Computational Mechanics of Cellular Automata

James E. Hanson

Department of Physics
University of California
Berkeley, CA 94720

Computational Mechanics of Cellular Automata

Contents

Preface

This work grew out of five years of collaboration with my advisor, Jim Crutchfield. I am and shall remain deeply indebted to him for his insight, guidance and support.

The suggested order of reading is to start with chapter 1, then skip immediately to part II and continue with chapter 6, referring back to the intervening material as necessary. To the reader familiar with dynamical systems, cellular automata and computation theory, chapters 2–5 serve to define terms, fix notation, and introduce important technical tools used in the second part. Readers unfamiliar with any of the above topics are urged to consult the following books (or their equivalents) for a more comprehensive introduction:

- *Dynamical systems:* S. Wiggins, *Introduction to Applied Nonlinear Dynamical Systems and Chaos*.[102]
- *Cellular automata:* S. Wolfram, *Theory and Applications of Cellular Automata*.[107]
- *Computation theory:* J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation*.[53]

PART I
Introductory

1 Introduction

Nonlinear dynamical systems are ubiquitous in nature and notoriously difficult to analyze. The last thirty years have seen tremendous progress in the understanding of the complex behavior of *low-dimensional* nonlinear dynamical systems, in which the system's state can be characterized by a few variables. These developments were to a large extent motivated by questions about high-dimensional, spatially-extended systems. In this context we need only mention the pioneering work of E. N. Lorenz on atmospheric flow[72,73] and the notion of *strange attractor* proposed by Ruelle and Takens as a mechanism for fluid turbulence.[85] Out of the work on low-dimensional systems has emerged, among other things, the important and useful notion of *temporal chaos*, which refers to disordered or apparently random behavior arising from a simple deterministic equation of motion.[88]

In recent years, increasing attention has been paid to extending the knowledge gained from the study of low-dimensional systems to spatially-extended systems.[1,16,84,95] For example, one finds defect-induced turbulence, among other things, referred to as "spatio-temporal chaos."[15] But the differences between low-dimensional and spatially-extended systems are manifest. Most obviously, the latter have the fundamental property of forming *spatial patterns*. This adds a whole new dimension to the problem, both literally and figuratively.

Spatially extended dynamical systems in general are described by state spaces of infinite dimensionality. If, however, some emergent pattern or regularity in a system's configurations can be found, this pattern can be used to reduce the evolution of the system to some effectively low-dimensional—and therefore tractable—dynamics. Discovery and identification of emergent patterns is therefore a crucial first step in the analysis of the system. If no such patterns can be found, the analysis must fall back on an approximate stochastic model of the evolution—though even that depends for its success on some *statistical* regularity in the behavior.

One of the barriers to fuller understanding of pattern formation in spatially-extended systems has been the lack of a sufficiently general, precise, mathematical definition of "pattern" in the first place. One consequence of this lack is that the tools for discovering general patterns in space-time data are far from being well developed. Spatially periodic patterns can be detected by means of the spatial Fourier transform. But no analogous technique exists for finding patterns in aperiodic or disordered configurations; often their identification relies on visual inspection alone. Another consequence of the lack of a general definition of pattern is that the precise representation of the patterns exhibited by a system—representation in some analytically useful form—is largely confined to periodic patterns.*

This work represents an attempt to place the discovery, measurement, and analysis of patterns in spatially-extended systems on a more systematic footing. To get at this

---

*One alternative representation is afforded by the recently developed wavelet transform.[30,37] This promises to be extremely useful in discovering and representing patterns with localized, possibly fractal or self-similar structure.

difficult problem, we have stepped back to a class of simple idealized systems, cellular automata (CA). Concepts and methods developed for CA are therefore to be regarded partly as useful in their own right—to the extent that CA are "useful"—and partly as starting points for generalization to more "physical" systems.

Our primary focus will therefore be on *methods* of analysis. The various and detailed results discussed in later chapters are interesting in themselves, but their most important function is to demonstrate the validity and utility of the methods used to obtain them.

1.1 Patterns in Spatially-Extended Dynamical Systems

To set the context, we begin with a short discussion of two physical systems, describing some of the patterns of behavior encountered in experiments. This is followed by a brief introduction to cellular automata. Fuller treatment of CA is deferred to chapter 2.

The experimental systems were chosen because they are the subject of intensive and ongoing study, so their phenomenology is both well known and extensively detailed. The list of alternative systems that we could have discussed is a long and varied one. An excellent pictorial survey of the spectacular variety of patterns in different types of fluid flow can be found in ref. [34]. Other systems under study for their spatio-temporal patterns include interacting surface waves,[79] reaction-diffusion systems,[96,100] nematic liquid crystals,[9] ferromagnetic fluids,[6,86,92] and magnetic garnet films,[4,5] to name just a few.

Rayleigh-Bénard Convection

One of the most famous and well-studied systems in fluid dynamics the Rayleigh-Bénard system. This consists of thermally expansive fluid filling a region between two flat horizontal plates, which are usually rectangular or circular in shape. The linear size of the plates is denoted by L, the distance between them by d. A carefully controlled temperature difference is maintained across the plates, so that the top plate is held at temperature T_0 and the bottom at $T_0 + \delta T$ with $\delta T > 0$. This temperature difference is the control parameter of the system, expressed in dimensionless form as

$$\mu = \frac{\delta T}{T_0} \tag{1.1}$$

The fluid near the bottom plate heats up, expands, and is therefore more buoyant than the cold, dense fluid above it. If μ is sufficiently small, the viscosity of the fluid overwhelms the buoyant force and there is no macroscopic fluid motion. This is the stationary or heat conducting state of the system.

For μ slightly above a threshold value μ_1, however, the buoyancy is sufficiently great to induce convection, as the warm fluid rises and the cold fluid sinks to replace it. The stationary state destabilizes, and in time, the fluid motion organizes itself into

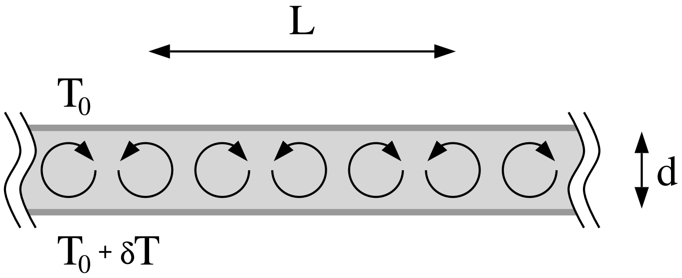

Fig. 1.1 Cross-section of part of the Rayleigh-Bénard system. Arrows indicate the direction of fluid flow. This system is in the spatially periodic roll regime $\mu_1 < \mu < \mu_2$.

spatially periodic roll patterns of a single (spatial) frequency. The velocity of fluid flow in this regime is time-independent. Fig. 1.1 shows a cross-section of part of the system in this regime.

