generations, we could then take the bounding polygons of the previous generation and apply the maps of the IFS to their individual vertices. This produces another polygon which is not necessarily convex, nor will it necessarily contain the true image of the original polygon under the map, since we have only applied the map to the vertices, not the entire boundary. However, the hope is that the distortion produced

by applying the maps of the IFS will not be too great, so that this new polygon will be of roughly the right shape to contain the true image, and will in fact be convex. This polygon will be called a *schematic patch*, since applying the maps to the vertices of the polygons (rather than to whole boundary) provides a method for quickly generating approximations to the patches of a generation and, therefore, for drawing generations schematically.

The idea is to construct such a schematic polygon and then test whether is it actually convex or not. If not, the program simply halts with a suitable message and the search for bounds is at an end. On the other hand, if the polygon *is* convex then it may be suitably scaled until it fits over the relevant set. Since, of course, linear rescaling and translation preserve convexity, the resulting polygon is then a suitable convex superset to use for getting bounds on the derivatives of the maps.

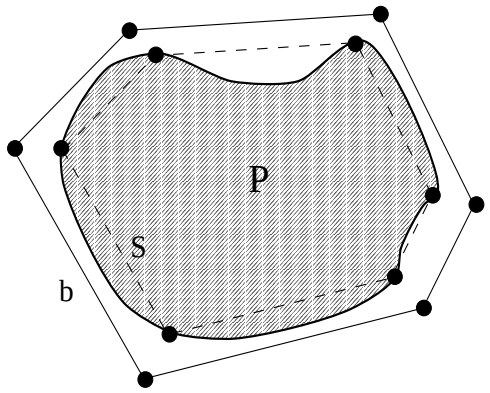

Figure 5.2: Constructing a bounding polygon b from a schematic S.

Figure 5.2 demonstrates this schematically. Suppose that the image of the initial domain under some map is the patch P in the figure. By mapping the individual vertices of the initial domain, we produce the "schematic" polygon S (represented by the dashed-lines in the figure). Notice that S has roughly the same shape as P. A suitably scaled and translated copy of S will contain P, this is the *bounding polygon* b. Notice that b must be convex if S is convex. If this is so, then an upper bound on the derivative of a map on the region bounded by b gives a Lipschitz constant of that map on the patch P.

For the inverse regular binary IFS with a suitable initial polygonal domain, it turns out that this process can be successfully repeated until the eleventh generation, producing quite a reasonable upper bound for the dimension. In addition, by the time this condition fails, the resulting patches are small enough that a simple bounding rectangle may be used instead. The program can therefore be made to automatically switch over to using a bounding rectangle if a difficulty arises.

One alternative possibility is to compute a *bounding circle* rather than a bounding polygon. The construction of such an object is very easy: choose some central point and simply calculate rigorously the maximum distance of the grains of the coarse-curve covering the relevant patch from this point. The resulting number may be used as the radius of a disc that is guaranteed to contain the set. It would have been nice if this method had been enough, however for the first few generations the resulting circles simply do not fit the

invariant set of the IFS tightly enough (in fact, when using the analytic continuations of the maps, the resulting circles are not always completely contained within the domains of definition of the maps, which is disastrous).

(Another alternative method which avoids the explicit construction of a convex polygon would be to simply take an existing convex polygon, for example the initial domain itself, and perform suitable translation and rescaling to fit it over the set in question.)

Constructing Fresh Coverings: An additional problem arises. We are using the coarse-curve construction to apply compositions of the maps of the binary IFS rigorously to the boundary of the initial domain. When a coarse-curve is first created to cover the boundary of the initial domain, we choose to use a large number of small grains to give tight bounds. After the application of a few maps, however, the image coarse-curves will contain grains that are somewhat larger (intuitively speaking, they have expanded in order to absorb the floating point errors and the uncertainty expressed by the radii of the function balls used). After a few generations the resulting grains are prone to "explode" due to absorbing the accumulated error bounds. (This problem is particularly apparent when using the analytic extensions of the maps for the accretive composition order.) The result is that the bounds are either not tight enough to be of any use (in fact, they are typically worse than those obtained in earlier generations), or that the grains are so large that they protrude from the domain of definition of one of maps, causing the program to exit immediately with a suitably righteous error message. Increasing the number of grains used to cover the initial domain up to the highest practical value improves the bounds, but typically after one or two generations more the problem occurs again.

The solution adopted for this problem is to apply the maps to the bounding polygons at each stage, rather than to the coarse-curve coverings used in the previous generation. Since the regions defined by the bounding polygons are (by definition) supersets of the regions bounded by the coarse-curves, this is strategy is justified. This means that for each generation a covering is created anew over the bounding polygon using small grains. This is demonstrated in figure 5.3(a), where we look at the fate of a single grain $\mathfrak{a}$ forming part of a coarse-curve covering some initial polygon. After several rigorous applications of a function ball, the resulting image grains explode. In figure 5.3(b), a new covering (of the bounding polygon) is constructed after the function ball is applied, using grains of a similar size to the original ones, preventing the accumulation of errors.

