

Linear Cellular Automata

Harold V. McIntosh
Universidad Autonoma de Puebla
Apartado Postal 461
(72000) Puebla, Puebla, MEXICO
Email:mcintosh@servidor.unam.mx

May 20, 1987
revised August 10, 1990

Contents

1	History	1
1.1	Early origins	1
1.2	Automata theory	2
1.3	The Gardner era	3
1.4	The Wolfram era	5
1.5	One dimensional automata	6
1.6	The vocabulary of automata theory	7
1.7	Probability theory	9
1.8	Graph theory and de Bruijn diagrams	11
1.9	Reversible automata	12
2	What to look for	15
2.1	General characteristics	15
2.2	Cycles	16
2.3	Periods	16
2.4	Ancestors	17
2.5	Subautomata	18
2.6	Factor automata	19
2.7	Product automata	19
2.8	Automata with memory	20
2.9	Idempotent rules	21
2.10	Membranes and macrocells	22
2.11	Totalistic rules	23
2.12	Two-cell neighborhoods	24
2.13	Blocking transformations	25
2.14	Tailor made automata	26
3	Cycles in space	27
3.1	Cycles in <i>Life</i>	27
3.2	Evolution for Rule 22 - one dimensional <i>Life</i>	28
3.3	Evolutionary diagrams and matrices	30
3.4	Evolution of a seven-cell ring	31
3.5	Cycles for a sixteen-member ring	33
3.5.1	Quiescent ring	33
3.5.2	Period 14	34

3.5.3	Period 12, 6 phases	35
3.5.4	Period 12, 12 phases	35
3.5.5	period 6	35
3.5.6	Period 4	37
3.5.7	Period 2	37
3.5.8	Period 1 (still life)	37
3.6	Cycles for Rule 22	38
3.7	The evolution matrix	40
3.8	The reduced evolution matrix	41
4	Periods in time	45
4.1	Characteristics of cycles	45
4.2	Overlapping of neighborhoods	46
4.3	The de Bruijn matrix	47
4.3.1	B_1	47
4.3.2	B_2	48
4.3.3	B_3	48
4.3.4	B_n as product and sum	49
4.4	Periods and other properties	51
4.5	Superluminal configurations for Rule 22	53
4.6	Periods for Rule 22	54
4.7	The gap theorems	58
5	The Garden of Eden	61
5.1	The subset construction	61
5.2	Excluded states for Rule 22	62
5.3	Symbolic equations	63
5.4	Arden's lemma	65
5.5	The use of symbolic equations	66
5.6	Systems of symbolic equations	69
5.7	Ancestorless states for Rule 18	71
5.8	Factors	73
6	The calculus of regular expressions	75
6.1	Derivatives	75
6.2	Ideals and factors	77
6.3	Rule 18	80
6.4	Factors for Rule 18	83
6.5	A geometrical representation	85
7	Probabilistic de Bruijn matrix	89
7.1	Block probabilities	89
7.2	Kolmogorov conditions in matrix form	91
7.3	Probabilistic de Bruijn matrix	92
7.4	Some properties of n -block probabilities	94
7.5	Some simple examples	96
7.6	Determinant and inverse	97

7.7	Chacteristic equation	98
7.7.1	merged product	98
7.7.2	trace	98
7.7.3	second coefficient	99
7.7.4	principal minors	99
7.7.5	determinant	100
7.8	Correlations	100
8	Probabilistic evolution matrix	103
8.1	Regularities and anomalies	103
8.2	Mean field theory	104
8.3	More refined theories	105
8.4	Local structure theory	106
8.5	Hartree-Fock approach	107
8.6	Kolmogorov consistency conditions	108
8.7	The vector subset diagram	111
8.8	Estimating the number of ancestors	112
8.9	Trivial solutions	114
9	Positive matrices	115
9.1	Gerschgorin's disks	115
9.2	Eigenvalues on the boundary	116
9.3	Minimax principle	118
9.4	Largest eigenvalue	119
9.5	Second largest eigenvalue	120
9.6	Averaging and convergence	121
9.7	Non-negative matrices	123
10	Zeta function	127
10.1	Counting loops	127
10.2	Traces, $\chi(t)$, $\zeta(t)$	129
10.3	Infinite de Bruijn martrix	131
10.4	Cluster expansion	133
10.5	Reduced evolution matrix	135
A	Cycles for Rule 22	139
A.1	Summary	139
A.2	$N = 1$	142
A.3	$N = 2$	142
A.4	$N = 3$	142
A.5	$N = 4$	142
A.6	$N = 5$	142
A.7	$N = 6$	142
A.8	$N = 7$	142
A.9	$N = 8$	142
A.10	$N = 9$	143
A.11	$N = 10$	143

A.12 $N = 11$	143
A.13 $N = 12$	143
A.14 $N = 13$	143
A.15 $N = 14$	143
A.16 $N = 15$	144
A.17 $N = 16$	144
A.18 $N = 17$	144
A.19 $N = 18$	144
A.20 $N = 19$	144
A.21 $N = 20$	145
A.22 $N = 21$	145
A.23 $N = 22$	145
A.24 $N = 23$	146
A.25 $N = 24$	146
A.26 $N = 25$	146
A.27 $N = 26$	147
A.28 $N = 27$	147
A.29 $N = 28$	148
A.30 $N = 29$	149
A.31 $N = 30$	150
A.32 $N = 31$	151
A.33 $N = 32$	152
A.34 $N = 33$	153
A.35 $N = 34$	154
A.36 A pair of isolated cells	155

Chapter 1

History

Up to now, the history of cellular automata seems to revolve around three outstanding events, or periods of heightened interest, although there has always been a fairly steady stream of contributions. The first of these was John von Neumann's self-reproducing automaton, the second, Martin Gardner's popularization of John Conway's game of *Life*, the third, Stephen Wolfram's classification of automata.

1.1 Early origins

Cellular Automata have been studied as a part of the abstract theory of computation since the time that John von Neumann became interested in the possibility of constructing self-reproducing automatic factories. Since actual factories and physical machinery involve a myriad of practical but non-essential details, he eventually followed a suggestion of Stanislaw Ulam that an abstract mathematical model would be more amenable to a demonstration of the possibilities of universal construction and self reproduction. He worked out a scheme for such an automaton, in terms of a cellular space occupying a two dimensional grid, in which each cell would be found in one of twenty nine states.

The details of von Neumann's construction remained unpublished at the time of his death in 1957, but were subsequently edited and published by A. W. Burks [90]. Even as he was working on his model, von Neumann realized that it was too literal an interpretation of the computing machines of the era, but he himself never attempted to carry out a complete revision of his original design. About a decade later, in the years 1964-65, E. F. Codd [25] worked out a variant which required only eight states per cell, still using the original five cell neighborhood.

Ulam's [111] work on functional iteration and his experiments on nonlinear mappings were reported in conference proceedings, and in the course of time cellular automata became a topic in the theory of abstract machines, along with the work of Edward F. Moore, Claude Shannon, and others. The principal results of the time were the demonstrations of the existence of universal constructors, and Moore's "Garden of Eden" theorem which showed the necessary existence of configurations in automata of a rather representative type which could only be initial states.

Such a pattern could never again be repeated during the course of the automaton's evolution.

Of course, there were even earlier beginnings to automata theory, for example in the studies of Warren S. McCulloch and Walter Pitts in 1943 on neural nets, followed in 1951 by an interesting mathematical abstraction to regular events by S. C. Kleene, and even in general ideas about the new subject of Cybernetics introduced by Norbert Wiener [115] in his famous book of 1948.

An interesting parallel development arose from an even older source, Henri Poincaré's emphasis of the qualitative aspects of classical mechanics in terms of stability, ergodic properties, and the recurrence of orbits; such topics nowadays constitute measure theory, topology, or symbolic dynamics, all of which are an outgrowth of various of Poincaré's ideas.

The latter, symbolic dynamics, elaborated by George Birkhoff [15] and others, was described by Walter Gottschalk and Gustav A. Hedlund in an American Mathematical Society *Colloquium Publication* [45] in 1959, and continues to be a subject of intense mathematical interest. Hedlund's very abstract summary [56] of 1969 contains a wealth of results applicable to cellular automata, although its orientation is entirely different and the relationship between the two concepts has not always been well appreciated.

1.2 Automata theory

Cellular automata are but a specialized instance of the general theme of automata theory; the difference lies in the fact that automata are driven by input signals and produce output signals. Cellular automata enjoy all the symmetries, mostly translational, inherent in their crystallographic layout; but they use the states of selected neighbors for input signals and are not generally considered to produce output.

Automata theory itself has an ancient history, if one thinks of automata as mechanisms capable of performing intricate movements; but if the actual apparatus is discarded in favor of the activity itself, such a theory more properly begins with the neurophysiological abstractions of McCulloch and Pitts. Their refinement into the theory of regular expressions by Kleene constitutes one of several viewpoints, which have gone on to include semigroups (or monoids) of mappings of a set into itself, or even the theory of grammars.

The semigroup aspect has been exploited by Kenneth Krohn and John L. Rhodes [68]; that was followed up with monoid theory by Samuel Eilenberg [38, 39]. Semigroup theory is much more intricate than the theory of groups, whose classification has been one of the more important mathematical accomplishments of recent times; but a general knowledge of the principles involved is still very useful for automata. These principles include the definition of one sided and two sided ideals, and their use to give a semigroup a standard structural form.

The grammatical approach is primarily a creation of the MIT linguist Noam Chomsky [20, 21, 22, 23], referring primarily to the sequence of transformations which a system can undergo. The emphasis is more on characterizing a set of symbolic equations describing the transformations, than on the transformations

themselves, which is the province of semigroup theory. The details to be resolved consist principally in establishing the existence and possibly the uniqueness of the solutions.

The decades of the fifties, sixties, and into the seventies, saw a tremendous amount of research on automata, languages, and similar topics; cellular automata as such were not given much special attention, although sometimes they were utilized as examples or illustrations.

Some of the themes treated, although not necessarily unique to cellular automata, but often emphasizing one dimension, were the “busy beaver” problem in which long lasting activity from simple initial configurations was sought, the “firing squad” problem in which evolution to a specified configuration was the goal, or the ability to achieve “universal computation” via the emulation of a Turing machine.

1.3 The Gardner era

Public awareness of cellular automata can mostly be attributed to John Conway’s interest in finding a simpler configuration than von Neumann’s and exploring its capabilities. Some of his results were presented in 1970 as an ecological game called *Life*, at a time when such concerns were popular, in Martin Gardner’s monthly *Mathematical Games* column in Scientific American. For a period of about three years Robert T. Wainwright maintained a quarterly newsletter [113] disseminating discoveries made by Martin Gardner’s readers, some of which were followed up in later columns in Scientific American. Many of the more interesting results were obtained at MIT’s Artificial Intelligence Laboratory with the help of the graphics facilities of their PDP-6 computer ¹.

When microcomputers began to attract popular attention, Conway’s game of *Life* became one of the early inspirations for an application; Cromemco’s “Dazzler,” a color video controller and one of the earliest peripherals, was frequently used to display the evolution of *Life* configurations. Early issues of Byte magazine contained some material on *Life*, but in one memorable issue many results which had appeared in Wainwright’s newsletter [113] a decade earlier, but still not reached mass circulation, were presented by some of their discoverers. Other magazines, such as Omni, also revived the topic [53, 85], and in recent years Scientific American has returned to the subject, most recently in A. K. Dewdney’s *Computer Recreations*, the current successor to Martin Gardner’s column.