The periodicity of the rolls means that representing the system in terms of its Fourier modes projects the infinite-dimensional state space down onto a space of two dimensions: the amplitude and phase of the rolls. The small nonlinearities in the evolution then show up as modulations in the amplitude, which are governed by an "envelope" equation.[75]

For large aspect ratios ($L \gg d$), in which the number of rolls is large, the system generically forms complicated patterns of rolls containing many "dislocations" and "grain boundaries." These are respectively points and lines at which the periodic roll pattern breaks down. The dislocations and grain boundaries move, interact, and annihilate on a timescale much longer than the timescale of the convection. They are *coherent structures*, and can be treated as quasiparticles with their own equations of motion.[10,75]

There is an orderly sequence of bifurcations as the control parameter μ is fixed at higher and higher values. At each bifurcation point μ_2, μ_3, μ_4, etc., the system undergoes a qualitative change in behavior. The bifurcation sequence is not to be considered a timeseries, but rather a sequence of separate experiments at different parameter values. At each value of μ, the system is allowed to evolve until transient behavior has died away before the asymptotic behavior is characterized. For μ in the region $\mu_1 < \mu < \mu_2$, the time-independence of the roll pattern breaks down as the boundaries of the rolls oscillate in a time-dependent periodic manner. Then for $\mu_3 < \mu < \mu_4$, the oscillation of the roll boundaries is governed by two independent frequencies. At successively higher values of μ, the temporal modulations of the roll pattern become more and more complicated.[97] It can also happen that the roll patterns become destabilized with respect to some other pattern, such as hexagonal cells. This is an example of *pattern selection*, in which two different patterns compete for dominance in the system.

At much higher temperature differences, the rolls completely disappear and fully-developed "hard" turbulence[50] is observed. This is characterized by a broad-band power spectrum similar to the Kolmogorov scaling law for turbulent flow. In the words of one author, “the flow seems totally unstructured.”[97] On the other hand, as refs. [50,109] and others have noted, coherent structures are common in the system; they have been described as “waves”, “spiraling swirls”, “mushrooms”, “thermals”, “plumes”, “jets”, and so on. Needless to say, these structures were identified purely visually; in some cases there is no mathematics to even describe them, let alone formulate their equations of motion.

Couette Flow

Another famous spatially-extended system, Couette flow, has been studied extensively for the rich variety of patterns it exhibits in different parameter regimes.[13] It is one of the classic systems in fluid dynamics. Like the Rayleigh-Bénard system, this system goes through an orderly sequence of bifurcations to more and more complicated patterns. The sequence is the subject of intensive and ongoing study. For our present purpose it is sufficient to pick out just a few of the many patterns this system can exhibit.

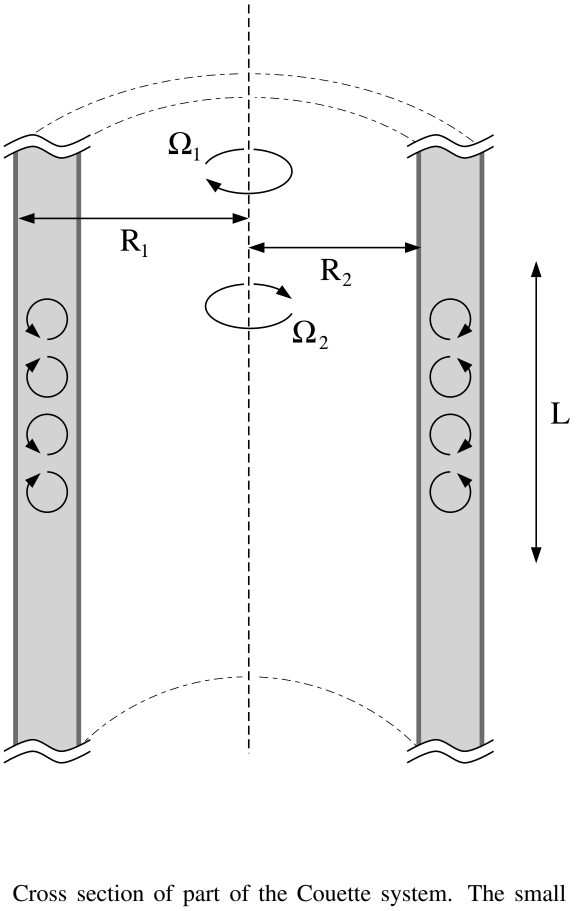

Fig. 1.2 Cross section of part of the Couette system. The small circular arrows indicate the motion of Taylor vortices. Only four vortices are shown.

Fig. 1.2 shows a diagram of the experimental setup for Couette flow. The system consists of a thin layer of fluid trapped between two concentric cylinders of radii R_1

and R_2. The cylinders are oriented so that their common axis is vertical. The system is sealed at the top and bottom. Each cylinder can be rotated on its axis independently at a range of speeds in either direction. The experimental control parameters are the two rotation rates Ω_1 and Ω_2 of the outer and inner cylinder, respectively.

At very low rotation rate Ω_1 with $\Omega_2 = 0$, the fluid flows uniformly around the cylinder. Small inhomogeneities in the flow decay over time until all of the fluid is moving in a simple, spatially-homogeneous laminar flow. This is the "Couette flow" regime from which the system as a whole takes its name.

When the rotation rate Ω_1 is increased, horizontal annular rolls or Taylor vortices appear. Just as for Rayleigh-Bénard system, in large aspect-ratio cylinders ($L \gg R_1 - R_2$) dislocations are present in the newly-formed roll pattern, which move and decay very slowly.[32] As Ω_1 is fixed at higher and higher values, the boundaries of the vortices oscillate, then more complicated structures appear. Finally, at high enough rotation rates, the flow becomes turbulent.

A surprisingly wide variety of different patterns appear at different values of Ω_1 and Ω_2. Ref. [3] gives a detailed map of different regimes observed in the flow at various rotation rates, including regimes for positive and negative Ω_2. The different regimes are given names expressive of their characteristic space-time patterns, and are listed in table 1.1. Many of the names given to the regimes contain in themselves essential information necessary for a proper mathematical representation. For example, "wavy interpenetrating spirals" would be parametrized by amplitude and frequency of the waves, size and rotation rate of the spirals, degree or geometry of interpenetration, and so forth. Other names contain references to existing mathematical representations—e.g., "Ekman vortices."

When the flow is turbulent, however, the differentiation among regimes is severely curtailed. Only four turbulent regimes are listed: intermittent turbulent spots, spiral turbulence, turbulent Taylor vortices, and featureless turbulence. Even among these, the differentiation is accomplished via the overall shape of the turbulent region's boundary—e.g., spiral turbulence. There is no identification of different patterns within turbulent flow, as is emphasized by the description of fully-developed turbulence as *featureless*. Yet the photograph of turbulent flow in ref. [3] shows anything but the flat uniformity of something "featureless." Quite the opposite, in fact: the image teems with a profusion of lines, curlicues and squiggles. The reason it was termed "featureless" was, presumably, that the authors were unable to find any regularity shared by all these features.