Solving the Partition Equations: Recall that the partition equations involve sums of powers of the contractivity, respectively coercivity, constants. The method chosen for solving these rigorously was very simple. First the partition equation is solved numerically (i.e. non-rigorously) by using a bisection algorithm to repeatedly refine numerical lower and upper bounds on the solution. The partition equations are then evaluated rigorously at the value produced (using rigorous routines to perform exponentiation), in order to show that it is a genuine upper or lower bound on the true solution. (If the rigorous check were to fail, it would suffice to increment or decrement the non-rigorous solution by a small amount until it succeeded.)

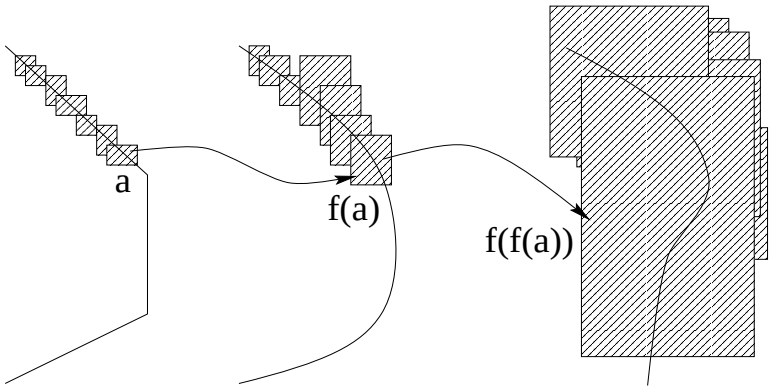

(a) Repeated application of a function ball,
resulting in an "explosion" of error bounds.

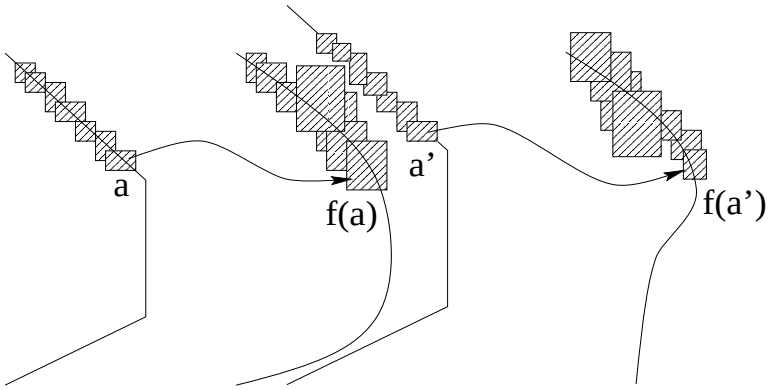

(b) Constructing a fresh covering at each stage,
avoiding accumulation of errors.

Figure 5.3: Re-covering boundaries with coarse-curves.

Lower Bounds

Lower bounds are much harder to obtain computationally, since it is required to estimate the magnitude of the derivative of a function over the *pre-image* of a convex set. This entails all of the problems mention above for the calculations of upper bounds, and more. In the time available, bounds were only found for the inverse case. It is hoped that bounds might be found for the accretive case in future (in particular, this would give a genuine lower bound for the dimension).

The strategy used is essentially the same as that presented above, with the exception that we now need to find a convex *co-bounding set* of each patch, whose image under the map contains a convex superset of the child of the patch. We then get a rigorous lower bound for the coercivity constant of the non-trivial map $\mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1}$ evaluated on the co-bounding set. This is done by obtaining a lower bound for the derivative on the boundary of the co-bounding set. The other map is either the anti-analytic map $\mathcal{B}^{-1}$ (accretive IFS) or $\mathcal{B}^{-2}$ (inverse IFS) and is simply an (anti) linear rescaling, so that it is trivial to get a coercivity constant.

In order to see how this difficulty may be handled, notice that a convex superset of a child patch is constructed anyway in the form of the child's bounding polygon (on which we estimate the derivative to get upper bounds). A method is therefore needed for finding a convex co-bounding set whose image under the relevant map contains the child's bounding polygon.

The method used was to start with the parent's bounding polygon and repeatedly expand it until its image under the map was found to contain the child's bounding polygon. The resulting polygon is convex, since linear scaling preserves convexity.

This strategy worked well for the inverse order, except for finding a co-bounding polygon for the initial domain, which had to be done by hand (the reason for the failure of the method was that the bounding polygon had to be expanded to the point where it protruded from the domain of definition of one of the maps).

For the accretive order a suitable set of co-bounding polygons was not found in the time available. The reason was that the initial domain that had to be chosen was too close to the edge of the domain of definition of the maps. The hope is that this difficulty may be overcome in future.

5.9.2 Results for the Inverse Composition Order

Recall that for the inverse regular IFS, the maps used are as follows:

$$\psi_L(z) = \mathcal{B}^{-2} = \alpha^2 C \alpha^2 C(z) = |\alpha|^4(z), \tag{5.113a}$$

$$\psi_R(z) = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1} = \alpha^2 C Q U \alpha^2 C(z), \tag{5.113b}$$

with the initial domain being a subset of D_V.