All of Martin Gardner’s columns of the early 1970’s have been reissued in a recent reprint collection [43], together with some of his reminiscences. He identifies three places as having been particularly active centers of interest, at all of which people were mainly concerned with collecting examples of one type of behavior or other, many of them reported in Wainwright’s newsletter. Many fanciful names were invented to describe the variety of configurations which were found, including two (glider gun and puffer train) due to Conway himself.

Conway had devised the rules of evolution of *Life* carefully, to avoid the extremes in which live cells proliferated and grew without bound, or in which live

¹for an account, see Levy’s book [72]

cells dwindled and eventually died. He was still not sure about the ultimate fate of his delicately balanced creation; there could always be some fairly uncommon combinations nevertheless capable of unlimited growth. Several small patterns were known which delayed thousands of generations before their final behaviour became evident.

A five-cell figure called a “glider” capable of diagonal movement had been discovered in the early stages of experimentation, as well as larger figures – “space ships” – capable of horizontal or vertical movement. An isolated “glider gun” that produced gliders periodically would be one structure with unlimited growth, a puffer train formed by a space ship which left permanent debris behind as it moved would be another. Whether such structures existed was not clear at first, but it was evident enough that regular structures would be required before theorems could be proved. Otherwise there were simply too many courses of evolution open to study them all.

At MIT glider guns were quickly found; also glider collisions were studied in great detail. Interestingly enough, puffer trains were also found, even some that included gliders in their residues. Through this combination a mass was found which, not really dense but still reasonably compact, violated Conway’s conjecture against an infinitely growing configuration in the worst possible way. That is, it occupied a diamond shaped region of constant density whose borders expanded at a uniform velocity – a quarter of the “velocity of light” – because of the prevalence of gliders in its makeup.

Glider guns created a steady stream of objects - gliders - which could be used as signals in something resembling an electrical network, which enabled Conway to design a configuration which fulfilled von Neumann’s original goals of creating self-replicating automata. In the process it was made clear that many questions about automata theory were undecidable because their answers would also have to describe the Turing machines which could be associated with the process. However, that conclusion can be reached much more directly by embedding linear cellular automata which emulate Turing machines in two-dimensional environments. However such machines use many more than the two states per cell which Conway’s construction achieves.

Work at MIT also revealed some interesting schemes for constructing very large configurations of period two, even some which completely vanished in the process of creating the new generation and which in turn disappeared as they recreated the original generation. Another group of students in Canada worked out large numbers of oscillators of different periods, on the basis of the strategic placement of structures which came to be called “eaters.”

It was evident from the first that some structures were much more stable than others, and that they could be used as building blocks to regulate the growth of more erratic or less stable combinations. Indeed, glider guns seem to have been a byproduct of having first constructed oscillators by confining colonies with promising but undisciplined growth, and then checking out collisions of the moving parts of such oscillators.

Of course, given moving objects such as gliders or space ships, it was natural to look into their collisions with stationary objects or with each other. Rather than

simply admiring the spectacle, one can take a scientific approach, systematically varying the relative positions of the reactants and recording the results. In one way or another, some extremely varied and highly interesting configurations were discovered.

Wainwright undertook a classification of all the noteworthy *Life* configurations which his correspondents reported, and one could fairly say that the “Gardner era” of cellular automata was characterized by an intensive search for “interesting” configurations in Conway’s *Life*. Some modest variants were considered by various persons; but *Life* itself was sufficiently challenging, and the variations exhibited by the variants insufficiently spectacular, to result in any great volume of reports.

The extent of the diverse results which were obtained was indeed impressive, but it was inevitable that the field would become saturated as the readily imaginable concepts of “interesting” were explored, and the combinations accessible to the computer technology of the day were exhausted. The next steps depended not only upon a new computer technology, but upon advancing concepts in information theory and even in formal language theory.

1.4 The Wolfram era

Professional scientific interest in cellular automata received a considerable impetus from the investigations of Stephen Wolfram [118], who undertook a computer based search through the properties of one dimensional automata, guided by some concepts from the realm of nonlinear dynamics and statistical mechanics. For anyone, in fact, the microcomputers, programming languages, and video displays which are currently available are sufficient for many experimental studies of cellular automata, not a few of whose results have considerable artistic merit.

A recent article exploring the aesthetic side of automata theory was entitled *Abstract Mathematical Art*, by Kenneth E. Perry [95], published in the December, 1986, issue of *Byte*. The article included a Basic program for use on IBM/PC compatible computers, with indications that a Pascal version was available from the magazine, and a statement that the author himself had used machine language to quickly seek out the hundred examples which he selected for the readers’ experimentation. Several of them were shown in striking color photographs illustrating the article.

The idea of a one dimensional cellular automaton is quite simple, and its evolution in time is ideal for a two dimensional presentation, as on a video screen. To start with, a cell is a region, even a point, with differing forms, called states. For convenience, these states are usually numbered with small integers beginning with zero, rather than described. For the purposes of automata theory the nature of the states does not matter, only their relation to one another, and the way they change with time according to their environment. Since they are abstract, they can just as well be represented by colored dots on a video screen, which is what makes them so dramatic when interpreted as an abstract artistic design.

Although both von Neumann and Conway were aware that alternative rules of evolution existed, incredibly large numbers of them in fact, they concentrated on one single rule which served their purposes, exploring its consequences in detail.

Wolfram, by contrast, was one of the first to compare the evolutionary histories of large numbers of different rules, with the intent of classifying them according to their long term behavior. Indeed he seems to have been inspired by work in dynamical systems theory, particularly Stephen Smale's discovery of "strange attractors," and the possible parallels they might have in automata theory.

In any event, he noticed that the evolutionary histories of linear cellular automata were quite varied, for which he proposed four classes. The first contained automata evolved automata evolving, usually fairly rapidly, into a constant field. Quiescence, of course, was a basic assumption of von Neumann, Conway, and others who needed a static background against which to perform their constructions. Wolfram took this class to be the analogue of those dynamical systems which evolve toward a point of stable equilibrium.

The second class was supposed to correspond to the limit cycles of nonlinear mechanics, that is, limiting behavior which is cyclic. With automata, many limiting fields are not quiescent, but nevertheless show little, no, or completely predictable, activity. *Life's* still lifes are a typical example, as are the blinkers, gliders, and other isolated periodic structures which characterize its long term evolution.

The third class is quite the opposite, consisting of chaotic fields whose behavior is never predictable; these would correspond to ergodic behavior in dynamical theory or even strange attractors.

The fourth, and remaining, class is the one which may be the most interesting; Wolfram wished to associate it with universal computation, but it is essentially the same class which held Conway's interest; the reasoning was quite similar. In spite of the ordering implied by Wolfram's numbering, the fourth class is probably best regarded as an interface between the second and third classes as suggested by Christopher Langton [71]. In any event, it is a class characterized by "islands of chaotic behavior in an ocean of quiescence." One does not find such behavior in binary automata of small radius; it is always fairly uncommon, yet far from rare.

1.5 One dimensional automata

To make a one dimensional automaton, a series of cells is strung out along a line, the simplest assumption being that all the cells have the same number of similar states, and that the rules of evolution will be the same for all of them. The idea of forming the cells into a line implies a linear order, but of course other arrangements are possible, both in terms of the dimension and the connectivity of the cells. This is the second element of definition for a cellular automaton – the relationship between the cells, or the kinds of neighborhoods which they form. Again, the simplest neighborhood would consist of a cell and its nearest two neighbors; generally speaking we would take a neighborhood to consist of r neighbors on each side, giving $2r + 1$ for the total number of cells in a neighborhood.

There are some small quibbles to be made. If a chain is finite, it has ends, which surely do not have the same neighborhoods as the interior cells. Either they can be treated differently, or the chain can be imagined to close into a ring, which would preserve the uniformity of all the neighborhoods. Also, it is considered preferable to work with symmetric neighborhoods, each one centered on its own cell, rather

than worrying about irregular neighborhoods. An interesting exception would be to work exclusively with a single neighbor, always on the same side, but that is another story.

Thus there arises Wolfram's notation (k, r) for a linear cellular automaton which has k states within each cell, and such that it, together with r cells on either side, is considered to form a neighborhood. In reality, a neighborhood does not have to be centered on its cell; for example when its length is an even number of cells, succeeding generations might be staggered by half a cell, and the radius assumed to be half-integral.

There is one final ingredient in the definition of a cellular automaton, which is the rule of transition by which the cell changes state from one generation to the next, conventionally assumed to be the same rule for each neighborhood. It is the judicious selection of a rule as much as anything that makes a particular automaton interesting or not.

Conway's game of *Life* was the result of a particular choice of rule for a two dimensional binary automaton – two states per cell – whose neighborhoods contained the cell in the center and the eight cells touching it, four of them laterally and four of them diagonally. The announced criteria by which that particular rule was chosen were that a field of cells should neither dwindle away to nothing – all zeroes – or eventually fill up completely – all ones. Reportedly, he examined many different rules before choosing the particular one which gave us his famous game; even so, so much variety was encountered with that one particular rule that years passed before many others were studied.

Wolfram's further work, mostly done at the Institute for Advanced Studies in Princeton, systematically examined all the possible rules for one dimensional automata. His recent book [121], a collection of reprints comprising much of his work on automata, contains an extensive appendix tabulating diverse properties of all 256 $(2, 1)$ automata.

In two dimensions there are far too many possibilities – vastly more so even than in one dimension – for there to be a chance to try everything, even with a very fast computer. Notwithstanding, Dewdney's column in a recent issue of *Scientific American* describes a three dimensional variant of *Life* devised by Carter Bays [6, 7, 8]. With twenty seven cells per neighborhood, its analysis has to be far beyond the reach of either present methods or present and foreseeable computers. However, it is not necessary to choose such an ambitious model to gain useful insight and obtain interesting results. So, even the one dimensional automata of type $(2, 1)$ constitute a reasonable starting point for systematic studies.

1.6 The vocabulary of automata theory

One of the heritages of the interest in *Life* is a rich vocabulary of such picturesque terms as still lifes, gliders and so on, describing typical artifacts encountered during evolution; even though of two-dimensional origin, they are freely used for other automata.

Similarly, Wolfram has created some basic descriptive terms by which automata can be classified according to their rules of evolution, not to mention their classifi-

cation into four phenomenological evolutionary types.

For (2, 1) automata, a neighborhood contains three cells, each of which exists in one of two states; altogether eight possible neighborhoods. Since there is nothing to require evolution of a given neighborhood to lead to one value of the new cell or another, there are 256 possible ways that each set of eight different neighborhoods can evolve into the next generation, starting with the possibility that everything evolves into zero and ending with the possibility that everything evolves into one.

The choice of the words “zero” and “one” assumes a decision to number the states of an automaton, which is certainly convenient for computer processing, or even a plain mathematical discussion. Yet the choice is quite arbitrary; Conway must have thought that the more vivid dichotomy between “live” and “dead” or “inactive” would make his game more appealing. Naturally the states of an automaton displayed upon a color television screen would be the colors of the pixels themselves.


As suggested by the choice of a numerical representation, the easiest way to enumerate (2, 1) neighborhoods is to make up a binary number whose eight digits tell how the neighborhoods 000, 001, 010, and so on evolve. Every evolutionary rule gets its own serial number in the process; it is convenient to always do this in a standard way and refer to the rule by the decimal equivalent of the resulting number. We might call the following choice its *Wolfram rule number*: Set out the possible sequences of cells in reverse lexicographic order:

111	110	101	100	011	010	001	000
0	0	0	1	0	1	1	0

The second row shows how the three-cell neighborhoods evolve for Wolfram’s Rule 22, the decimal equivalent of the eight-digit binary number 00010110. This ordering of neighborhoods is natural if one thinks in terms of “high order first” but is just the opposite of the one used in Perry’s article [95].