This again demonstrates how discovering patterns is the starting point and basis for all further analysis. Once a particular pattern—such as convection rolls or Taylor vortices—has been identified, it can be cast into mathematical form by the assumption of a canonical waveform. Much progress in understanding the different modes of motion has been made this way. If no pattern can be found, the analysis must content itself with stochastic models. Furthermore, for each new pattern, the analysis must

Abbreviation	Description
AZI	Azimuthal laminar flow with weak Ekman vortices
CKS	Corkscrew
INT	Intermittent turbulent spots
IPS	Interpenetrating spirals
MWV	Modulated wavy vortices
RIP	Ripple
SPI	Spiral vortices
SPT	Spiral turbulence
TRA	Transition region
TTV	Tubulent Taylor vortices
TUR	Turbulent flow (featureless)
TVF	Taylor vortex flow
TWI	Twisted vortices
WIB	Wavy inflow boundary
WIS	Wavy interpenetrating spirals
WOB	Wavy outflow boundary
WVF	Wavy vortex flow
WVL	Wavelets

Table 1.1 Principal regimes observed in flow between independently rotating cylinders (transcribed from ref. [3]). The order is alphabetical.

start over from the beginning. There is no overall theoretical apparatus for identifying or analyzing the general patterns of the system.

Cellular Automata

A similar diversity of behavior appears in a much simpler class of spatial system, one-dimensional cellular automata (CA).[107] CA are one class of fully discrete spatially-extended dynamical systems. Space and time are indexed by integers, and the value of each spatial site is taken from some finite alphabet of formal symbols. In CA, the spatial coupling between sites is local, and the site-update rule is expressed as a spatially and temporally invariant lookup table.† A space-time plot of the evolution of a randomly chosen initial condition under a particular CA is shown in fig. 1.3, with cell values printed as integers. The spatial region shown is a part of a much larger lattice. The values of the sites in the "parent neighborhood" shown in gray in the top

---

†Other types of fully discrete spatially-extended systems have been investigated. One such system, closely related to CA but more powerful, is defined in ref. [23].

row of the figure determine the value of the “child cell” shown in gray in the second row. The CA shown has nearest-neighbor coupling.

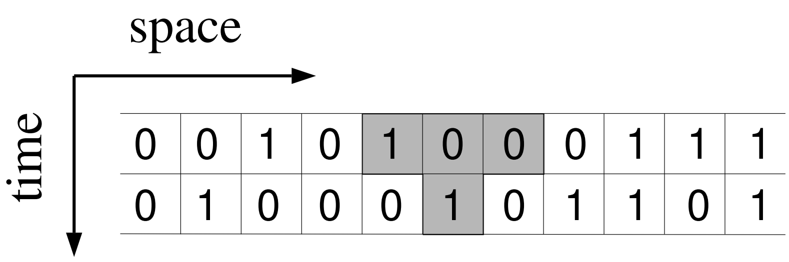

Fig. 1.3 Part of a space-time plot of a cellular automaton. At each time t, the state s_t is displayed horizontally; evolution is down the page. A parent neighborhood and its child cell (for a nearest-neighbor CA rule) are shown in gray.

Despite their extreme simplicity, CA exhibit an extremely wide range of behavior, from spatially uniform stationary states to spatially and temporally periodic evolution to spatial disorder to fully-developed, spatio-temporally “turbulent” behavior. As we will see, many CA evolve to patterns which can be decomposed into two components: (i) a spatio-temporally uniform, periodic, or turbulent background, and (ii) “particles” or “dislocations” moving against the background. In the same way as for the other systems above, the particles are defined *in terms of* the background pattern: they are points at which the pattern breaks down.

CA are the dynamical systems that will be investigated in part II of this work. They were chosen because they are at once simple systems and systems exhibiting the full range of qualitative behavior found in nature.

1.2 Pattern Bases and Coherent Structures

In the preceding examples we have indicated the common methodology in analyzing nonlinear spatially-extended systems: first the characteristic pattern(s) into which the system evolves are identified, then a representation that fits these patterns is formulated, and finally the characteristic patterns are used as the underlying structure in terms of which the system is understood. Another way of saying this is that the characteristic patterns emerging in a system’s evolution comprise a *pattern basis* that governs the system’s behavior and guides its analysis. We have also seen that this methodology breaks down in the face of spatio-temporal disorder, by failing at the very first step: no characteristic patterns can be identified.

Coherent structures are points in the system’s configuration that do not conform to the pattern basis: defects, singularities, walls, etc. The hierarchy is important here: *first* we identify the pattern bases, *then* we analyze the coherent structures in terms of the pattern bases. This definition of coherent structure is somewhat different from another one in common use in fluid dynamics. According to that definition, a coherent structure is some persistent structural feature that lasts longer than the typical timescale for dissipation or diffusion. But persistence is a purely temporal phenomenon, and therefore should not be the sole defining property of a coherent structure. Coherent structures are more properly defined in terms of spatial patterns. As we will see, pattern bases—and therefore coherent structures—are identified according to both spatial and temporal properties.

Expansion in Fourier modes, the common choice of pattern basis for continuum systems, implicitly assumes that harmonic waves are appropriate for representing the system's spatial configurations. In many instances, this assumption is justified. For example, the periodic rolls of Rayleigh-Bénard convection are the pattern basis for that system, and the dislocations and grain boundaries are the coherent structures defined by breakdowns in the roll patterns. The Fourier basis is less appropriate for describing the coherent structures themselves, however, since any spatially localized structure necessarily has an infinite number of Fourier coefficients. Thus for example an isolated soliton or dislocation, as represented in the Fourier basis, is far from being a single "entity"; rather, it is described by an infinite number of Fourier modes.

If the system is turbulent, a further difficulty arises. The diversity of observed configurations is so great that one is forced into the position of disregarding the details of the evolution in order to adopt some sort of stochastic model, such as in the case of fully-developed "hard" turbulence. The problem with these models is that they virtually ignore the possibility of *any* regularity existing in the system, effectively masking it with purely stochastic dynamics. It is by now well-appreciated that for chaotic systems, this modeling procedure throws away important structural information; in particular, it fails to capture the possibility that coherent structures may be present in the system. What is needed instead is a *turbulent pattern basis*: a way of representing spatial configurations that captures the structural properties of the dynamics—e.g. via an appropriate notion of pattern—but that still allows for sufficient stochasticity that it is not burdened with describing the unpredictable details of high-entropy-rate "turbulent" behavior. In this way, coherent structures, if they exist, would show up as points at which the *turbulent* pattern breaks down. The system's evolution can then be decomposed into a turbulent background and coherent structures evolving against it. Recent work on two-dimensional vortex turbulence has shown how effective this type of decomposition can be.[11,12,54]

In developing a general idea of pattern bases, we are immediately presented with a fundamental question: what is a pattern? The basic notion is that the members of a particular set of configurations are similar in some way, in that they share some common structural feature. The *pattern* is that which they have in common. Conversely, the configurations are the *realizations* of the pattern. Note that this definition of pattern does not admit of choosing arbitrary sets of configurations: there must be some particular feature, however insignificant, shared by all members of the set.

One way of implementing this notion of pattern is to use the “structural feature” as a rule for constructing examples of the pattern: i.e., an algorithm. This is pattern *generation.* Another implementation builds the structure into a rule for classifying existing configurations by deciding whether a given configuration is an example of a given pattern. This is pattern *recognition.* Both of these are made possible only by a prior step, however: pattern *discovery.*

Intuitively, then, a pattern is a structural feature or collection of features. There are many different senses of the phrase “structural feature,” each one of which identifies a different kind of pattern. In part II of this work, we will explore the application of a powerful and extensive class of patterns to the spatial configurations generated by cellular automata.