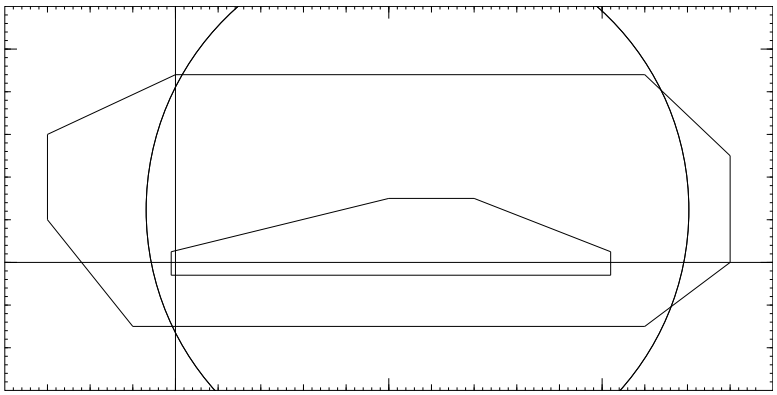

Figure 5.4: Initial domain for the inverse regular IFS in relation to the fundamental domain and the domain D_V.

Figure 5.4 shows a suitable initial domain in relation to the domain D_V and the fundamental domain. This polygon was found after much experimentation and is thought to be pretty optimal. (In particular, the images of the polygon under each of the maps are found to be contained inside the polygon, fitting reasonably closely.)

Figure 5.5 shows the schematic patches for some of the generations of the IFS. Notice, in particular, the resemblance between the final picture and the critical invariant curve shown in figure 1.9 in section 1.4.