The working of the rule of evolution can be seen in the following sample of ten generations of evolution of Rule 22 from a randomly chosen initial configuration. A cyclic ring is assumed, so the first cell in each row is the right neighbor of the last cell, just as the last cell is the left neighbor of the first cell.

Every neighborhood is repeated numerous times throughout the diagram; evolution is referred to central cells, whose new values can be found in the next line, just below the old cell.



A table of transitions explicitly describes the function mapping each neighborhood of a cell into the cell's next value, but for general discussions or proving theorems it is better to use common mathematical notation. The symbol $\varphi(a, b, c)$, or simply φ , is often reserved to denote this *transition function*. Of course, such a definition serves only for nearest neighbors; the function would have $2r + 1$ arguments if it belonged to a (k, r) automaton.

Either Perry's or Wolfram's representation makes clear how many rules there are, because, given a (k, r) automaton, any one of the k states can result from the evolution of each of the k^{2r+1} different neighborhoods. Because of the exponent involved, the number of possibilities increases drastically if either the number of states for a single cell increases, or the number of neighbors increases. Thus a $(3, 1)$ automaton – three states with nearest neighbors – has $3 \times 3 \times 3$ or 27 neighborhoods, and the total number of possible rules would be 3^{27} , somewhere on the order of 10^{12} , so they are not soon going to be studied one by one. Alternatively, the combination $(2, 2)$ would have 32 different neighborhoods, and thus 2^{32} different rules, which is “only” in the range of 10^{10} .

To obtain a reasonable sampling of even the smaller linear automata, Wolfram used the idea which was already implicit in Conway's statement of his rule, that the evolutionary criterion should depend on the number of cells in the neighborhood, but not on their particular arrangement. Talking in such terms reveals some hidden assumptions about our vocabulary. In Conway's binary game, zero represented a dead cell, one a live cell, and his rules were stated in terms of the number of live cells in the neighborhood. It is a simple extension of this idea to assign numbers (weights, if you wish) to the states, and make the transition depend only on their sum. A rule gotten this way is called a totalistic rule; not all rules are totalistic, but they lead to a more manageable sampling of all the possible rules.

Curiously, rule 22 cited above is also a totalistic rule. Three binary digits can add up to zero, one, two, or three. If these possible sums are arranged in descending order as before and paired with the corresponding value of the evolved cell, a *totalistic rule number* can be derived (which in this case would be decimal 4).

sum	3	2	1	0
new state	0	0	1	0

Even the use of the number of distinct sums can lead to very large numbers; also there is a question of whether the evolving cell should have its own weight included or excluded from the sum. For binary automata (but for higher dimensions) Bays has introduced the notation $w/x/y/z$ to mean: if the cell is one (alive) and there are between w and x neighbors, it survives; while if it is zero and has between y and z neighbors, a cell will be born. In this notation, Conway's original game was a $2/3/3/3$ *Life*.

Doubtless other notations will evolve as significant new combinations requiring them come to light.

1.7 Probability theory

The enthusiasm which greeted Martin Gardner's exposition of *Life* was initially focussed on tracing evolutions of two-dimensional configurations that struck one's fancy, or in trying to design configurations which would have some interesting characteristic. The original challenge was to either produce a configuration that would grow without limit, or to demonstrate that none existed. Conway proposed two mechanisms, a glider gun or a puffer train and examples of each were soon found – contrary to his expectations in spite of having conjectured them, apparently.

However, as David J. Buckingham [18] remarked in his Byte article, interest eventually turned to information-theoretic aspects of automata theory – something which would probably surprise anyone not familiar with this branch of engineering (or mathematics.) Even from the beginning, there was attention paid to the statistical properties of *Life*, Martin Gardner having reported some results of Robert Wainwright on the behaviour of "primordial soups" and their eventual evolution at fairly constant density [43, p. 237].

Indeed, any prolonged observation of a *Life* field leads to the conclusion that there are three phases of evolution: The first phase is a relatively short transient phase – at most ten or tens of generations – in which excessively high or low initial densities adjust themselves; the second phase may last for thousands of generations in which nothing seems to be definite; followed by the third and final phase in which isolated groups of cells go through predictable cycles of evolution.

Simple ideas of probability yield approximately correct but undeniably discrepant results. The evident explanation that assumptions about independence and exclusivity – after all, if the neighborhoods of two adjacent cells overlap by 50% or more, their evolutions might not be independent – identifies the likely source of error, but says little about what to do about it. M. Dresden and D. Wong [35] made some probabilistic calculations concerning the evolution of *Life* in 1975; L. S. Schulman and P. E. Seiden [102] slightly later in 1978.

It has recently become apparent that old ideas about probability nests, quite highly developed in other branches of mathematics, are ideally suited to contending with this problem. It is a theory which can be especially well developed for one dimensional systems; W. John Wilbur, David J. Lippman and Shihab A. Shamma [116], then Howard A. Gutowitz, Jonathan D. Victor and Bruce W. Knight [51] have all published interesting results using these techniques.

Not only has there been considerable interest in resolving the statistical properties of cellular automata, there is a very definite current running in the opposite direction, by which known automata are taken as the starting point for exploring the statistics of other systems whose rules of evolution are much complicated. In this way one returns to Poincaré's views of approximating the properties of a dynamical system. The underlying hope is that valid statistical conclusions will arise from dynamical systems that are much simpler than those governed by differential equations.

Regarding a discretized differential equation as a cellular automaton is too much of a simplification unless the cells have a rather large number of states; still there are some significant aspects of nonlinear systems which can be modelled by relatively small automata. In one well studied case, three characteristic regions are

recognized for the variable governed by a differential equation — quiescent, active, and transitional.

Such a system is usually quiescent; but an activity can be initiated which must proceed to completion before calm can be restored; especially in two or more dimensions some interesting self-sustaining activation chains can be found, which have been taken as models of chemical or neurological behavior [46, 76].

1.8 Graph theory and de Bruijn diagrams

Looking in another direction, the cells of one dimensional automata, by definition, form linear chains. But the neighborhoods themselves form another kind of chain, wherein they necessarily overlap. The result is an arrangement which has a very close connection with shift register theory [44], which is a systematic study of the properties of overlapping strings and how the overlaps happen to come about. In particular, the form of graphical representation known as the de Bruijn diagram [98] enters into many discussions, and can be used to organize a major portion of the theory.

The application to automata theory arises from labelling links in a de Bruijn diagram in two ways — as a neighborhood N or as the evolutionary image of that neighborhood, $\varphi(N)$. Including or excluding links according to various properties of either the neighborhood or its image results in a considerable variety of subdiagrams, or graphs, to be used for further study.

Graphs provide a finite frame of reference for describing the multitudinous paths corresponding to the actual configurations of an automaton; an economy of presentation which fully justifies linking graph theory with even the theory of general automata. Graphs in turn have such diverse representations as their connectivity matrix or a system of symbolic equations expressing their connectivity, not forgetting simple paper sketches of graphs which are not overly complex.

In turn the collection of paths through a diagram is readily described by such entities as polynomials in the connectivity matrix or even by regular expressions. The former permit the definition of a zeta function; the latter follow from the symbolic connectivity equations, and can be seen as an opportunity to use formal language theory to obtain further results. Both approaches have been discussed in the literature.

Still further aspects of graph theory are available to describe automata and their evolution — such as dual diagrams, cartesian product diagrams, or mappings between diagrams. For example, product diagrams serve to compare two or more paths through the underlying graph, a comparison which can profitably be used to decide questions of uniqueness or of ambiguity.

So in one way or another it pays to be familiar with the standard results of graph theory. As it is, Masakazu Nasu [86] refers the properties of injective and surjective evolutionary functions to de Bruijn and related diagrams; Wolfram [120] himself used them to express evolutionary properties, and Erica Jen [62] has used them to calculate ancestors.

Other examples can be found in the literature. Moreover, it is always possible for a retrospective knowledge of the appropriate diagrams to facilitate the under-

standing of those publications whose authors did not avail themselves of graph theory. Or at the very least, to hope that the resulting insight will lighten the burden of obtaining further results.

1.9 Reversible automata

Supposing that von Neumann's goal of exhibiting an automaton capable of universal construction or of universal computation has been realized, additional questions come to mind. For example, an automaton which contains an embedded Turing machine has to suffer some of the same limitations as the Turing machine, the most notorious of which is undecidability.

It would appear that a cellular automaton ought to be more versatile than a Turing machine; for example it could have numerous read heads and work on many computations simultaneously. Nevertheless, since any automaton can presumably be modelled within some Turing machine, the computational powers of the two artifacts must be coextensive. One seeming paradox arose when it was found that some automata had configurations which could not be the product of any evolution, the "Garden of Eden" states of Edward F. Moore [82, 84].

The reconciliation of such states with a universal constructor lies in the realization that the constructor is not required to fill space with arbitrary designs, but rather to create specific objects (including copies of itself) according to the instructions which it has received. Universality refers to whether arbitrary descriptions can be followed, not whether arbitrary constructs can be produced.

Nevertheless, there has been considerable interest in ascertaining whether or not there are restrictions on the long term behavior of an automaton, both in the remote past and in the remote future. Such restrictions could manifest themselves in unattainable configurations such as the Garden of Eden, as well as in certain limiting configurations which would only develop fully as limits.

But there is also a middle ground, consisting of rules or even of configurations within a given rule, which never end up in some inaccessible region, either past or future. For this to be true, it is especially important that there be no barrier such as the Garden of Eden, devoid of a previous existence in any form. Equally, although the future always exists in some form or other, there might be reasons to avoid an approach to extremely complicated limits.

In other words, there is an interest in time reversal symmetry, or even a simple equivalence between past and future, whereby past configurations would be just as recoverable from a given configuration as the future configurations that are deducible from the same data. Quite trivial instances in which the states shift sideways, remain unchanged, or complement themselves between generations readily come to mind; naturally real interest centers on whether there are others, besides.

S. Amoroso and Y. N. Patt [4] searched for possible reversible rules and discovered eight nontrivial $(2, 3/2)$ rules in 1972; Tommaso Toffoli [109] found a way to generate whole classes by an increase in dimension in 1977, and Edward Fredkin discovered another scheme which has been reported in Toffoli and Margolis's book [110]. Fredkin employed evolutionary functions extending over two generations, evidently extendible within the same framework to three or more generations. Yet

another alternative employs some of the states in a cartesian product to create a sort of memory.

It is interesting to observe that the whole idea of reversible automata falls within the province of dynamical systems (once the connection is realized) just as it was already reported in great detail by Hedlund [56] in 1969. A recent survey article by Roy Adler and Leopold Flatto [2] contains a good exposition of the relationship of symbolic dynamics to flows through a graph.

The reason that symbolic dynamics has such a special relationship to reversible automata lies in the fact that a dynamical system can be given a topology, with respect to which it is much easier to discuss infinite systems and the existence of limits. This in turn brings up such concepts as continuity, the multiple valuedness of mappings, and the existence of inverse functions.

An essential element of the discussion hinges upon the fact that the evolution function for automata is not invertible in and of itself, but the discrepancies between counterimages can be pushed to remote regions of the automaton. With the help of the topology they can then be made to vanish in the limit, providing a context in which the reversibility of evolution can be discussed.