1.3 Qualitative Pattern Dynamics

The theory of qualitative dynamics, developed primarily for low-dimensional systems, attempts to describe the system in terms of important geometric features of state space that govern the system’s temporal behavior.

Applying the qualitative approach to spatially-extended systems such as fluid flow has proved difficult, however. The traditional approach has been to treat the system as a microscopically continuous medium, then to associate a degree of freedom with each spatial point. The result is an uncountably infinite dimensional system—a function space. This is the approach taken in deriving partial differential equations (PDEs), for example. In many cases, however, *deriving* a PDE is a long way from *solving* it, or from determining the qualitative behavior of the flow it generates. The geometry of orbits in infinite-dimensional spaces is sufficiently complex that at present for PDEs, “the beginning of a qualitative theory is still lacking.”[15] There do not even exist general theorems for the existence and uniqueness of the solutions of an arbitrary PDE.[97]

Further, treating each point in space as a separate dimension obscures the essential fact of spatial extension—i.e., that some points are closer, some are farther, from any given point, in various directions. In the infinite-dimensional space, the dimension representing each point is “equally” orthogonal to every other dimension. A coherent structure, far from being identified as a single entity, is a sequence of “active” subspaces of the state space. As the coherent structure changes shape and location, the particular dimensions that make up the active subspaces are in constant flux.

We have seen that for spatially-extended systems, the *spatial* patterns play a fundamental organizing role in the system’s evolution. Therefore in analyzing the qualitative dynamics of a spatially-extended system, it is the temporal behavior of the spatial patterns that we want to investigate. This includes discovering which patterns emerge during the system’s evolution, which patterns are reached asymptotically, what sort of patterns are stable to perturbations, etc.

The name for this type of analysis is *qualitative pattern dynamics*. The key element, on which all subsequent development depends, is the proper representation for the patterns—i.e., the pattern bases. In many cases, the pattern bases allow the system to be projected down onto a space of much lower dimension. The spatial features are built into the analysis from the start.

1.4 Computational Mechanics

Computational mechanics, as the term is used here, was introduced in refs. [19,21,22] as the name for the investigation of the *computational structure* of dynamical systems. Computational structure refers to the use of models and techniques developed in theoretical computer science: i.e., automata and formal languages. This is entirely different from the usage of "computational physics" to mean the development of effective techniques for performing numerical calculations of physical quantities.

One connection between computation and dynamics rests on the fact that any idealized computer is itself a deterministic dynamical system. This particularly clear in the case of analog computers, in which the "computation" is performed by the evolution of a system of interacting operational amplifiers, multipliers, and so forth. Computational mechanics takes this connection and adds to it the reverse analogy: it investigates the ways in which a given dynamical system is in effect performing a computation. Besides providing a new kind of information about dynamical systems, general methods of making this connection may lead to applications in the design of new practical methods of performing computations.

Computational mechanics seeks to discover and quantify the *intrinsic* computation in a dynamical system. Intrinsic computation refers to the computational properties that are generic or typical of the system's evolution. This is distinct from the "designed-in" or "best-case" computational capacity of a system, which describes the types of computation that the system can in principle be set up to perform.

Computational mechanics is the next step beyond the use of information theoretic ideas in dynamics, which are included under the general heading of *statistical* mechanics. Information theory has been used with great success to characterize the diversity of behavior of a system, its instability, the rate at which predictability of or information about the system's state decays.[90] However, the notion of "information" is insufficient to help identify or describe structural features or patterns in the system's evolution. This lack is precisely what computational mechanics is designed to fill.

When applied to spatially-extended systems such as CA, computational mechanics, addresses the intrinsic *parallel* computational properties of the system. It effectively treats the system as a parallel computer, for example by associating the evolution of spatial patterns with the performance of a parallel computation. It can therefore be regarded as contributing to the foundations of a theory of parallel computation in physical systems.

By far, the best-developed mathematics of patterns is the theory of formal languages and automata. These are defined over the domain of discrete data, i.e., sequences of formal symbols drawn from a finite set. To use this framework, a system's spatial configurations must be converted into a sequence of discrete symbols at the outset. In practice, this conversion is implicit in the very act of taking measurements. Since any measuring instrument has a finite level of resolution, experimental data is inevitably in the form of a sequence of finite-precision—i.e., discrete—values. In principle, the conversion to symbolic data is done using the techniques of symbolic dynamics.[2] A *generating partition* is employed to make the conversion. A generating partition is a discretization of a real-valued state variable—i.e., a way of partitioning the variable's range into discrete bins—in such a way that there is a finite-to-one correspondence between the possible values of the state variable at a given time and the possible infinite-length time sequences of discrete or symbolic measurements at that point.

In carrying out the program of computational mechanics for cellular automata, the conversion to discrete symbols is unnecessary, since the system's state variables are already discrete. This is another reason why CA are an attractive developing-ground for the theory of computational mechanics in spatially-extended systems generally.

2 Cellular Automata: Definition and Survey of Related Work

As we saw in chapter 1, cellular automata (CA) are spatially-extended dynamical systems with the particular feature of being as discrete as possible: in space, in time, and in site value. The rule governing a CA's evolution is extremely simple, being a spatially and temporally invariant, locally coupled, parallel update of the lattice. Yet CA exhibit the same wide range of behavior, and many of the same qualitative features, as considerably more complicated "physical" dynamical systems.

Because they are so discrete, CA are ideally suited to fast simulation on digital computers. The use of CA as a simulation tool for partial differential equations (PDEs), such as those of fluid dynamics, is a vigorous and growing field. These systems have two or three spatial dimensions, usually on a hexagonal lattice, in which the cell values are used to store information about the location and velocities of particles obeying specific microscopic interaction rules. They usually go by the name of "CA fluids" or "lattice gases."[33,36]

On a deeper level, they have been presented as a fundamental *alternative* to PDEs in the analysis of spatially-extended physical systems.[98,99] In deriving a PDE for some physical system, one makes the essential simplifying assumption that the system can be treated as being microscopically continuous without seriously impairing the validity of the results. This is clearly not the case for real systems, yet PDEs are remarkably successful in aiding the analysis of a wide variety of systems. In modeling the system as a CA, one makes an alternative assumption, namely, that at some scale the system can be treated as being completely discrete.

A third point of view, which will be taken here, is that CA provide a computationally and analytically simple developing ground for general-purpose ideas. Spatially-extended systems are inherently difficult to analyze because the state space is infinite dimensional. The discreteness of CA is expected to make the analysis simpler. The ideas developed in this simple setting would then be used as a guide in generalizing to more continuous systems such as coupled map lattices and (eventually) to fully continuous systems such as fluid flow. One possible route for this generalization uses a generating partition to discretize the state space, as discussed in section 1.4.