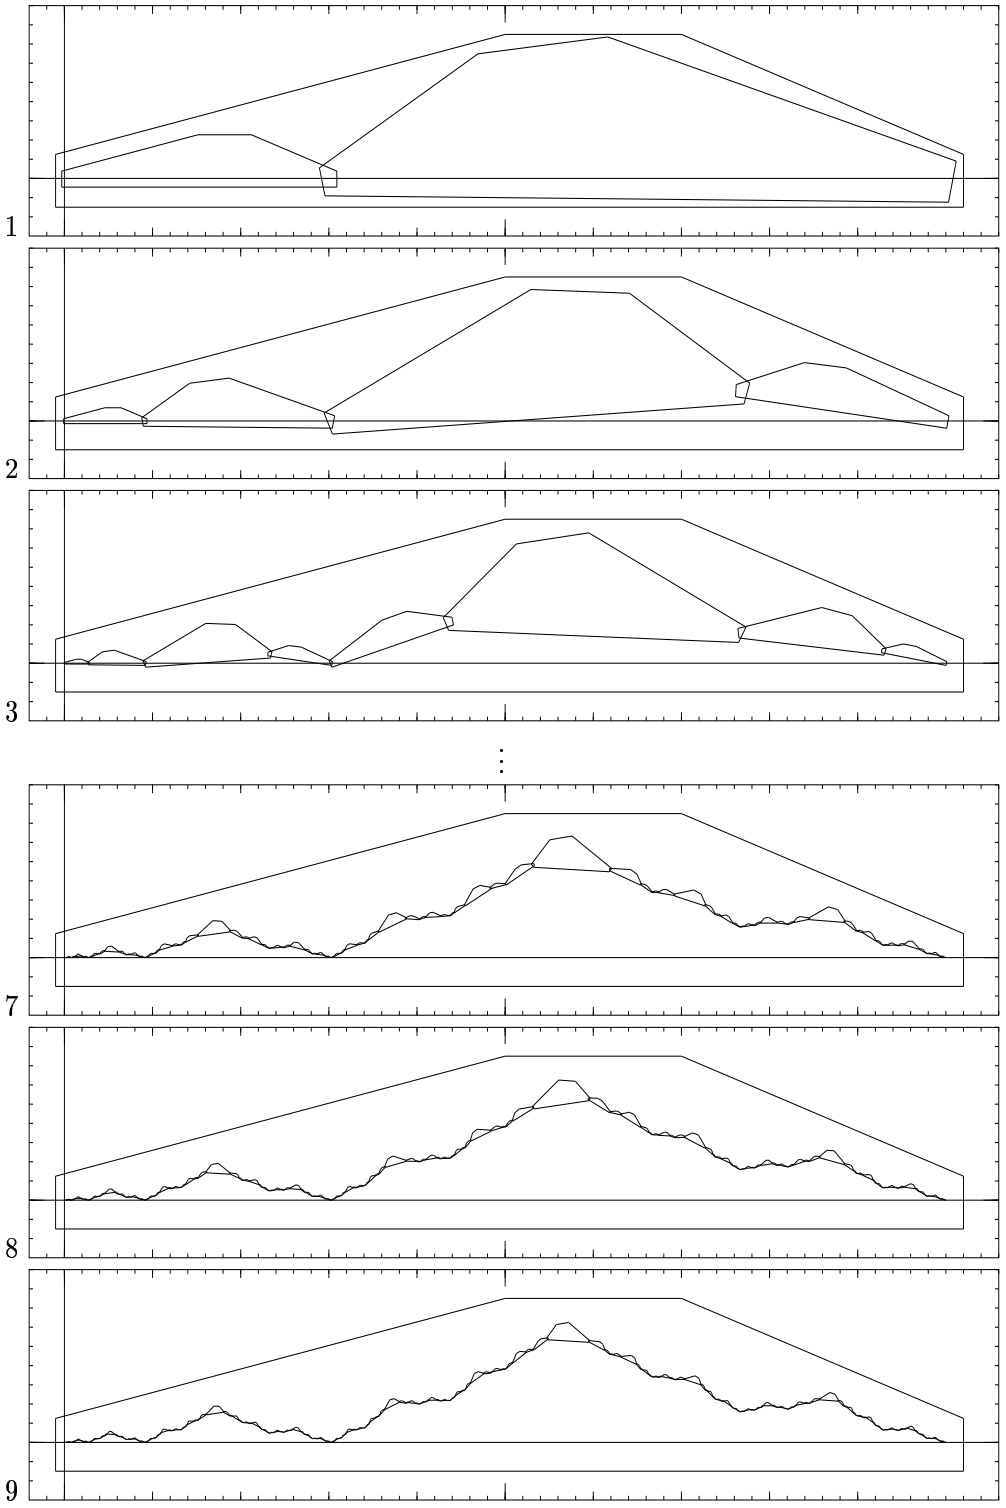

Figure 5.5: Schematic picture of the generations of the inverse regular IFS.

j	Upper bound (a)	Upper bound (b)	Time (hh:mm:ss)
1	1.33807	1.21975	00 : 00 : 15
2	1.31507	1.20298	00 : 00 : 30
3	1.26769	1.17154	00 : 00 : 56
4	1.23477	1.15548	00 : 01 : 50
5	1.20544	1.13904	00 : 03 : 39
6	1.18253	1.12748	00 : 07 : 12
7	1.16353	1.11694	00 : 14 : 49
8	1.14819	1.10864	00 : 29 : 28
9	1.13546	1.10124	00 : 58 : 57
10	1.12491	1.09515	01 : 54 : 58
11	1.11605	1.08975	03 : 50 : 31
12	1.10858	1.08523	07 : 49 : 09

Figure 5.6: Upper bounds for the inverse composition order

Figure 5.6 shows the upper bounds that were obtained by solving the corresponding upper partition functions (a) for a less suitable polygon, (b) for the optimised polygon shown in Figure 5.4. In each case, 2048 grains were used to cover the boundary of each patch and to evaluate the necessary derivatives. The time taken for each generation is shown (the program was run on a 143 MHz. Ultra Sparc). (Recall that generation j contains 2^j patches and that the corresponding upper partition equation contains 2^j terms.) Figure 5.7 shows these estimates graphically.

When finding lower bounds, the program is very sensitive to the strategy used for finding the co-bounding polygons. Reliable bounds were produced for the first seven generations of the IFS, and are shown below (after which, the method that was being used to find the co-bounding polygons became prohibitive).

j	Lower bound
1	0.75355
2	0.80751
3	0.84265
4	0.87065
5	0.89175
6	0.90855
7	0.92189

Figure 5.8: Lower bounds using the inverse regular IFS

Unfortunately, no lower bound greater than 1 was found. It is hoped that this result will be improved in future by using a better strategy to find bounding and co-bounding sets, and that the open set condition for this composition order will be verified.

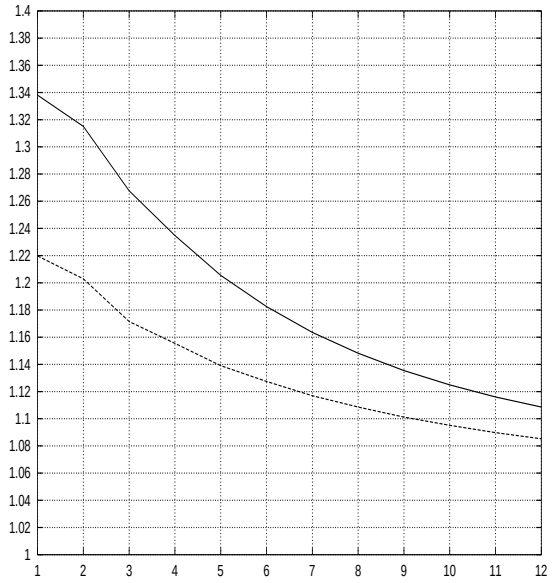

Figure 5.7: Upper bounds on the dimension

5.9.3 Results for the Accretive Composition Order

Recall (section 5.3) that the maps for the accretive regular IFS are given by

$$\psi_L = \mathcal{B}^{-1} = \alpha^2 C, \tag{5.114a}$$

$$\psi_R = \mathcal{B}^{-1}\mathcal{T}\mathcal{B}^{-1} = \alpha^2 CQV\alpha^2 C. \tag{5.114b}$$

In fact, for the map ψ_R, we must work with the analytic continuations for V, each of which contains several functional compositions. This makes the situation somewhat hopeless for rigorous numerics.