Chapter 2

What to look for

Before studying the properties of linear cellular automata in detail it may be worth a quick review of a few of their most visible features, as well as a summary of some of their more important features. This includes the Wolfram classes, some traditional elements of mathematical structure, specialized rules of evolution, and relationships between different automata. Everything will be discussed again at greater length, very likely in its own chapter.

2.1 General characteristics

To watch the evolution of an arbitrary linear automaton from a random initial configuration is to see a great deal of confusion. Gradually – in some cases quite quickly – it becomes apparent that each rule of evolution has its own personality, and that as rules and types of automata are varied, similarities are as apparent as differences. Presumably this led Conway to seek out rules for which random configurations eventually settled down to simple activity rather than disappearing entirely, remaining motionless, or filling up the entire space. However, there seem to be automata for practically every taste.

Wolfram laid down a serviceable classification into four categories

- Class I – evolution to a uniform state
- Class II – evolution to isolated cyclic states
- Class III – evolution to comprehensive cyclic states
- Class IV – evolution to complex isolated states

which were derived from some classifications in nonlinear mechanics. His attention was particularly attracted to the Class IV states. It seems that these are to be found for rules whose de Bruijn diagrams contain certain loops. These can be readily detected, for short periods at least.

A good starting point, having selected a specific rule, is to work out a table of periods (time repetition) and cycles (space repetition) in which a given row shows the number of cycles of given period in rings of length given by the columns. It

may be a bit disappointing that so few rows can be obtained within the limits of computer memory and running time that presently exist. The exponential growth of resources required ensures that rows or columns will only be added one at a time, and gradually at that, as computer power increases. Still, the first few rows and columns can actually be done; the information obtained can be quite informative.

2.2 Cycles

There are two ways to obtain the cycles for a given automaton. The first is to enumerate all the rings of the desired length, and follow up the evolution of each. In doing so various shortcuts can be taken, such as generating the configurations in Gray code order so that only a single cell changes state from one to the next. Still lifes can be detected very quickly this way. Numerical comparison of successive generations means that whenever the new generation is smaller, it will already have been examined and need not be pursued further.

The second way is more systematic and is worth the bookkeeping effort involved. A graph whose links are determined by evolution is prepared, following which a path enumerating procedure is followed to locate all the loops, whose lengths will give the periods of all the cycles of that length. Cycles of length up to ten can be obtained easily, twenty with effort, but passing thirty requires dedication; for binary automata it is slightly easier, increasingly more difficult for others.

Certain theoretical conclusions do not depend on practical limitations. For example, we might select an initial ring and examine successive generations of its evolution. The new generations can all be distinct, or at some point we might find that an earlier generation has repeated itself. For a finite ring the first possibility cannot occur, because there are only k^n possible different combinations of states in a ring of length n – 32 in a binary ring of length 5, for example. If no repetitions were allowed, the supply of distinct rings would eventually be used up, so repetition is the only alternative. In practice, the number of generations elapsing before one of them repeats may be fairly large, but is usually much smaller than the exponential bound of the worst possible case.

Once repetition has occurred, we need to recall that evolution is uniquely defined – there is just one successor to a given ring in each new generation – so the tail end of the evolution has to repeat over and over again in a fixed cycle, starting with the generation which first repeated.

Statistics of interest concerning the cycles include: the number of cycles and their lengths, the height of the transient trees leading into the cycles, and the convergence factor at each node of these trees.

2.3 Periods

The rows of the period-cycle table can be found from de Bruijn diagrams in the same way that the cycles can be found from the evolution diagram; since $2r+1$ cells are needed to deduce a generation of evolution, only about half as many periods as cycles can be worked out for a given amount of effort. This anomaly is really

an artifact of the way r parameterizes the neighborhood, and would disappear if half-integral increments were taken for r .

Similar theoretical conclusions are possible, since the periods are gotten from a subset of the de Bruijn diagram. A $2r$ -stage de Bruijn diagram for k symbols has k^{2r} nodes; k times as many links. Once this number of links has been used up in constructing a path through the diagram, one of them would have to be repeated. Thus there is also an exponential upper bound in the rows of the period-cycle table. For example, if an automaton has a cycle of period 2, it must already show up in some short ring; if it has not appeared in rings below a certain limit, it will never appear in longer rings.

Similar statistics can be compiled for the de Bruijn diagrams: the number of periods and their lengths, the number of transients leading into loops, and the corresponding convergence and divergence factors. Failed loops are still interesting; they correspond to strings which are periodic, but which dwindle away with each generation, eventually leaving a residue which is no longer periodic.

Both the cycle diagrams and the period diagrams may have intersecting loops; this simply means that choices are present at certain junctures in working up a chain of cells with a certain property, leading to a greater variety of sequences than would otherwise occur. This choice is particularly vivid in the de Bruijn diagram when one of the loops consists of a single node of identical cells embedded in another loop, since these can be strings of Wolfram's Class IV.

Not only can the de Bruijn diagram be used to reveal the periodicities of a given cellular automaton, it can also be used in automaton synthesis. That is, desired loops can be marked out first in the diagram, and then a rule chosen which respects the links. For period 1 properties the mapping has to be direct, since each link in the de Bruijn diagram corresponds to a distinct neighborhood in the automaton. Including or excluding it from the period diagram either determines or limits the value of the transition for that neighborhood. For longer periods the evolution corresponds to a composite rule, which may or may not be factorizable in the required manner.

2.4 Ancestors

Another use of the period diagram is to obtain ancestors of a given chain. The simplest application is to find the ancestors of uniform chains, which follows readily from the fact that each link represents the evolution of the central cell in one neighborhood. All the links evolving into zero determine the chains which must evolve into zero; conversely demanding that given loops evolve into zero determines the rules for which such an evolution is possible. For a binary automaton, that is enough information to determine the rule uniquely.

An interesting exercise is to show that for totalistic automata, the ancestor diagram for constant chains consists of pure loops, without any transients at all. One consequence of this result is that any finite chain which maps into zero can be embedded in a still longer chain which also maps into pure zeroes.

Only slightly more complicated is the determination of the static chains, or still lifes. Rather than choosing neighborhoods which evolve into a constant value,

neighborhoods for which the central cell evolves into itself must be chosen. Quite a few other variations on the theme are also possible.

Historically, de Bruijn diagrams were created to solve the problem of finding all the distinct sequences of certain symbols [98]. This idea can be applied to a period diagram, by asking whether all possible sequences of cells can appear as possible evolutions. Since the period diagram is a restriction of the de Bruijn diagram, it may be suspected that they may not; this confirms the existence of “Garden of Eden” states for cellular automata. These are chains of cells which can only be seen as initial configurations for an automaton, because they have no ancestors and cannot arise during the course of evolution from any other states.

Searching for a path in a diagram can be a tedious process, but if no more is required than knowing of its existence, Moore’s subset diagram [81] provides a way to systematize the search. It is a new diagram, whose nodes are subsets of states; each subset is linked to the union of the states which are linked to its members. Let the evolutionary function, which links nodes in the dew Bruijn diagram, be φ . Further let A be a subset whose typical member is the node a . Then

$$\Phi(A) = \bigcup_{a \in A} \{\varphi(a)\}$$

defines the linkage between subsets.

One usually starts from the full set, supposing that it does not matter where a path begins, continuing as long as possible. Having arrived at the null subset, it is certain that there was no path of the same characteristics in the original diagram. An even more elaborate subset diagram retaining all details can also be constructed, from which the original paths with all their multiplicities can be extracted.

2.5 Subautomata

Binary automata may be judged to be less interesting because they “don’t do anything” or fall into Wolfram’s Classes I and II; but there are other ways in which automata with larger numbers of internal states can fall into a pattern of restrictive behavior. For many rules, watching the screen display for a while will reveal that one of the colors has disappeared. This would be especially noticeable for a rule in which a certain value never appeared in the rule, because it would have to be absent in all lines after the first.

A mathematician would describe this situation by saying that he was dealing with a subautomaton – one for which a subset of states could be found which was closed under evolution. What that means is that states in the given subset would evolve only into each other and into no others. Just as mathematical definitions tend to include many apparent quibbles as extreme cases of some general proposition, it might be remarked that a good many automata actually exhibit one extreme example of subset behavior. For these automata dead states evolve only into dead states, using Conway’s biological metaphor.

Here, the extreme subset is the one consisting of just the quiescent state, so that the subset automaton, strictly speaking, would be a monary automaton; a category which we might have thought that we would never need to use. In reality we have

happened upon the concept of an automaton with a quiescent state (a more elegant adjective than “dead”), and seen its characterization by saying that the quiescent state belonged to a subautomaton.

At the opposite extreme, equally a convenient quibble, the whole set of states can be considered as a subset of states forming its own automaton. The value of setting up such a vocabulary becomes apparent when we have more complicated automata in which a whole hierarchy of subautomata can be perceived and we want to make systematic comparisons between the members of the hierarchy.

2.6 Factor automata

There is another mathematical concept related to subsets, which is the idea of equivalence relations. According to this concept, two or more states of the automaton might be regarded as being interchangeable. If not actually identical, there is no essential difference between them. Watching the evolution of certain automata on the screen, there sometimes seems to be a wash of color laid over an underlying pattern. The pattern seems to endure, while the color overlay has a life of its own. This is an example of a factor automaton, in which the overlaid colors are equivalent. They form one equivalence class, black (the general background color) another; the automaton could just as well be binary and could be viewed on a monochromatic screen.

Finally, there are mappings from one automaton to another. One of the simplest examples would be to complement all the cells of a binary automaton. The complemented automaton would probably not evolve according to the same rules, but it might. For automata whose states are television colors, interchanging the colors would be such a mapping. Aesthetically the difference is striking; mathematically it is the same automaton.

Let f be a function between the state sets of automaton A and automaton A' , with equal neighborhood size and evolutionary functions φ and φ' . Functions which satisfy the condition

$$\varphi'(f(x), f(y), f(z)) = f(\varphi(x, y, z))$$

are of a special kind, evidently more compatible with the two automata than arbitrary functions. Counterimages of individual states by such a function are the prototypes of equivalence classes; equivalence means having the same image.

2.7 Product automata

The identification of subautomata and factor automata is an analytic process, in which structure is to be sought in a previously existing system. Synthesis, the converse process, seeks to build up more complex objects by combining simpler parts in prescribed ways. The most common synthetic procedure is to try to endow a cartesian product with structure analogous to that of its constituents. As would be expected, the process works well with automata.

Suppose that A is one $(k, 1)$ automaton, with transition function α , and that B is a second, with transition function β , and that (x, y) is a pair of states. We could define the new transition function φ for paired states by setting

$$\varphi((a, b), (c, d), (e, f)) = (\alpha(a, c, e), \beta(b, d, f)).$$

In other words, corresponding members of the pairs are neighbors of one another, and transitions in the two halves take place independently, even according to different rules if such is the case.

With pairs for states, a common neighborhood width, and the product transition function, we have created a new automaton, $A \times B$ which is called the product (or cartesian product) of the two automata A and B . Although the definition given applies to $r = 1$, a similar definition would apply for automata with wider neighborhoods. It is apparent that such definitions multiply the number of states in the automaton, while keeping the same neighborhoods. Thus, a $(2, 1) \times (2, 1)$ automaton would really be a $(4, 1)$ automaton, if the states were relabelled properly.

Two applications of product automata come to mind at once. One is to see how the same data evolves according to two different automata; we would set $x = y$ to see simultaneous evolution by two different rules. The second is to set $A = B$, to see how different initial data evolve for one single automaton.

It is a standard result of structure theory that a product automaton has factor automata corresponding to each of the two automata used to construct it, and that their equivalence relations are complementary. Likewise, we would want to use the notation $\varphi = \alpha \times \beta$.