We will be considering CA in one spatial dimension only. The theoretical framework for higher-dimensional CA presents specific difficulties and is considerably less well developed. One of the most immediate problems with higher dimensional CA is the fact that a state can no longer be treated as a *sequence* of cells. This fact severely complicates the connection between CA and formal language theory that, as we will see in part II, is currently one of the cornerstones of the computational mechanics of CA. The extension to CA of higher spatial dimension is a subject of ongoing research.

2.1 Definitions

Cellular Automata

A **cellular automaton** consists of a one-dimensional countable array of discrete sites or **cells** i and a discrete-time **local update rule** ϕ operating in parallel on local neighborhoods of a given radius r. At each time t each site takes on a value σ_t^i in a finite alphabet $\mathcal{A}$ of primitive symbols: $\sigma_t^i \in \{0, 1, ...k-1\} \equiv \mathcal{A}$. The local site-update operation is written

$$\sigma_{t+1}^i = \phi(\sigma_t^{i-r}, \ldots, \sigma_t^{i+r}). \tag{2.1}$$

The space index i and the time index t are omitted when they are unimportant.

The string $\eta_t^i = \sigma_t^{i-r} ... \sigma_t^{i+r}$ of $2r+1$ symbols is the **parent neighborhood** of σ_{t+1}^i, its **child symbol**. The local update rule is therefore also written $\sigma_{t+1}^i = \phi(\eta_t^i)$. The set of distinct ordered pairs (η, σ) with $\sigma = \phi(\eta)$ is the **rule table**, of which the η's are the **input portion** and the σ's are the **output portion**.

The **state** (or configuration) s_t of the CA at time t is the configuration of the finite or infinite spatial array: $\mathrm{s}_t \in \mathcal{A}^N$, where $\mathcal{A}^N$ is the set of all possible site values on an array of N cells. The array is also called the CA's **lattice**.

The **extended state space**, denoted by $\mathcal{A}^*$, is the union of all states of any N:

$$\mathcal{A}^* = \bigcup_{N \geq 0} \mathcal{A}^N \tag{2.2}$$

with $\mathcal{A}^0 = \emptyset$.

The CA **global update rule** $\Phi : \mathcal{A}^N \to \mathcal{A}^N$ applies ϕ in parallel to all sites in the lattice: $\mathrm{s}_t = \Phi \mathrm{s}_{t-1}$. For finite N it is also necessary to specify a boundary condition. In the following, we will always use periodic boundary conditions.

The simplest class of cellular automata are the "**elementary**" CA (ECA) for which $(k, r) = (2, 1)$; that is, nearest-neighbor interactions and binary cell values, giving a total of $k^{2r+1} = 8$ possible parent neighborhoods. Since for each such parent neighborhood there are k possible values of the child cell, there is a total of $k^{k^{2r+1}}$ different local update rules ϕ. Thus there are $2^8 = 256$ different elementary CA rules. For the class of binary CA with the next-larger radius, binary range-2 CA (BR2CA) in which $(k, r) = (2, 2)$, there are 32 different parent neighborhoods, hence 2^{32} different rules.

A numerical index $\#(\phi) \in \mathbb{Z}^+$, the **rule number**, is assigned to each function ϕ (and therefore to each Φ), using a conventional numbering scheme as follows.[103] The different parent neighborhoods η are regarded as numbers in base k, and arranged in decreasing numerical order, say from left to right. Immediately beneath each parent neighborhood, its child symbol $\sigma = \phi(\eta)$ is written. The sequence of $2r+1$ child symbols is then regarded as a number, again in base k. The ECA having rule index 18, for example, has the following rule table:

Neighborhood	$\underline{111}$	$\underline{110}$	$\underline{101}$	$\underline{100}$	$\underline{011}$	$\underline{010}$	$\underline{001}$	$\underline{000}$	$\underline{\sigma_t^{i-1}\sigma_t^i \quad \sigma_t^{i+1}}$
Next site value	0	0	0	1	0	0	1	0	σ_{t+1}^i

(2.3)

In some CA applications, it is useful to designate a particular symbol $\hat{\sigma} \in \mathcal{A}$ (usually 0) as a "**quiescent**" symbol, and consider other members of $\mathcal{A}$ as somehow denoting "excited" cell values. In lattice gases, for example, a cell with value $\sigma = 0$ is an "empty" cell that does not contain any gas particles. Other applications interpret the value of a cell as the velocity of fluid flow at that point; then $\sigma = 0$ denotes a singularity in the velocity field. It is usual to require that the CA rule leave blocks of quiescent cells invariant; i.e.,

$$\phi(\hat{\sigma}, \hat{\sigma}, ..., \hat{\sigma}) = \hat{\sigma}. \tag{2.4}$$

Ensemble evolution

As defined above, the CA update rule Φ operates on individual configurations. Alternatively, we can look at how a CA processes entire sets of spatial configurations. This is analogous in dynamical systems theory to the study of the evolution of state ensembles, rather than individual orbits. To this end we define the ensemble evolution operator $\mathbf{\Phi}$ by

$$\mathcal{L}_t = \mathbf{\Phi}\mathcal{L}_{t-1} = \mathbf{\Phi}^t \mathcal{L}_0 \tag{2.5}$$

where $\mathcal{L}_0 \subseteq \mathcal{A}^*$ is any set of spatial configurations used as initial conditions, and the ensemble evolves according to

$$\mathcal{L}_t = \{\mathrm{s}_t \,:\, \mathrm{s}_t = \Phi \mathrm{s}_{t-1},\ \forall\, \mathrm{s}_{t-1} \in \mathcal{L}_{t-1}\} \tag{2.6}$$

Note that no restrictions are made on the size of the lattice; a single ensemble may contain states with different lattice sizes.

Many CA are irreversible or dissipative in the sense that the global update rule is many-to-one, in which case $\mathcal{L}_t \subset \mathcal{L}_{t-1}$. As we will see, irreversibility is a typical feature of CA exhibiting "chaotic" behavior (which will be defined), though reversible CA can also behave in this way.

Sparse sets

It will turn out that for finite lattices, there are a few values of N for which the evolution of the CA is radically different from its usual behavior on other lattice sizes. Since we are interested in the *typical* behavior of a given CA, we define the notion of a **sparse set** of lattice sizes, on which the behavior is allowed to be atypical, but which are themselves so rare as to be insignificant. This allows us to investigate the qualitative dynamics of the CA without worrying about complications introduced by rare number-theoretic fluctuations in the behavior. Such fluctuations are an interesting subject in themselves, and will be touched on in chapter 13.

Definition. Given a set $R = \{N_1, N_2, ...\}$ of distinct CA lattice sizes ordered so that $N_i < N_j$ for $i < j$, let $I_R(n)$ be the number of elements $N_i \in R$ such that $N_i \leq n$. Then R is said to be **sparse** if and only if

$$\lim_{n \to \infty} \frac{I_R(n)}{n} = 0 \,. \tag{2.7}$$

If a property P holds for all but a sparse set of lattice sizes, we will say that P holds for "**almost all**" lattices, or equivalently, for "**typical**" lattices.