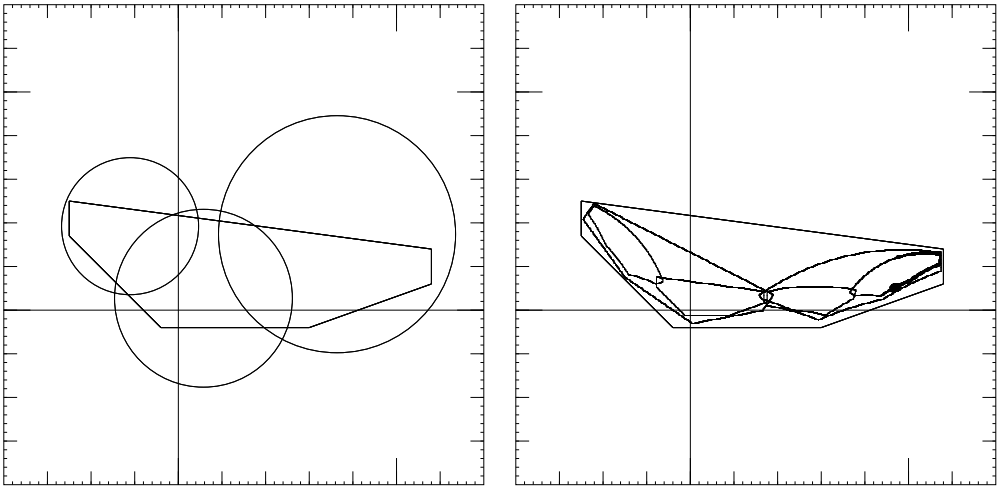

(a) Initial Domain of IFS (b) First and Second Generations

Figure 5.9: The Accretive Regular IFS

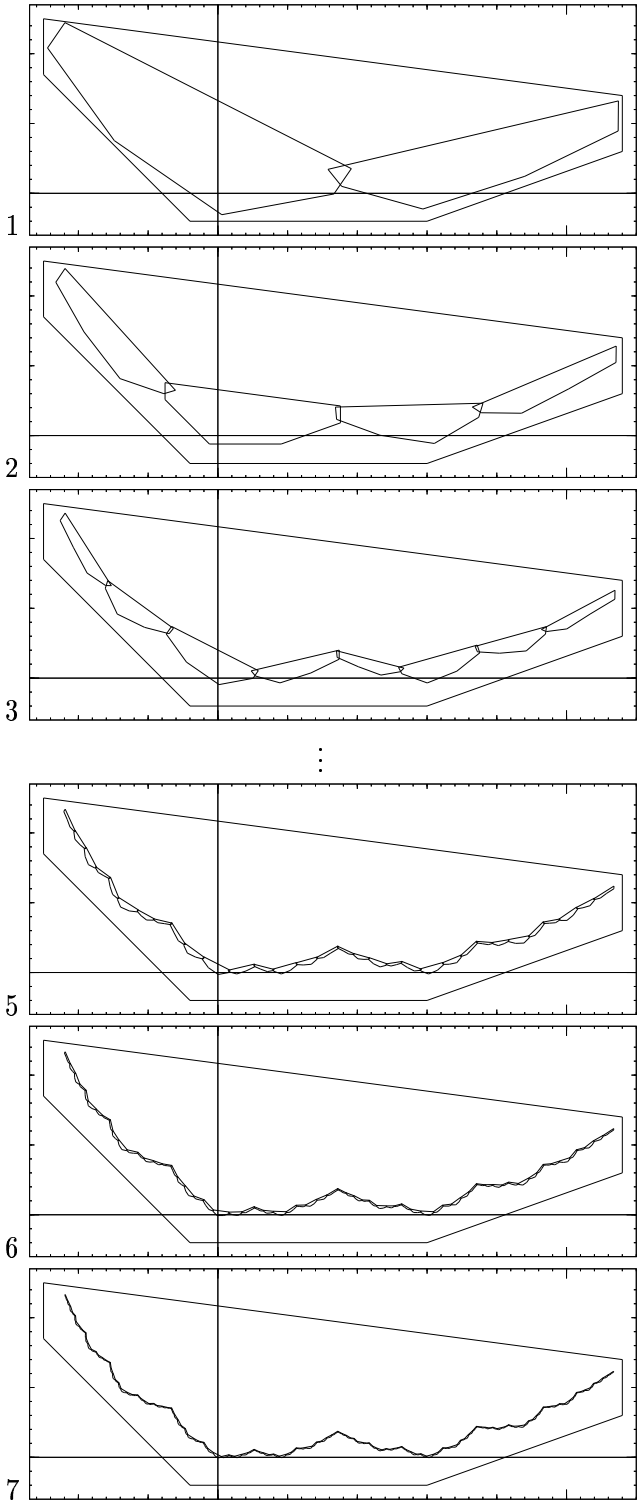

Figure 5.10: Schematic picture of the generations of the accretive regular IFS.

Figure 5.9 shows (a) the initial polygon that was found for the accretive order in relation to the discs making up the domain Δ_V, and (b) the first two generations of the IFS, calculated using coarse-curves (a schematic representation of several generations is shown

in figure 5.10). Simply proving that the images of this polygon under the two maps are contained within it requires the use of a very large number of grains in each coarse-curve (of the order of a million). Evaluating the derivative on the resulting patches requires even more. In addition, the initial polygon has to be taken very close to the edges of the domains of definition of the maps, which means that it is difficult to find useful convex bounding polygons on which to evaluate the derivative (the situation for finding co-bounding polygons for the lower bounds is even more hopeless). As a result, useful bounds were only obtained for the first and second generation, namely upper bounds of

j	Upper bound
1	1.203
2	1.160

(It is worth noting that the computation time to obtain these bounds was of the order of two or three days, so that calculating higher generations would be impractical using this method.)