2.8 Automata with memory

There is another way to use the cartesian product to create a composite automaton. Begin with the $(2, 1)$ automaton whose transition rule is α . Define, for pairs such as (x, y) (\otimes is exclusive or),

$$\begin{aligned}\varphi((a, b), (c, d), (e, f)) &= (d, c \otimes \alpha(b, d, f)) \\ \varphi^{-1}((a, b), (c, d), (e, f)) &= (\alpha(a, c, e) \otimes d, c).\end{aligned}$$

The left part of the pair simply remembers the right part from the old generation while the right part combines the new and old values of the cell. Thus both items of information are always present and can be extracted if needed. The inverse rule, whose structure is similar, follows the same procedure.

States a and e are of no consequence in defining the transition, serving simply as a one-generation memory which will be forgotten by the second generation. Nevertheless their presence is essential, likewise the fact that φ must ignore them.

As with the cartesian product, we have a $(4, 1)$ automaton with the square of the number of states and the same width. Still another way to induce an automaton in the cartesian product would be to omit the exclusive or:

$$\psi((a, b), (c, d), (e, f)) = (d, \alpha(b, d, f)),$$

from which the previous generation with respect to α can always be recovered, but not with respect to ψ . By contrast all of φ 's previous generations can be recovered, but it must not be thought that the same applies to α , even though it participates in the definition of φ .

For example, suppose that $\alpha(b, d, f) = 0$, which is not reversible at all. Then $\varphi((a, b), (c, d), (e, f)) = (d, c)$, indeed reversible, but of no help whatsoever for obtaining α^{-1} .

2.9 Idempotent rules

For some automata it can be seen that states repeat themselves; indeed this behaviour is a very prominent feature of Wolfram's Class II automata. Formally, a $(k, 1)$ automaton with transition rule φ is idempotent if

$$\varphi(\varphi(a, b, c), \varphi(b, c, d), \varphi(c, d, e)) = \varphi(b, c, d).$$

Even though the algebra is messy, it can be verified that for such a transition rule, a cell remains unchanged from the second generation onward.

It is a temptation to state the idempotency condition in the form $\varphi^2 = \varphi$; in a sense it is true that iterated φ acts like φ ; but φ has three arguments and only one value, and so it is better to give the precise definition shown above.

It is too strong a requirement to insist that $\varphi(a, b, c) = b$ always, because only the identity rule fulfils this requirement. There are also many evident variations on the theme – evolution could stagnate after the third generation rather than the second, for example. Likewise, the states could shift sideways rather than stagnating:

$$\varphi(\varphi(a, b, c), \varphi(b, c, d), \varphi(c, d, e)) = \varphi(a, b, c),$$

for example. Other variations would have the states undergoing permutations leaving the entire configuration stagnant only in those generations in which the full cycle of the permutation had run its course.

The evolution of idempotent automata, or those equivalent by one of the variations, is not particularly interesting from a dramatic point of view. Nevertheless, they form a class susceptible to a closed analytical treatment, and they are of rather common occurrence when all the possible types of rules are taken into account. There is an interesting borderline class of automata whose rules are not idempotent, but monotonic. For example, one of the $(2, 1)$ automata has the rule (& is boolean and)

$$\varphi(a, b, c) = a\&b\&c$$

which is the logical **and** of the three binary neighbors. The rule is Wolfram's # 128, and has the property that zeroes persist. Unless one began with an infinite or cyclic chain of ones, all ones must eventually disappear, and the rule looks idempotent to the zeroes left behind. It is convenient to call these rules *asymptotically idempotent* and classify them together with idempotent rules.

2.10 Membranes and macrocells

More subtle than Wolfram's Class IV automata are another class which seem to exhibit natural barriers, not just the quiescent regions which characterize Class IV. Areas of independent evolution are observed, which seem to go about their business independently of what is happening in other regions [16].

Nor are the barriers always invariable; they may go through their own relatively simple cycle of evolution while still separating regions of more complex activity. Let us call such regions *macrocells* to distinguish them from the individual cells of which the automaton is composed. When their width is short their behaviour must be periodic of relatively short duration. There is the slight difference that their evolution will be subject to different boundary conditions than if the macrocells were cyclic or part of an infinite cyclic pattern. Wider macrocells have a correspondingly longer – exponentially longer – period for repetition.

An explanation can be found by consulting the de Bruijn diagram; it is found that there are certain nodes, for which all incoming or outgoing links survive in the period diagram, even if they do not continue on to form closed loops. It is only necessary that there is an unbroken chain from a node with complete incoming links to another (possibly the same) with complete outgoing links. These chains form the prospective cell membrane; the environment of the terminal cells ensures that the membrane will evolve properly, no matter the content of the macrocells which they enclose.

Slight perturbations are possible; a membrane may be stable in most environments and attacked by configurations in others. Once it starts to dissolve, the membrane will diminish by two cells each generation, and no long term declarations can be made about the evolution. Membranes may form spontaneously when conditions are right; it is also possible to have periodic membranes or shifting membranes. With moving boundaries, the de Bruijn diagram need only be adjusted so that all the pertinent links correspond to a consistent translation. However, the positioning of the guard strings must be changed; instead of flanking the boundary membrane they must run ahead of it ensuring an adequate anticipation of the glider which they are protecting.

A final variation, which is also often encountered, is to find that not all the links beyond the terminal nodes in the de Bruijn diagram are guard links, but that the macrocell will only place cells near the membrane whose links are guarding. As an extreme case, a wall of one color could confine cells of different colors, but still dissolve when confronted with cells of its own color.

Since the critical requirement for a membrane is the presence of guard links at its extremes, its internal structure can be simple or complex as the occasion demands. Thus, there could be loops in the chain defining the membrane, allowing macrocells to be separated by membranes of varying structure.

2.11 Totalistic rules

A *totalistic rule* is one for which the transition depends upon the sum of the weights of the neighborhood. Such a rule, for a $(k, 1)$ automaton, would take the form

$$\varphi(a, b, c) = \tau(a + b + c) \bmod k$$

where τ is a function of the integers modulo k . In fact, the simple modular sum is a good example of a totalistic rule, yielding a rule for which the computations of linear algebra suffice to determine its evolution.

While the choice of a totalistic rule simplifies evolutionary calculations, it is clear that a given sum can be formed in many different ways. Thus every neighborhood possessing a fixed sum must necessarily evolve into the same state. This should not cause undue concern; since there are always more neighborhoods than states, any rule will necessarily have clusters of neighborhoods mapping into the same state. It is just that the multiplicity follows the degeneracy of sum formation for totalistic rules, giving them a statistical property which is useful.

There are more ways to form sums in the middle range than for the extremes; one might think in terms of the binomial distribution. Thus the values assigned the middle range will be relatively influential in determining the overall behavior of the automaton, while the extremes can be used for fine adjustments. One extreme determines what happens to long sequences of zeroes, the other to long sequences of $k \Leftrightarrow 1$'s, and in both cases to sequences in which these extremes dominate.

In practice, a sum of zero can only arise in one way, from a string of zeroes. Equally, the high sum can only be formed from a string of cells of the highest weight. Whatever value is assigned to these extremes will influence the color of the background, just as will the values assigned to the evolution of any other constant strings. Thus the sequencing of long strings of solid color can be read off directly from the totalistic rule. In prompting for the rule, the programs mark off the sums which would correspond to constant values.

A bit more effort is required to work out the sequencing of repeated pairs, but it is not too hard to do mentally. It is an interesting observation that for totalistic rules, the sequencing of iterated triples follows the same sequencing as for constant sequences, but in terms of the sum corresponding to the triple.

In conclusion, the values assigned the extreme sums influence the cycling of background colors and can be deliberately chosen for this effect, but the drama of this change should not overshadow the fact that the remainder of the evolution is little changed and that the change will hardly be noticeable unless the rule tends toward large areas of constant color.

The next-to-extreme sums tend to determine what happens to the fringes of regions dominated by the middle values, and it is often preferable to step through these values rather than the extremes themselves to get small changes in the patterns of evolution.

Although totalistic rules form a special class of rule, they are general enough to be representative in the sense that any other rule can be transformed into a totalistic rule, although the new rule may have a considerably larger number of states. Albert and Culick have demonstrated [3] that all that is really required is

to express the (k, r) neighborhood as a r -digit number relative to the base $k + 1$, but in such a way that the overlap between neighborhoods is taken into account. Thus, suppose that the configuration

$$\dots a b c d e f \dots$$

is rewritten

$$\dots 000a \ 00b0 \ 0c00 \ d000 \ 000e \ 00f0 \ \dots .$$

Now, suppose that the sum U of three four digit numbers $X + Y + Z = U$ has the base $k + 1$ representation $U = wxyz$. The reason for changing to $k + 1$ is that the new representation needs a zero which cannot be confused with one of the states of the original automaton. If the original states were numerical, they should be shifted to accommodate the new base; if they were not numerical, they should now be assigned numerical values. Then, define the auxiliary transition function τ to be

$$\tau(U) = \begin{cases} 00\varphi(z, y, x)0 & w = 0 \\ 000\varphi(w, z, y) & x = 0 \\ \varphi(x, w, z)000 & y = 0 \\ 0\varphi(y, x, w)00 & z = 0 \end{cases}$$

and finally, the new transition function Φ by

$$\Phi(X, Y, Z) = \tau(X + Y + Z).$$

With this definition, we have embedded a $(k, 1)$ automaton with an arbitrary rule in a $((k + 1)^4, 1)$ automaton with a totalistic rule, although it is necessary to code the cells before making the simulation and to decode them afterwards. If strict adherence is made to only necessary values for cell states, one non-zero digit in the four-digit base $k + 1$ expansion, the number of states could be considerably reduced, improving the state economy. Otherwise, transitions for the unused combinations get arbitrary definitions.

2.12 Two-cell neighborhoods

Amongst the variety of transformations which can change the appearance of cellular automata are those which modify the neighborhoods. For example, it is possible to show that linear cellular automata do not need to have large neighborhoods; in fact, as Albert and Culick [3] showed, it is sufficient to have just two cells in a neighborhood. The essential idea is to replace the entire neighborhood by a single cell; of course the new cell must have enough states to describe every possible content of the original neighborhood.

Even if the cells are grouped in clusters, there will still be some interaction at the boundary, so that a new rule of evolution has to be devised. Also, since the interaction has to be one-sided if there are only pairs of cells in the new neighborhood, the original two-sided interactions can be recovered by extracting information from

alternate generations. Since three cells would form a second generation neighborhood, it can be contrived that we think of the central cell as evolving, rather than one of the edge cells.

Thus, for lack of a better notation, let us define a $(k(k+1), 1/2)$ automaton A with the rule Φ in terms of a given $(k, 1)$ automaton a whose rule of evolution is φ . The states of A are to be the $k \times k$ ordered pairs (i, j) together with the k states i . We shall only define $\Phi(I, J)$ partially through the table

		(j', k)	j
(i, j)	{	$\varphi(i, j, k)$	$j = j'$
		undefined	otherwise
i		\dots	(i, j)

There is an enormous number of possible definitions of A , since the mixed transitions involving pairs and singlets are left undetermined. In a given case, some guidance may be available to complete the table, but for the present construction the completion is irrelevant.

To recover the operation of a from the initial configuration $\dots ijk\dots$ we need to see that A generates first $\dots (i, j)(j, k)\dots$ and then $\dots \varphi(i, j, k)\dots$. Thus every second line contains the evolution of a , but we have to be observant of how it is positioned inasmuch as A always leaves its result in the left cell.