2.2 Phenomenology

As even the most cursory inspection shows, CA exhibit a remarkable variety in behavior. Space-time plots of different CA range from being entirely ordered to spatio-temporally disordered. At least one CA has been shown to be usable as a random number generator.[61,106]

Nevertheless, after an extensive survey, Wolfram[105] proposed that CA rules can be classified in four qualitative classes, based on a visual examination of space-time patterns exhibited by the CA at long times:

1. Spatio-temporally uniform state.
2. Separated simple or periodic structures.
3. Chaotic space-time pattern.
4. Complex localized structures.

An example of each of these classes is shown in fig. 2.1. It is important to emphasize that this classification scheme is phenomenological at root. The scheme was originally motivated by the hierarchy of behavior in dynamical systems theory (see section 3.2)—i.e., the sequence: fixed point, limit cycle, quasiperiodicity, chaos.

Classes 1 and 2 are more or less self-explanatory. Note, however, that temporally periodic evolution of a spatially disordered configuration ("spatial chaos") falls into class 2. In class 3, "chaotic" is taken to mean "spatio-temporal disorder", in the sense that the number of distinct space-time patches grows exponentially with the linear size of the patch along both the spatial and temporal axes. Class 4 has been the subject of speculation about its possible computational capacity. It has been suggested that CA in this class may be capable of universal computation, that is, of implementing a universal Turing machine.[105] Class 4 was defined by the presence of spatially isolated transients of many shapes evolving for arbitrarily long times. The conjecture was that using these propagating "blobs" a model of universal serial computation could be constructed. Examples of one-dimensional CA implementing universal Turing

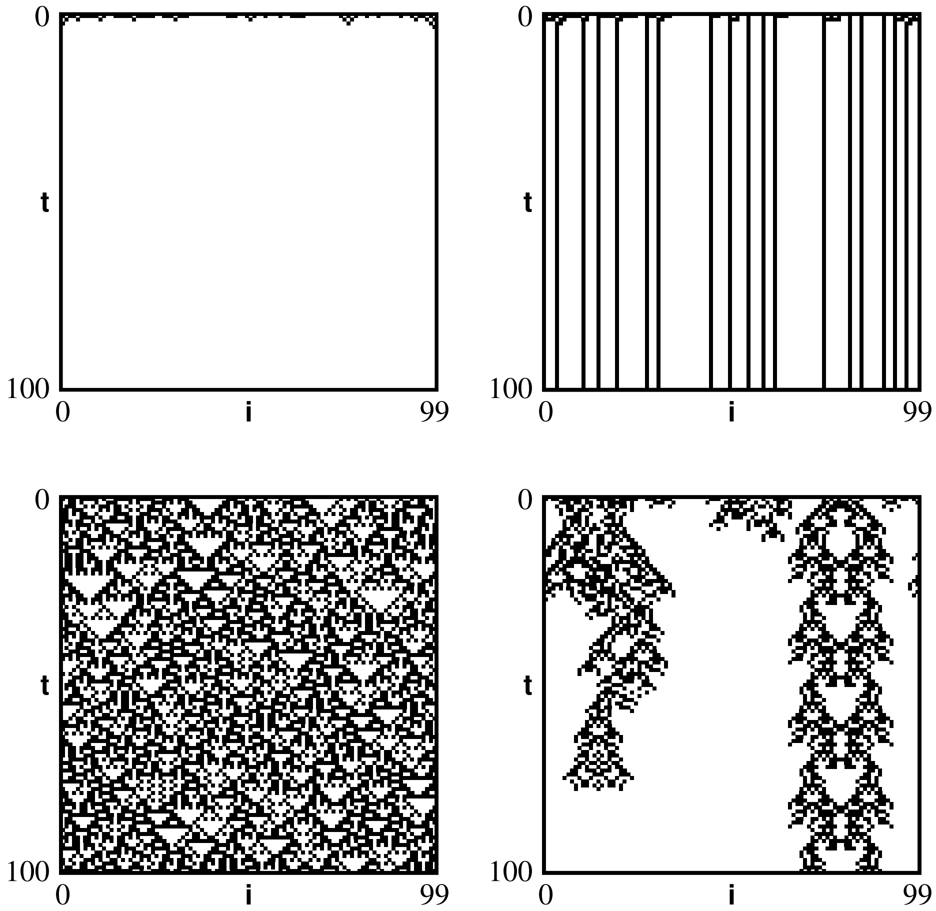

Fig. 2.1 Examples of Wolfram's four qualitative classes. Top to bottom, left to right, they are: (a) Class 1: ECA 32; (b) Class 2: ECA 44; (c) Class 3: ECA 90; (d) Class 4: Binary range-2 $(k = 2, r = 2)$ CA 1771476584.

machines have in fact been found.[93] These are in class 4 by construction, since the computational capability was built in by means of propagating space-time structures; in effect, they use the CA lattice as a substrate for serial computation.

Further work on the classification of CA developed the "lambda parameter"[69] as a diagnostic tool. The lambda parameter is defined for CA rules possessing a "quiescent" cell value, and is given by the number of non-quiescent symbols in the output portion of the rule table. It has been proposed as a heuristic to help in finding CA rules in one or another of the four Wolfram classes. A refinement of this approach is the hierarchical mean-field analysis,[45] which augments the lambda parameter with a series of more precise statistics that are essentially the number of non-quiescent symbols in the output portions of various subsets of the rule table.

2.3 Transients and Cycles on Small Lattices

Because the state space is discrete, on a finite lattice of any given size the number of possible configurations is necessarily finite. This means that, in principle at least, all possible configurations of the system can be enumerated (say, as points on a sheet of paper), and the step-by-step evolution can be represented by connecting those points with directed segments. Once this has been done for each state in the state space, the resulting graph is a "state transition diagram," introduced in ref. [76]. The state transition diagram is a directed graph, often consisting of multiple disconnected components, with each node represents a unique configuration. Because the CA rule is deterministic, there is exactly one edge leaving each node. The majority of CA rules are dissipative, which shows up as multiple edges entering nodes. The branching in the trees is a measure of the amount of dissipation in the system. The so-called "Garden-of-Eden" configurations, which have no predecessors, have in-degree 0. After a possible "transient", a CA evolving from any initial condition must ultimately enter a cycle.

The state transition diagrams for all of the elementary CA and a number of non-elementary CA, on a variety of small lattices, have been mapped in ref. [108]. The cycles are called "attractors" because they represent the set of configurations asymptotically reached by the system. Each component of the graph has a single attractor, and is therefore called the "basin" of that attractor. The entire state transition diagram, consisting of all attractors and their basins, is the "basin-of-attraction field."

The practical difficulty in constructing the state transition diagram is that the number of configurations grows exponentially with lattice size as

$$\left\| \mathcal{A}^N \right\| = k^N \tag{2.8}$$

so that a complete enumeration can only be carried out on very small lattices. Typical examples studied in the literature have dealt with lattices of size on the order of $N \sim 10^1$ sites. This is to be compared with (two dimensional) lattices of order $N \sim 10^4 - 10^6$ commonly used in fluid dynamics simulations.[33,36]

As has long been recognized, the shape of the state transition diagram is extremely sensitive to the number-theoretic properties of the lattice size.[46,66] The state transition diagram of a CA with a given lattice size does not generalize to other lattice sizes. The algebraic analysis, outlined in the next section, helps to explicate this fact.