Even numerically (i.e. non-rigorously) it is difficult to obtain a good lower bound on the dimension using this composition order. In fact, it is necessary to go to the eleventh generation to get a numerical lower bound on the dimension that exceeds 1 [Sti]. In view of these difficulties, it would probably be necessary to re-program the rigorous framework using discs instead of rectangles in order that the error bounds accumulated during functional compositions are minimised.

5.10 Future work on dimension

The previous sections attempted to establish both lower and upper bounds for the Hausdorff dimension of the Siegel disc boundaries. Recall that the results for the inverse case are extremely good, in particular a very good upper bound of 1.08523 was rigorously established. For lower bounds, the corresponding open set condition has not been verified. It would be nice if this could be done. In addition, for the accretive order, it might be possible to find a satisfactory initial domain and produce a useful lower bound. As mentioned above, it would be advisable to re-write the rigorous computer framework using discs instead of rectangles. (However, it is possible that the numerics are sufficiently badly-behaved for the accretive composition order that finding a lower bound above 1 is impractical using rigorous methods.)

The invariant object associated with an IFS is actually a *measure*. It turns out that, in this case, the invariant measure of the IFS is identical with the dynamically invariant measure on the Siegel disc boundary. This means that it may be possible to use a similar approach to analyse the so-called *multifractal spectrum* of the dynamically invariant measure. This problem is, however, rather more complicated than dealing with the Hausdorff dimension.

Bibliography

[Ahl79] Ahlfors L V (1979): *Complex Analysis.* 3rd ed. McGraw-Hill.

[Arn63a] Arnold V I (1963): *Proof of A N Kolmogorov's theorem on the preservation of quasiperiodic motions under small perturbations of the Hamiltonian.* Russ. Math. Surv. **18**, (5) 9–36.

[Arn63b] Arnold V I (1963): *Small divisor problems in classical and celestial mechanics.* Russ. Math. Surv. **18**, (6) 85–192.

[Arn78] Arnold V I (1978): *Mathematical Methods of Classical Mechanics.* Springer-Verlag.

[Arn83] Arnold V I (1983): *Geometrical methods in the theory of ordinary differential equations.* Springer-Verlag.

[Bar88] Barnsley M (1988): *Fractals Everywhere.* Academic Press, San Diego.

[Bea91] Beardon A F (1991): *Iteration of Rational Functions.* Springer-Verlag, New York.

[Bri91] Briggs K (1991): *A Precise Calculation of the Feigenbaum Constants.* Mathematics of Computation **195**, 435–439.

[Bri] Briggs K: *Complex Feigenbaum Functions.* Preprint.

[BS95] Burbanks A D and Stirnemann A (1995): *Hölder continuous Siegel disc boundary curves.* Nonlinearity **8**, 901–920.

[CG93] Carleson L and Gamelin T W (1993): *Complex dynamics.* Springer-Verlag.

[CC96] Celletti A and Chierchia L (1996): *On the stability of realistic three-body problems.* Preprint.

[CL94] Cline M P and Lomow G A (1994): `C++` *Frequently Asked Questions.* Addison Wesley.

[CE80] Collet P and Eckmann J-P (1980): *Iterated Maps on the Interval as Dynamical Systems.* Birkhäuser, Boston.

[CEK81] Collet P, Eckmann J-P, and Koch H (1981): *Period doubling bifurcations for families of maps on* $\mathbb{R}^N$. J. Stat. Phys. **25** (1), 1–14.

[Cre27] Cremer H (1927): *Zum Zentrumproblem.* Math. Ann. **98**, 151–163.

[CM83] Cvitanović P and Myrheim J (1983): *Universality for period n-tuplings in complex mappings.* Phys. Lett. **94A**, 329.

[Cvi89] Cvitanović P (1989): *Universality in Chaos.* 2nd ed. Adam Hilger.

[CM89] Cvitanović P and Myrheim J (1989): *Complex universality.* Commun. Math. Phys. **121**, 225.

[Die75] Dieudonné J (1975): *Grundzüge der modernen Analysis.* Band 1, Friedr. Vieweg u. Sohn, Braunschweig, zweite Auflage.

[DH82] Douady A and Hubbard J H (1982): *Iteration des Pôlynomes Quadratiques Complexes.* CRAS Paris **294**, 123–126.

[Dou86] Douady A (1986): *Julia Sets and the Mandelbrot Set.* In: *The Beauty of Fractals.* Peitgen H O and Richter P H. Springer-Verlag.

[EE86] Eckmann J-P and Epstein H (1986): *On the Existence of Fixed Points of the Composition Operator for Circle Maps.* Commun. Math. Phys. **107**, 213–231.

[EKW82] Eckmann J-P, Koch H, and Wittwer P (1982): *Existence of a Fixed Point of the Doubling Transformation for Area-Preserving Maps of the Plane.* Phys. Rev. A, **26**. 720–722.

[EKW84] Eckmann J-P, Koch H, and Wittwer P (1984): *A computer assisted proof of universality for area-preserving maps.* Mem. Am. Math. Soc. **47**.

[EW85] Eckmann J-P and Wittwer P (1985): *Computer methods and Borel summability applied to Feigenbaum's equation.* Lecture notes in physics, **227**. Springer-Verlag.