In order to guarantee the substitution of a two-cell neighborhood for a three cell neighborhood it has been necessary to augment the states by state pairs. For a binary automaton this means that we have six states rather than two. We defined the rule of evolution Φ for just enough arguments to ensure the recovery of the automaton a , which means that there are really many automata capable of making the simulation. There may be some reason for selecting the remaining transitions in a certain way, perhaps to simplify the automaton somehow. It also shows that one automaton may be capable of simulating another if only one can detect the proper embedment.

2.13 Blocking transformations

There are no doubt countless ways of embedding an automaton in another whose neighborhoods are constructed differently, or which has a different state set; but one mapping which might be among the first to be tried out would be to group several cells together and treat them as a single unit. It is not as useful a procedure as it might seem at first because there are necessarily redundancies in the evolutionary rules. This is because cells which lie beyond the range of interaction after the grouping cannot influence the evolution, so that blocked neighborhoods in which they differ must have the same rule of evolution. Of course there is a small possibility that this might turn out to be an advantage going in the opposite direction, whereby the number of states might be reduced by dividing the cells into pairs.

For simplicity, suppose that a $(k, 1)$ automaton undergoes an evolution described by the following sequences of cells

$$\begin{array}{cccccccc} \dots & x & a & & b & & c & & d & y & \dots \\ & & & & \phi(a, b, c) & & \phi(b, c, d) & & & & \end{array}$$

Then the transition rule for the blocked $(k^2, 1)$ automaton is defined by

$$\Phi((xa), (bc), (dy)) = (\varphi(a, b, c), \varphi(b, c, d))$$

Conversely, any $(k^2, 1)$ automaton for which a coding of its states could be exhibited satisfying this rule could be regarded as a $(k, 1)$ automaton by splitting its states.

2.14 Tailor made automata

Two contrasting ways to approach automata theory are to begin with an automaton to see how it evolves, or to think of an interesting evolution and try to select the rules accordingly; the latter were von Neumann's and Conways approaches. From either aspect it is convenient to have a random number generator available; in the first instance for following the evolution of typical configurations, in the second it can even provide the rule to study.

A more deliberate approach begins specific plans for the evolution, and attempts to deduce a suitable rule. Suppose that the objective is a slow glider, that advances two cells in four generations. There is at *least* a 5×5 rectangle to be considered, to accommodate the generations, the displacement, and the neighborhood length. The evolution of at least 25 neighborhoods must be defined, but fewer may suffice on symmetry grounds; at least five if every row is the same.

Neither $(2, 1/2)$, $(2, 1)$, nor $(3, 1/2)$ automata look promising with 4, 8, and 9 distinct neighborhoods, respectively; but any other state-radius combination beginning with the 16-neighborhood $(2, 3/2)$ ought to have rules with such a glider. These figures may not be precise, but they lend plausibility to the conjecture that there is an automaton of some minimal complexity for almost any task that can be described by a sample evolution.

One typical task which is quite instructive is to emulate a Turing machine; Albert and Culick [3] suggest one approach. Just as there are simpler Turing machines than the universal machine, there are numerous simple automata to be designed. Binary counters are a good example; their existence shows that there are automata with very long transients, also with very long cycles. A challenge might be to come as close as possible to the theoretical upper limit, a sort of variant on the "busy beaver" problem for Turing machines.

Another traditional exercise is the "firing squad problem," wherein there is a reserved state, "fire" which all the cells of an automaton of arbitrary length are to enter simultaneously, none having occupied it previously [83, 114]. It is a whimsical one-dimensional variant of one of the components of von Neumann's constructor; which was supposed to activate each finished object, as a finishing touch when it was finally completed and turned loose on its own.

Chapter 3

Cycles in space

Fundamental to finite automata is eventually periodic evolution, prolonged to the full number of states in the exceedingly rare event that their sequence is cyclic. The number of “states” of a cellular automaton is the number of configurations (not states per cell), providing an exponentially large bound relative to the automaton’s length. In practice, many short cycles usually predominate over a few long ones, almost always reached through transients.

Longer automata admit longer cycles and longer transients too; the infinite limit may lack cycles. Cyclic boundary conditions locate behavior repeating over a finite range, leaving truly aperiodic configurations for a separate study.

3.1 Cycles in *Life*

When Martin Gardner first announced Conway’s game of *Life*, there was much experimentation with simple designs, some of the simplest consisting of just a single row, column, or diagonal of live cells. Columns tended to grow shorter and fatter until they formed a diamondlike configuration which in one remarkable case—a column of fifteen cells—retreated back into a column and became periodic. It was one of the first complex oscillators to be discovered.

However, instead of forming diamonds, very long columns exhibited a rather curious behaviour not unlike a binary counter. That is, except for the ends, a single column survived while two new live columns flanking it were born. The next generation saw the birth of still another pair of flanking columns, but the central three could no longer survive, leaving the new pair of columns isolated by a gap of three cells.

Then in the fourth generation two triple columns separated by a vacant column produced a single pair of columns separated by a distance of seven cells. In subsequent generations these columns repeated the behaviour already established until a collision occurred between the flanks expanding into the central region. This time a single pair of columns separated by a gap of fifteen cells was left in the ninth generation. Continuing this sequence on through further generations inevitably suggests a binary interpretation of the evolution.

Infinitely long constant columns reduce a two dimensional automaton to one dimension; all the cells in any given column behave alike, leaving all the significant information to be gleaned from a row acting as a cross section. Such configurations were called “ripples” in Wainwright’s newsletter. Adaptation of Conway’s rule shows that Wolfram’s Rule 22 for a $(2,1)$ automaton gives the appropriate description of the result, which might reasonably be referred to as *Life in one dimension*.

A great advantage of working with one dimensional cellular automata is that their time evolution can be shown on a two dimensional chart, whereas the evolution of a two dimensional structure would require a third dimension. While not impossible to show, there is too much information involved to keep the presentation from becoming extremely cluttered. Even so, an infinite, or even a very long, line is hard to manage. A systematic study could easily start with short lines, folded around to form a ring, avoiding end cells whose rules of evolution would differ from the interior cells. For short enough rings, the complete evolution of all possible configurations can be calculated.

3.2 Evolution for Rule 22 - one dimensional *Life*

Let us study some examples.

The shortest ring has one cell, which is consequently always its own neighbor. A zero must evolve into zero, while a one evolves into zero, which must repeat itself thereafter.

This result applies specifically to Rule 22, but it is equally applicable to any other rule for which uniform neighborhoods have the same transitions. General binary rules admit other possibilities, four altogether, since a field of zeroes can either be quiescent or switch over to ones. A field of ones could do the same, so either: 1) both fields are quiescent, 2) they alternate parity in a cycle of length two, 3) zero is quiescent and one vanishes in a single generation, or 4) one is quiescent but absorbs zero after the first generation.

Naturally the more states there are, the more sequences of evolution that could be followed by uniform fields; ascertaining which one should be the first order of business in analyzing any given automaton.

There are four rings of two cells: 00 evolves into 00, 01 into 01, 10 into 10, but 11 evolves into 00. Again there are many formats for the general case to follow, the more so the more states in the automaton; it is a matter of how many graphs can be constructed using k^2 nodes with an outgoing link for each.

There are eight rings of three cells, the shortest ring for a $(2,1)$ automaton in which left and right neighbors can be distinct. Shorter rings always lack certain neighborhoods, in this case, those which are nonsymmetrical. Therefore classifying short ring evolution is useful in establishing common features that could eventually distinguish classes of automata.

The transitions for three cells are:

000	→	000	100	→	111
001	→	111	101	→	000
010	→	111	110	→	000
011	→	000	111	→	000

It is already evident that there are four classes with cyclic symmetry, namely $\{000\}$, $\{001, 010, 100\}$, $\{011, 110, 101\}$, and $\{111\}$, and that it would be sufficient to describe the mappings of the classes into one another.

For this purpose it is convenient to select one typical element from each class, but a class representative will not often evolve into another representative. Therefore the mapping between representatives should be supplemented with an indication of the discrepancy. A sequence of cells is naturally interpretable as a binary number, making the least number in the class a convenient representative.

$$\begin{array}{l} \{000\} \rightarrow \{000\} \quad \{011\} \rightarrow \{000\} \\ \{001\} \rightarrow \{111\} \quad \{111\} \rightarrow \{000\} \end{array}$$

Of course, there are sixteen rings containing four cells. They form six symmetry classes: $\{0000\}$, $\{0001, 0010, 0100, 1000\}$, $\{0011, 0110, 1100, 1001\}$, $\{0101, 1010\}$, $\{0111, 1110, 1101, 1011\}$, and $\{1111\}$. Individually, the transitions are

$$\begin{array}{llll} 0000 \rightarrow 0000 & 0100 \rightarrow 1110 & 1000 \rightarrow 1101 & 1100 \rightarrow 0011 \\ 0001 \rightarrow 1011 & 0101 \rightarrow 0101 & 1001 \rightarrow 0110 & 1101 \rightarrow 0000 \\ 0010 \rightarrow 0111 & 0110 \rightarrow 1100 & 1010 \rightarrow 1010 & 1110 \rightarrow 0000 \\ 0011 \rightarrow 1100 & 0111 \rightarrow 0000 & 1011 \rightarrow 0000 & 1111 \rightarrow 0000, \end{array}$$

so that the economy of listing the transitions by class becomes more and more evident:

$$\begin{array}{ll} \{0000\} \rightarrow \{0000\} & \{0101\} \rightarrow \{0101\} \\ \{0001\} \rightarrow \{0111\} & \{0111\} \rightarrow \{0000\} \\ \{0011\} \rightarrow \{0011\} & \{1111\} \rightarrow \{1111\} \end{array} \quad \begin{array}{l} \text{(rotated right)} \\ \text{(rotated two right)} \end{array}$$

Writing down lists of transitions, either between rings or between symmetry classes, does not show the structure of the transitions to very good advantage. It is better to group them into transitive chains, in which an evolutionary sequence is shown until it repeats. Only maximal chains should be shown, which means beginning with a ring which has no ancestor if the chain has one. Sometimes there are cycles which have no ancestors outside the cycle. Since evolutionary sequences tend to converge, it may be necessary to mark some chains as continuing in another chain that has already been presented.

If we give this treatment to the thirty two chains of five elements, the result is surprisingly compact; there are eight symmetry classes, seven of which form a sequence terminating with pure zeroes, and another which evolves to zero independently (enclosing a string of classes in parentheses signifies that the string forms a cycle).

$$\begin{array}{cccccccc} 1111 & (00000) & & & & & & \\ 01011 & 00001 & 00111 & 00011 & 01001 & 01111 & (00000) & \end{array}$$

There are sixty four chains of length six, which fall into thirteen symmetry classes (the confluence of one chain with another is indicated by enclosing the junction in square brackets).

```

010111    000001    000111    000101    011011    (000000)
001111    000011    001001    111111    (000000)
001011    [000111]
(010101)
011111    (000000)

```

Several formal representations of sequences of this kind exist; two of them are through graphs and their connectivity matrices.

3.3 Evolutionary diagrams and matrices

These evolutionary sequences could be represented by diagrams, which makes them considerably easier to visualize. According to the presentation, either the rings of cells or their symmetry classes are represented by nodes in a diagram. Nodes are linked according to whether the node at the head of the arrow has evolved from the node at the tail of the arrow or not. The result is what would technically be called trees rooted on cycles. Zero, one, or more arrows may enter a node according to the number of its ancestors. Only one arrow can leave each node, because evolution is unique, and so is the symmetry class of the descendant. However, cycles are possible because a ring could evolve either into itself or one of its ancestors.