2.4 Algebraic methods

The algebraic analysis of CA [63–65,76] considers finite CA with periodic boundary conditions. The possible site values are taken to be elements of a finite commutative ring $\mathcal{R}_k$ with k elements. Typically $\mathcal{R}_k$ is chosen to be $\mathbb{Z}_k$, so that site values are treated as integers modulo k.

On a lattice of size N, each configuration $\mathrm{s}_t \in \mathcal{A}^N$ is represented by a characteristic polynomial

$$A^{(t)}(x) = \sum_{i=0}^{N-1} a_i^{(t)} x^i \tag{2.9}$$

where $x \in \mathbb{R}$ is a formal parameter and $a_i^{(t)}$ is the numeric value of cell σ_t^i.

The local update rule $\phi\big(\sigma_t^{i-r}, ..., \sigma_t^{i+r}\big)$ is then expressed as a function

$$a_i^{(t)} = F\Big(a_{i-r}^{(t-1)}, ..., a_{i+r}^{(t-1)}\Big) \tag{2.10}$$

"Additive" (or "linear") CA are those in which the function F has the form of a simple linear combination of its arguments, i.e.,

$$a_i^{(t)} = \alpha_{-r} a_{i-r}^{(t-1)} \; + \; \alpha_{-r+1} a_{i-r+1}^{(t-1)} \; + \; ... \; + \; \alpha_{+r} a_{i+r}^{(t-1)} \tag{2.11}$$

where each coefficient α_j is a fixed element of $\mathcal{R}_k$ and all arithmetic is performed in $\mathcal{R}_k$. For such CA, the global evolution $\Phi(\mathrm{s}_t)$ can be represented by the multiplication of the characteristic polynomial $A^{(t)}(x)$ by the fixed "dipolynomial" (i.e., a polynomial with both positive and negative powers of x)

$$\mathbb{T}(x) = \alpha_{-r} x^{-r} + \alpha_{-r+1} x^{-r+1} + ... + \alpha_r x^r \tag{2.12}$$

according to the relation

$$A^{(t)}(x) = \mathbb{T}(x) A^{(t-1)}(x) \quad \mathrm{mod}\Big(x^N - 1\Big) \tag{2.13}$$

In this way, the evolution of an additive CA is expressed in the language of linear algebra. Additive CA obey a linear superposition principle in the sense that the site-by-site sum (mod k) of two configurations evolving under the CA rule is itself a configuration evolving according to the same rule. Thus for any two states $A^{(0)}(x)$ and $B^{(0)}(x)$ (expressed by their characteristic polynomials), the sum $C^{(0)}(x) = A^{(0)}(x) + B^{(0)}(x)$ evolves such that $C^{(t)}(x) = A^{(t)}(x) + B^{(t)}(x)$ for each time t. In particular, this means that the evolution of an arbitrary initial condition s_0 can be expressed as a superposition of the evolution from "single-seed" initial conditions, each consisting of a single nonzero cell.

One of the most extensively studied additive CA is ECA 90, for which the evolution is governed by the dipolynomial

$$\mathbb{T}(x) = x + x^{-1} \tag{2.14}$$

For example, an initial condition $\mathrm{s}_0 =$ ...00100... containing a single nonzero cell evolves after t time steps to a state with characteristic polynomial

$$\mathbb{T}(x)1 = \big(x + x^{-1}\big)^t = \sum_{i=0}^{t} \binom{t}{i} x^{2i-t} \tag{2.15}$$

Note that this is the formula for Pascal's triangle modulo 2. Fig. 2.2 shows a space-time plot of the evolution of ECA 90 on a periodic lattice of size $N = 1000$, starting from a "single-seed" initial condition with a single nonzero cell at site $i = 500$. Only the central portion of the lattice is shown. Because ECA 90 is linear, the evolution of an arbitrary initial condition is given by the site-by-site sum (mod 2) of copies

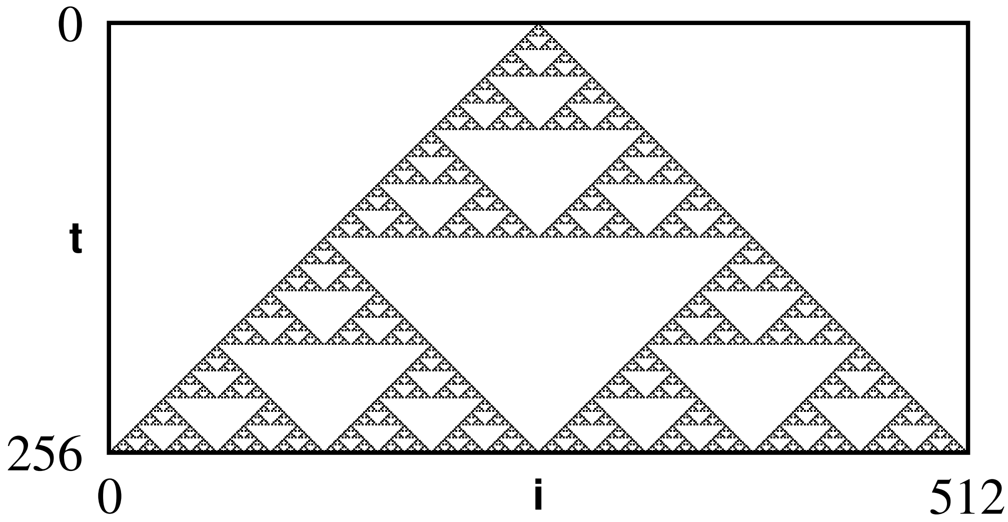

Fig. 2.2 Space-time plot of ECA 90 evolving from an initial condition consisting of a single nonzero site. The space-time pattern is just Pascal's triangle modulo 2.

of the evolution of this initial condition, shifted to the left or right as appropriate. Random initial conditions evolving under ECA 90 exhibit spatio-temporal disorder or "chaos", despite the linearity of the rule. The disorder arises from superposing many shifted copies of the space-time pattern of fig.2.2. Ref. [76] contains extensive calculations of the number and lengths of transients and cycles for ECA 90 on finite periodic lattices of size N.

For nonadditive ("nonlinear") CA, the algebraic analysis is considerably more difficult. Jen[63–65] has developed a "linearization" technique that can help in a number of cases. In essence, linearizing a CA consists of transforming both the state s_t and the rule Φ in such a way that the transformed system consists of a new state $\mathrm{s}_t{}'$ evolving under an *additive* rule Φ'. Such a transformation must be found for each nonadditive rule. The most extensively studied case is ECA 18, which can be linearized to ECA 90.