[EW87] Eckmann J-P and Wittwer P (1987): *A complete proof of the Feigenbaum conjectures.* Jour. Stat. Phys. **46**, 455–477.

[EMO91] Eckmann J-P, Malaspinas A, and Oliffson-Kamphorst S (1991): *A software tool for analysis in function spaces.* In: *Computer Aided Proofs in Analysis*, editors Meyer K and Schmidt D. The IMA Volumes in Mathematics, **28**, 147–166.

[EL81] Epstein H and Lascoux J (1981): *Analyticity Properties of the Feigenbaum Function.* Commun. Math. Phys. **81**, 437–453.

[Eps86] Epstein H (1986): *New Proofs of the Existence of the Feigenbaum Functions.* Commun. Math. Phys. **106**, 395-426.

[Fal90] Falconer K (1990): *Fractal Geometry, Mathematical Foundations and Applications.* Wiley, New York.

[Fat19] Fatou P (1919): *Sur les Equations Functionelles.* Bull. Soc. Math. Fr. **47**, 161–271 and **48** (1920), 33–94, 208–314.

[FL85] Fefferman C and de la Llave R (1985): *Relativistic stability of matter.* Preprint, Princeton University.

[Fei78] Feigenbaum M J (1978): *Quantitative Universality for a Class of Nonlinear Transformations.* J. Stat. Phys. **19**, 25–52.

[Fei79] Feigenbaum M J (1979): *The Universal Metric Properties of Nonlinear Transformations.* J. Stat. Phys. **21**, 669–706.

[FKS82] Feigenbaum M J, Kadanoff L P, and Shenker S J (1982): *Quasiperiodicity in dissipative systems: a renormalization group.* Physica **5D**, 370–386.

[Fei83] Feigenbaum M J (1983): *Universal Behaviour in Nonlinear Systems.* Physica **7D**, 16–39.

[GHJV94] Gamma E, Helm R, Johnson R, and Vlissides J — aka. "the gang of four" (1994): *Design Patterns: Elements of Reusable Object Oriented Software.* Addison-Wesley Professional Computing Series.

[GMVF81] Greene J M, MacKay R S, Vivaldi F, and Feigenbaum M (1981): *Universal behaviour in families of area-preserving maps* Physica **21D**, 468–486.

[GH83] Guckenheimer J and Holmes P (1983): *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* Springer-Verlag, New York.

[Hen93] Henrici P (1993): *Applied and Computational Complex Analysis.* Volumes I – III. Wiley Classics Library.

[Her79] Herman M R (1979): *Sur la Conjugaisons Differentiable des Diffeomorphismes du Cercle à des rotations.* Publ. I.H.E.S. **49**, 5.

[Her85] Herman M R (1985): *Are there critical points on the boundaries of singular domains?* Commun. Math. Phys. **99**, 593–612.

[Hof87] Hofstadter D R (1987): *Gödel, Escher, Bach — an Eternal Golden Braid.* Penguin.

[Hut81] Hutchinson J (1981): *Fractals and Self-similarity.* Indiana University Journal of Mathematics **30**, 713–747.

[IEEE85] IEEE Standard Board (1985): *The IEEE standard for binary floating-point arithmetic. ANSI/IEEE Std. 754-1985.* (IEEE Standard Board, 345 East 47th Street, New York, NY 10017, USA.)

[JBB83] Jensen M, Bak P, and Bohr T (1983): *Complete devil's staircase, fractal dimension and universality of the mode-locking structure on the circle map.* Phys. Rev. Lett. **50**, 1637–1639.

[Jul18] Julia G (1918): *Mémoire sur L'iteration des Fonctions Rationelles.* Journal de Math. Pure et Appl. **8**, 47–245.

[Knu81] Knuth D E (1981): *The Art of Computer Programming.* Vol 2, 2nd Edn. Addison-Wesley.

[KSW] Koch H, Schlenkel A, and Wittwer P: *Computer-Assisted Proofs in Analysis and Programming in Logic: A Case Study.* (Preprint)

[KW86] Koch H and Wittwer P (1986): *Rigorous Computer-Assisted Renormalization Group Analysis.* In: *VIIIth International Congress on Mathematical Physics*, editors Mebkhout M and Sénéor R. World Scientific.

[Kol54] Kolmogorov A N (1954): *On conservation of conditionally periodic motions under small perturbations of the Hamiltonian.* Dokl. Akad. Nauk. SSSR, **98**, 527–530.

[Lan82] Lanford III O E (1982): *A computer-assisted proof of the Feigenbaum conjectures.* Bull. Am. Math. Soc. **6**, 427 – 434.

[Lan86] Lanford III O E (1986): *Computer-assisted proofs in analysis.* Proc. Int. Congr. of Mathematicians, Berkeley, California.

[Lan91] Lanford III O E (1991): *Formal theory of fixed points. Preprint.*

[Man82] Mandelbrot B B (1982): *The Fractal Geometry of Nature.* W H Freeman and Co., New York.

[MN83] Manton N S and Nauenberg M (1983): *Universal scaling behaviour for iterated maps of the complex plane.* Commun. Math. Phys. **89**, 555–570.

[Mar90] Marmi S (1990): *Critical Functions for Complex Maps.* J. Phys. A: Math. Gen. **23**, 3447–3474.