A numerical, or at least matricial, representation of these sequences is also possible. The dimension of the square matrix is equal to the number of nodes of the diagram; its elements are to be zeroes or ones. Call the $N \times N$ matrix M , its matrix elements $M_{i,j}$; then its elements will be zeroes or ones according to whether the row nodes are linked to the column nodes. A variant of the Kronecker delta represents this alternative numerically.

$$\begin{aligned}
 M_{i,j} &= \begin{cases} 1 & i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \\
 &= \delta(i \rightarrow j)
 \end{aligned}$$

Such matrices are sparse, meaning that the non-zero elements are usually few and far between. Nevertheless, even if they are not constructed explicitly, they are very useful for expounding certain properties of diagrams. To begin with, the well-known formula for matrix multiplication, together with the property that zero annihilates any product in which it participates, shows that

$$\begin{aligned}
 M_{i,j}^2 &= \sum_{k=1}^N M_{i,k} M_{k,j} \\
 &= \sum_{k=1}^N \delta(i \rightarrow k) \delta(k \rightarrow j) \\
 &= \sum_{k=1}^N \delta(i \rightarrow k \rightarrow j).
 \end{aligned}$$

Thus M^2 has zeroes except for those elements for which the row index is connected to the column index by two links. Generally, the elements of the p^{th} power of M show where there are connecting paths of exactly p steps. Of course, the non-zero elements of some of these powers may be integers greater than one, but that simply means that there are multiple paths between those pairs of nodes.

Other properties of M and its powers:

- the row sums give the number of chains leaving a node,
- the column sums give the number of chains entering a node,
- the diagonal elements tell how many loops contain that node,
- the trace gives the total number of loops.

Normally neither M nor its powers are symmetric matrices, which complicates using Sylvester's theorem to represent them. Nevertheless, since all their elements are non-negative, the classical results of Frobenius and Perron[41] apply. The principal conclusions are that:

- a largest positive eigenvalue is bounded by both row sums and column sums;
- there is a nonzero eigenvector belonging to the largest eigenvalue, none of whose elements is negative;
- under certain circumstances this eigenvalue is unique, with a unique normalized eigenvector.

These results can be used for estimating the behavior of large powers of M , and thus the general characteristics of long chains.

3.4 Evolution of a seven-cell ring

We could then set up the matrix M for the eighteen symmetry classes of the seven-cell rings. To obtain manageable indices, let us express each binary class representative in decimal. Since these numbers are not consecutive anyway, we might as well improve the appearance of M by ordering the indices as they appear in maximal chains. Figure 3.1 shows the evolutionary diagram together with its connectivity matrix corresponding to a ring of length seven.

There are seven classes without ancestors, corresponding to the zero columns. There are two static classes, represented by ones on the diagonal. Every class has just one successor, shown by the single one to be found in each row. Seven classes have unique ancestors; these are the ones with a single one in their column. Multiple ones within the same column mean multiple ancestors. All of this information can be obtained by inspection from any evolution matrix; much of it is discernible numerically by examining the product of M with a vector all of whose components are ones.

The lengths of transients can also be read off from the matrix, as well as the lengths of cycles, but this information is harder to perceive. If, as has been done here, the row indices follow the evolutionary sequence insofar as possible, a chain of ones will be found on the superdiagonal of the submatrix corresponding to the

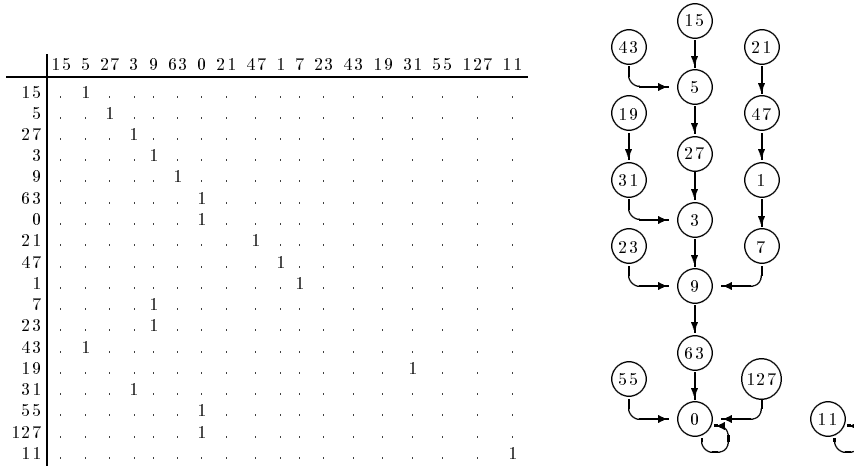


Figure 3.1: Evolution of a seven-cell ring under (2,1) Rule 22.

cycle or transient. Otherwise the chain may be fairly well hidden among the other matrix elements.

Powers of the connectivity matrix, being less sparse, are often easier to interpret because the remaining zeroes indicate disconnected groups of nodes more clearly. In any event, the elements M_{ij}^p specify the number of paths of length p running from node i to node j ; the diagonal elements therefore disclose the number of loops passing through a given node. The trace in turn counts the total number of loops; each loop is counted once for each of its nodes, but not for multiple passes through the same node.

Some results which have been calculated for rings up to length 12 are

N	rings	symmetry classes	Gardens of Eden	longest chain	cycles
1	2	2	1	2	0
2	4	3	2	2	1
3	8	4	2	3	0
4	16	6	4	3	$2x1$
5	32	8	2	7	0
6	64	13	3	6	1
7	128	18	8	7	$1x7$
8	256	30	11	9	$1x1, 1x2$
9	512	46	19	13	$3x4, 6x6$
10	1024	78	34	9	$38x4$
11	2048	126	56	12	$1x1, 24x4$
12	4096	224	105	24	$42x5$

There are 2^N rings of length N , just the quantity of binary numbers of that length; but the number of classes is a more complicated group theoretical result[44,

pp. 118-122]:

$$K = \frac{1}{|{}^oG|} \sum_{g \in G} I(g).$$

In this formula, the symmetry operations g of the group G of order $|{}^oG|$ have $I(g)$ fixed points. When the symmetry is rotational, the fixed points are sequences which repeat sooner than the full length of the ring; for reflections they are palindromes.

For rotational classes, Golomb transforms this formula into an expression involving Euler's function ϕ which counts divisors. The main point of interest is that the number of classes grows exponentially, although for longer rings it lags behind the number of configurations by a factor which fluctuates around $2N$, the size of the dihedral group of rotations and reflections.

The remaining statistics have to be obtained on a case-by-case basis for each rule and each cycle length, although some rules admit special techniques — for example if the rule of evolution is addition in a finite field. An extensive source of this kind of information is Wolfram's reprint collection[121], especially the tables of data in the appendix.

3.5 Cycles for a sixteen-member ring

The data eventually becomes so voluminous that it seems only to be worthwhile tabulating the cycles into which the rings of length N evolve, ignoring all the transients. Below, we show one particular case, $N = 16$, in detail. The ring is long enough to show a variety of interesting behaviour, but yet not so long that it cannot still be readily displayed.

The notation $x.y$ describing each cycle means that the period of the cycle is x , but that there are only y distinct phases within the cycle. This could either mean that the original pattern has reflected after y generations and will be completed by running through the mirror images, or that the pattern has been translated by a certain amount. Thus x must be a multiple of y , cycles of the form $2y.2$ almost always (but not exclusively) resulting from reflection. Sometimes an even more explicit notation, such as $16:14.7(\text{ref})$ or $15 : 20.4(3\ell)$ is required, in which the length of the ring is shown along with an indication of the direction and distance of displacement. Since reflectivity could be inferred if a displacement were not shown, it is not always indicated.

First, we show the smallest numerical representatives for each of the eight distinct symmetry classes of cycles for a ring of length $N = 16$:

Next, the eight figures which follow show a full cycle of evolution for every one of these patterns. In preparing each illustration, every line is copied twice to improve perception of the region which would otherwise be broken up by the boundary.

3.5.1 Quiescent ring

A sequence of pure zeroes is a still life whenever zero is quiescent; the resulting redundancy is generally avoided by omitting the quiescent configuration from all cycle censuses. Figure 3.3 contains the sixteen cell ring satisfying this description.




Figure 3.2: (2,1) Rule 22 rings of length 16 have eight symmetry classes of cycles.



Figure 3.3: By definition, a quiescent configuration is a still life.

3.5.2 Period 14

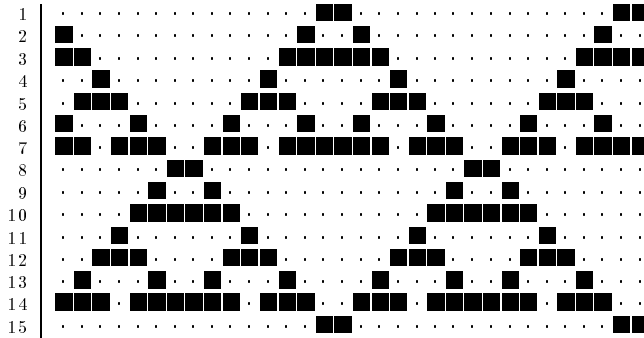


Figure 3.4: $(g.p) = 14.7$: Evolution from a joined (gap 0) pair.

Many evolutionary rules are distinguished by the presence of numerous inverted triangles in their space-time diagrams. The evolution of pure cycles shows them to especial advantage, often exhibiting a fractal structure[117]. Indeed, strings of any arbitrary format can be expected to show up occasionally; but quiescent strings are of particular interest because they can only disappear gradually, as live cells encroach on their boundaries. Triangles are the visual evidence of the eventual regrowth of the live region, evident in Figure 3.5.

3.5.3 Period 12, 6 phases




Figure 3.5: $(g,p) = 12,6$: Evolution from a gap 1 pair.

Rule 22 is a rule which acts somewhat like a binary counter, in the sense that isolated cells expand, generating regions whose interiors periodically vanish, leaving the frontier cells to expand anew. These new sites of expansion eventually collide, possibly leaving the whole new interior vacant. Properly synchronized, the binary counter effect ensues. Rule 22 has an especial tendency towards binary counters, notably amongst configurations evolving from single pairs of live cells. Figures 3.5 and 3.6 show boundary collisions leaving gaps of 1 and 2 respectively

Compatibility with a finite cycle length implies constraints which exclude certain mutual separations; most work, the rest vanish after a while or convert into other patterns. The majority of the cycles of length 16 fit this pattern; scrutiny of the figures reveals a variety of interactions amongst advancing boundaries, also the systematic depletion of the interiors.

3.5.4 Period 12, 12 phases

Some cycles run their course without returning to the same symmetry class; others may shift the original configuration several times before returning to their original form. Alternatively, they may pass through a reflected image before repeating; most cycles of the form $(2p,p)$ operate by reflection but a shift by half the length of the ring could produce the same result. Sometimes there is no distinction between translation and reflection. If N and p are relatively prime, quite long periods can result.

Amongst the last two examples, all the phases of Figure 3.6 are distinct, whereas Figure 3.5 manifests reflective symmetry. Given the way that boundaries interact, a pair with one gap may evolve into a pair with another gap, and then evolve back into the original pair, although this phenomenon is not present on a ring of length sixteen.

3.5.5 period 6




Figure 3.6: $(g,p) = 12.12$: Evolution from a gap 4 pair; no phase is shifted nor reflected relative to any other.




Figure 3.7: $(g,p) = 6.3$: A eight cell cycle repeated twice to get sixteen cells.