2.5 Global Evolution

The global evolution of a CA is given by the time-sequence of iterates Ω_t of the set of all states in the extended state space $\Omega_0 = \mathcal{A}^*$. Wolfram [104] introduced a method of tracing the global evolution of a CA that avoids enumeration of that infinite set. The calculation method uses computation-theoretic ideas introduced in chapter 4, so a full treatment here would be premature. Very briefly, the method is used to construct the finite automaton (or *machine*) $\mathrm{M}(\Omega_t)$ representing the set of all such iterates at each successive time step. In this section, we will consider only one of the properties of Ω_t, the *size* V_t of $\mathrm{M}(\Omega_t)$.* Roughly speaking, this can be regarded

---

*In ref. [107], this is called the *regular language complexity*. We will reserve the name *complexity* for a different quantity, defined in chapter 4.

as the amount of memory required to describe the state space. (As a baseline for comparison, the size of $\mathrm{M}(\Omega_0)$, with $\Omega_0 = \mathcal{A}^*$, is $V_0 = 1$.) In later chapters, we will develop these ideas more precisely.

Unfortunately, for many CA with nontrivial behavior (and especially the most interesting and complex ones), V_t grows extremely rapidly with time, to the extent that in some cases the calculation has to be suspended after as few as three iterations because of practical limitations on computational resources. For the nonlinear, chaotic ECA 18, for example, $V_2 = 47$ and $V_3 = 143$; V_4 was estimated to be greater than 2×10^4.[55]

Tracking the evolution of V_t can be problematic in other ways, as well. For example, in ECA 44, the "class 2" CA of fig. 2.1, V_t grows in a regular, "stair-step" fashion, approximately linearly, apparently without bound. This is not at all evident from the space-time plot of fig. 2.1, which indicates that after a relatively short time, the configuration s_t settles down to an invariant pattern. Fig. 2.3 shows a plot of V_t vs. t for ECA 44 in the range $0 \leq t \leq 100$.* This discrepancy between the global evolution and the space-time phenomenology is one of the central motivating factors of the present work.

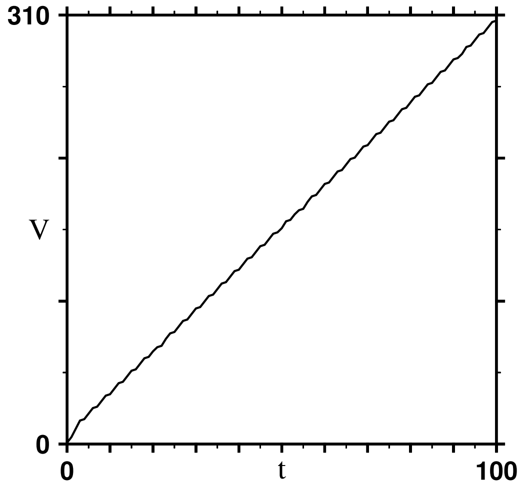

Fig. 2.3 Size V_t of global iterates Ω_t of ECA 44 vs. t.

The discrepancy between the calculation of the global evolution and the observed behavior is not particularly surprising. In terms of dynamical systems theory, the evolution of the extended state space describes the entire attractor-basin portrait (see chapter 3). In many cases, a complete description is prohibitively difficult to construct; that is in fact one of the central motivations behind the very development of the qualitative approach in dynamics. Such complete global descriptions are rarely pursued in dynamical systems theory. The main exceptions to this come in the study of flows on manifolds of constant negative curvature, uniform hyperbolic systems, and some very simple periodic systems—all of which are highly idealized systems.

---

*The figure was generated using the "FME algorithm" defined in chapter 6.

2.6 Limit Sets

Another object related to the global evolution is the **limit set** Ω_∞.[56] This is the set of all configurations that can be present in the state space for infinite time. Equivalently, it is the intersection of the iterates of the global language $\Omega_0 = \mathcal{A}^*$ in the limit of infinite time:

$$\Omega_\infty = \bigcap_{t=0}^{\infty} \Omega_t. \tag{2.16}$$

Since CA are typically dissipative, in many cases this set is extremely rich in structure. Ref. [56] contains examples of CA with limit sets at different levels in the hierarchy of computational complexity, including an example of a limit set that is not recursively enumerable—i.e., Ω_∞ is not capable of being generated by a universal Turing machine.[†]

The limit set is in fact different from the limit of global evolution, $\lim\limits_{t \rightarrow \infty} \Omega_t$. For example, we have seen that the size of the global machine in ECA 44 grows without bound. The limit set of this system, on the other hand, can be represented by a machine of finite size $V_\infty = 3$, as we will see in chapter 10. The global language grows without bound because of the different kinds of transient behavior, as we stated above. In the limit set, all of these transients are excluded, resulting in a simpler expression and a smaller machine size.

[†] See ref. [39] for a discussion of computational complexity.

3 Qualitative Dynamics

This chapter presents the briefest of reviews of the concepts we will need from the qualitative theory of dynamical systems. It is not intended to be a comprehensive introduction to this extensive and rapidly expanding field. Rather, it largely presents a compilation of the standard definitions and terminology that will be used or referred to later on. For a more complete treatment, see any text on the subject; we will primarily follow the development in ref. [102].

The general aim of qualitative dynamics is to characterize the geometry of the behavior of a given dynamical system. This a less detailed description of the system than the derivation of closed-form expressions for the motion. In most nonlinear systems, such closed-form solutions have remained unattainable. Rather, the system is characterized in terms of important structural features determining its behavior.

3.1 Dynamical systems and state space

A **dynamical system** consists of a **state space** $\mathcal{X}$ and a **dynamic** $\mathcal{T}$, both of which are specified at the outset of analysis. $\mathcal{X}$ is the space of all possible **states** x of the system and $\mathcal{T}$ is the **rule** determining the way in which a state x evolves in time. In the most common formulation of the theory, x is taken to be a finite set of real-valued quantities—i.e., an element of $\mathbb{R}^N$ (though this is not necessary and is not the case for cellular automata)—so that $\mathcal{X} \subseteq \mathbb{R}^N$. In systems with discrete time, the state at a given time t is written x_t and the dynamic is a mapping $\mathcal{T} : \mathcal{X} \to \mathcal{X}$ implemented by a function f such that

$$\mathrm{x}_t = f(\mathrm{x}_{t-1}). \tag{3.1}$$

A sequence of states $\{\mathrm{x}_0, \mathrm{x}_1, \mathrm{x}_2, ...\}$ with $\mathrm{x}_t = f(\mathrm{x}_{t-1})$ is an **orbit** or **solution** of the system, with x_0 the **initial condition**. In general, f need not be invertible. If f noninvertible, evolution is **irreversible**. The evolution of an ensemble of states $U \subseteq \mathcal{X}$ is written

$$V = f(U) = \{v : v = f(u), u \in U\}. \tag{3.2}$$

In continuous time, the state is written $\mathrm{x}(t)$ and dynamic is represented by the differential equation

$$\dot{\mathrm{x}} \equiv \frac{d\mathrm{x}}{dt} = f(\mathrm{x}). \tag{3.3}$$

An orbit (or solution, or **flow**) is then a continuous curve parametrized by time and initial condition $\mathrm{x}(0)$, written $\mathrm{x}(t) = \phi(t, \mathrm{x}(0))$. Ensemble evolution is written, for any $U \subseteq \mathcal{X}$,

$$V = \phi(t, U) = \{v : v = \phi(t, u), u \in U\}. \tag{3.4}$$