[Mac82] MacKay R S (1982): *Renormalization in area-preserving maps.* PhD thesis, Princeton. (Ann Arbor, MI: Univ. Microfilms Int.)

[Mac83] MacKay R S (1983): *A renormalization approach to the breakup of invariant circles in area preserving maps.* Physica **7D**, 283–300.

[MP85] MacKay R S and Percival I C (1985): *Converse KAM: theory and practice.* Commun. Math. Phys. **98**, 469–512.

[Mar67] Markushevich A I (1967): *Theory of Functions of a Complex Variable.* Englewood Cliffs, NJ: Prentice Hall.

[Mes85] Mestel B D (1985): *A computer-assisted proof of universality for cubic critical maps of the circle with golden mean rotation numbers.* PhD thesis, Warwick University.

[MO89] Mestel B D and Osbaldestin A H (1989): *Renormalization in implicit complex maps.* Physica **39D**, 149–162.

[Mil91] Milnor J (1991): *Dynamics in one complex variable.* Institute for Mathematical Sciences, SUNY, Stony Brook NY, Preprint 9-5-91.

[Moo66] Moore R E (1966): *Interval analysis.* Prentice Hall, Series in automatic computation.

[Mos62] Moser J (1962): *On invariant curves of area-preserving mappings of an annulus.* Nachr. Akad. Wiss. Gottingen Math. Phys. Kl., **2**, 1–20.

[Mos73] Moser J (1973): *Stable and random motions in dynamical systems with special emphasis on celestial mechanics.* Ann. Math. Studies **77**. Princeton University Press.

[Oli89] Oliffson Kamphorst Leal da Silva S M (1989): *Scaling properties of the boundary of Siegel disks: a computer study.* Thèse 2317, Université de Genéve.

[ORSS82] Ostlund S, Rand D A, Sethna J, and Siggia E D (1982): *A universal transition from quasi-periodicity to chaos in dissipative systems.* Phys. Rev. Lett. **49**, 132.

[PM80] Pomeau Y and Manneville P (1980): *Intermittent Transition to Turbulence in Dissipative Dynamical Systems.* Commun. Math. Phys. **74**, 189–197.

[Ran87] Rand D A (1987): *Universality and Renormalization.* Review appearing in *New Directions in Dynamical Systems.* London Mathematical Society Lecture Note Series **127**, 1–56. Editors Bedford T and Swift J. Cambridge University Press.

[Rog95] Rogers J T (1995): *Critical Points on the Boundaries of Siegel Discs.* Bulletin of the American Mathematical Society, Volume 32, Number **3** (July 1995), 317–321.

[Rud80] Rudin W (1980): *Real and Complex Analysis.* McGraw-Hill.

[Rud82] Rudin W (1982): *Functional Analysis.* McGraw-Hill.

[Sch70] Schröder E (1870): *Über unendlichviele algorithmen zur auflösung der gleichungen.* Math. Ann. **2**, 317–365.

[Sch71] Schröder E (1871): *Über iterirte Funktionen.* Math. Ann. **3**, 296–322.

[She82] Shenker S J (1982): *Scaling behaviour in a map of the circle onto itself: empirical results.* Physica **5D**, 405–411.

[Sie42] Siegel C L (1942): *Iteration of analytic functions.* Ann. Math. **43**, 607–612.

[SM71] Siegel C L and Moser J K (1971): *Lectures on Celestial Mechanics.* Springer Classics in Mathematics.

[Sti92] Stirnemann A (1992): *Renormalization for golden circles.* Dissertation ETH No. 9843. Swiss Federal Institute of Technology, Zurich.

[Sti93a] Stirnemann A (1993): *A renormalization proof of Siegel's theorem.* Nonlinearity **7**, 943–958.

[Sti93b] Stirnemann A (1993): *Existence of the Siegel disc renormalization fixed point.* Nonlinearity **7**, 959–974.

[Sti93c] Stirnemann A (1993): *Renormalization for golden circles.* Commun. Math. Phys. **152**, 369–431.

[Sti] Stirnemann A: (Private Communications.)

[Str93] Stroustrup B (1993): *The* `C++` *Programming Language.* 2nd ed. Addison Wesley.

[Str94] Stroustrup B (1994): *The Design and Evolution of* `C++`. Addison Wesley.

[Sul85] Sullivan D (1985): *Quasiconformal Homeomorphisms and Dynamics I.* Ann. Math. **122**, 401–418.

[Lei90] Tan Lei (1990): *Similarity Between the Mandelbrot and Julia Sets.* Commun. Math. Phys. **134**, 587–617.

[Wid83] Widom M (1983): *Renormalization group analysis of quasi-periodicity in analytic maps.* Commun. Math. Phys. **92**, 121–136.

[Yoc] Yoccoz J-C: *Theoreme de Siegel, Polynomes Quadtratiques et Nombres be Brjuno.* (Preprint)

[Yoc88] Yoccoz J-C (1988): *Linéarisation des germes de difféomorphisms holomorphes de* $(\mathbb{C}, 0)$. C. R. Acad. Sci. Paris **306**, 55–58.

[Zyg59] Zygmund A (1959): *Trigonometric Series.* Vol I & II 2nd Edn. Cambridge University Press.