27			
28		1	1	1	3	1	3	.	.	.	3	3	6
29		1
30		1	1	1	.	.	1	.	.
31		1	.	.
32		1	1	.
33	
34		1

In addition, there were many cycles in the range $21 \leq N \leq 34$ with longer periods than convenient for a table, which are listed below.

N	P	N	P	N	P	N	P
21	77.11	25	55.55	27	1215.135	30	1070.107
23	138.6	25	150.6	29	72.36	31	248.8
24	50.25	25	150.6	29	667.23	31	248.8
24	54.9	26	78.6	30	40.4	32	41.41
24	72.6	26	78.6	30	70.35	32	84.42
24	72.6	26	90.45	30	86.43	33	138.69
24	80.40	26	546.21	30	120.60	34	52.52
24	384.32	27	162.6	30	240.8	34	4590.135
25	50.2	27	459.17	30	240.8		

3.7 The evolution matrix

The evolution of an automaton can be described in matrix form, as well as by the evolution function φ . The matrix required has to be rectangular, since there are boundary cells at the ends of the string whose evolution cannot be computed from the information available in the string itself. The simplest matrix in the series describes the cells of the second generation in terms of the cells comprising their neighborhoods; two values $\{0, 1\}$ can evolve from eight neighborhoods $\{000, 001, 010, 011, 100, 101, 110, 111\}$, so the required matrix has the dimension 2×8 ; for Rule 22 it would take the form:

$$G(1 : 2) = \left[\begin{array}{c|cccccccc} & 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \\ \hline 0 & 1 & . & . & 1 & . & 1 & 1 & 1 \\ 1 & . & 1 & 1 & . & 1 & . & . & . \end{array} \right],$$

where the evident formula for the matrix elements is

$$G(1 : 2)_{i,jkl} = \delta(i, \varphi(j, k, \ell))$$

Since the matrix is not square, successive generations of the evolution cannot be obtained by raising it to powers. However, by iterating φ it is possible to obtain the matrix for further generations, and thus to relate any pair of generations. Unfortunately such a procedure corresponds to none of the commonly recognized operations on matrices (such as the tensor product).

For example, we might work out, again for Rule 22, the matrix for one generation of evolution of the five-cell neighborhoods which will produce the three-cell neighborhoods which evolve into single cells. Such a matrix would be

$$G(2 : 3) = \left[\begin{array}{c|cccccccccccc} & 000000000000000111111111111111 & & & & & & & & & & & & \\ & 00000000111111110000000011111111 & & & & & & & & & & & & \\ & 00001111000011110000111100001111 & & & & & & & & & & & & \\ & 00110011001100110011001100110011 & & & & & & & & & & & & \\ & 010101010101010101010101010101 & & & & & & & & & & & & \\ \hline 000 & 1000000000001110000001100010111 & & & & & & & & & & & & \\ 001 & 01000000000000000000000000101000 & & & & & & & & & & & & \\ 010 & 00010000000000000000000010010000000 & & & & & & & & & & & & \\ 011 & 00100000000000000000000010001000000 & & & & & & & & & & & & \\ 100 & 00000011000100001000000000000000 & & & & & & & & & & & & \\ 101 & 00000000010000001000000000000000 & & & & & & & & & & & & \\ 110 & 00000100100000000001000000000000 & & & & & & & & & & & & \\ 111 & 00001000010000000010000000000000 & & & & & & & & & & & & \end{array} \right]$$

and finally, going one step further,

$$G(1 : 3) = \left[\begin{array}{c|cccccccccccc} & 000000000000000111111111111111 & & & & & & & & & & & & \\ & 00000000111111110000000011111111 & & & & & & & & & & & & \\ & 00001111000011110000111100001111 & & & & & & & & & & & & \\ & 00110011001100110011001100110011 & & & & & & & & & & & & \\ & 010101010101010101010101010101 & & & & & & & & & & & & \\ \hline 0 & 1010110011100111011110111011010111 & & & & & & & & & & & & \\ 1 & 01010011000110001000010010101000 & & & & & & & & & & & & \end{array} \right]$$

According to this scheme, one has $G(1 : 3) = G(1 : 2)G(2 : 3)$, and a clear procedure for advancing through further generations. Since the dimension of these matrices increases rapidly – doubling each generation – and since they have very few non-zero elements, they are of more interest for theoretical discussions than as a practical means of computation. In keeping with the fact that they represent a function, there is exactly one non-zero element in each column, but since they are not square, some rows must necessarily have *several* non-zero elements. The exact distribution will vary from rule to rule and to a great extent will characterize the rule involved.

3.8 The reduced evolution matrix

The matrices $G(i, i+1)$ are rectangular rather than square because of the incomplete neighborhoods at the ends of a chain of finite length, which preclude calculating the evolution of all the cells of the chain. This incompleteness was fully understood in defining such matrices, but it would still be useful if the evolution of a block of cells could be related to the block itself. One approach would be to close the block into a ring. Variations on the same theme would involve extending the block in various ways before closing it into a ring, somewhat in the way that a Möbius strip is formed by giving a strip of paper a twist before joining its ends to form a cylinder.

Another approach to compensating the lack of knowledge of what lies beyond the limits of a finite chain would be to characterize the possible extensions of the chain numerically and either storing this information in a square matrix, or possibly annotating the links of a connectivity diagram.

For example, let the rows of the matrix be indexed by the 2^N sequences of N bits, and the columns by the 2^{N+2} central sequences of $N+2$ bits which can evolve

into their indices via one of the (2,1) rules. The number of counterimages of each row index can range from zero to four, and this number could be recorded as the matrix element. Let us call such a matrix a *reduced evolution matrix*, denoted by E^N . Then

$$E_{i,j}^N = \text{card}\{a, b \ni \varphi^{-1}(i) = ajb\}$$

For example, choose $N = 3$, and use Rule 22; then,

$$E^3 = \begin{array}{c|cccccccc} & 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \\ \hline 000 & 1 & 1 & . & 3 & . & 1 & 2 & 4 \\ 001 & 1 & 1 & . & . & . & 1 & 2 & . \\ 010 & . & 1 & 1 & . & 1 & . & . & . \\ 011 & . & 1 & 1 & . & 1 & . & . & . \\ 100 & 1 & . & . & 1 & . & 1 & . & . \\ 101 & 1 & . & . & . & . & 1 & . & . \\ 110 & . & . & 1 & . & 1 & . & . & . \\ 111 & . & . & 1 & . & 1 & . & . & . \end{array}$$

The column sums are always 4, but the row sums could range between 0 and 32, according to the rule involved. A zero sum would signify a Garden of Eden state, because the row index would then have no chains from which it could be formed. The high extreme would signify a pathological rule favoring one single sequence, for example Rule 0 mapping all chains into the zero chain. Diagonal elements do not portend chains of period one, but a still life could not exist if all diagonal elements were zero.

Unfortunately the square of E does not correspond to the number of two-generation ancestors producing a given chain even though it arises from two steps of evolution. The reason is precisely the correlations introduced by the order in which cells occur. Squaring E implies *some* ancestor for each step of the evolution, but does not guarantee that they coincide where they overlap. An easy counterexample is obtained from chains of length 1. E_1 for rule 22 is

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix},$$

so its square is

$$\begin{bmatrix} 10 & 9 \\ 6 & 7 \end{bmatrix}.$$

The table for two generation evolution according to Rule 22 is

00000	00001	00010	00011	00100	00101	00110	00111
000	001	011	010	111	110	100	100
0	1	0	1	0	0	1	1
01000	01001	01010	01011	01100	01101	01110	01111
110	111	101	100	001	000	000	000
0	0	0	1	1	0	0	0
10000	10001	10010	10011	10100	10101	10110	10111
100	101	111	110	011	010	000	000
1	0	0	0	0	1	0	0
11000	11001	11010	11011	11100	11101	11110	11111
010	011	001	000	001	000	000	000
1	0	1	0	1	0	0	0

which produces a two-generation reduced evolution matrix

$$\begin{bmatrix} 10 & 11 \\ 6 & 5 \end{bmatrix},$$

which is clearly different from the square of the one-generation matrix.

There is an evident relation between the evolution matrices $G(i, j)$ and the reduced evolution matrices E_i : If we define a matrix R , for which

$$R(3) = \left[\begin{array}{c|cc} & 0 & 1 \\ \hline 000 & 1 & 0 \\ 001 & 1 & 0 \\ 010 & 0 & 1 \\ 011 & 0 & 1 \\ 100 & 1 & 0 \\ 101 & 1 & 0 \\ 110 & 0 & 1 \\ 111 & 0 & 1 \end{array} \right]$$

is typical, we have

$$E_j = G(j : j + 1)R(j + 1)$$

and the evident definition

$$[R(j + 1)]_{p,q} = \delta(p, \text{bin}(q) = apb).$$

In such terms it is clear that $G(1 : 2)G(2 : 3)R(3)R(2) \neq [G(1 : 2)R(2)]^2$ as a general proposition.

Chapter 4

Periods in time

The constraint of a fixed spatial periodicity on a cellular automaton produces an assortment of time periods, into one of which the evolution must eventually fall. The converse problem would be to select a time interval, with the intention of enumerating the lengths of all the possible rings whose evolution would regenerate after the given delay.

4.1 Characteristics of cycles

The course of evolution of any finite automaton can be traced through a diagram constructed for the purpose, or such an equivalent representation as the connectivity matrix of the diagram. Features of interest surely include:

- the cycles in which all evolutions terminate,
- nodes with no ancestors – the Garden of Eden configurations,
- the length of the longest transient,
- the average and standard deviation of transient lengths,
- the amount and type of branching in the diagram.

The first item is the only one of importance for the operation of an automaton after an extremely long time, but it is still useful to understand the short term behaviour, particularly as it affects the choice among different long term alternatives. Furthermore, what is a long term for a ring of a hundred cells is still very short term for a ring of a thousand cells.

The technique of tracing out the full evolution of a ring of length N yields all the configurations in an infinite ring which repeat themselves after translation through a distance N , whatever their period of repetition in time; in other words all the patterns with a given spatial periodicity. In the process of observing the results one learns to recognize certain patterns and to be able to predict their periodicities. It would be nice to have a procedure which would yield all the configurations with a given period, irrespective of the length of the ring on which they might occur,

beginning with a method for finding all the still lifes - patterns which never change with time.

4.2 Overlapping of neighborhoods

Neighborhoods in which the central cell does not change can be read off from the evolutionary rule. The complication in finding sequences which do not change lies in the overlap between neighborhoods of two successive cells. Fortunately there already exists a diagrammatic technique for dealing with overlapping sequences of symbols, which is commonly encountered in the theory of shift registers[44]. The underlying combinatorial theory goes back at least to the last century[98].

To see how this works, let us calculate the still lifes for Rule 22. The table below shows quite a bit more information than this; each transition is classified as to whether the new cell is the same as the old, as its left or right neighbor, whether it has evolved into zero or one, or the complement of its old value. In general the new cell could be tested against any Boolean function of the members of its neighborhood.

neighborhood	000	001	010	011	100	101	110	111
new cell	0	1	1	0	1	0	0	0
same	x	.	x	.	.	x	.	.
left	x	.	.	x	x	.	.	.
right	x	x	x	.
zero	x	.	.	x	.	x	x	.
one	.	x	x	.	x	.	.	.
complement	.	x	.	x	x	.	x	x

Three neighborhoods qualify for producing still lifes, namely 000, 010, and 101. 000 can overlap only with itself in either direction, so that it can only participate in a chain of pure zeroes. Such a chain does in fact remain zero, so it qualifies as a still life. The other two neighborhoods can overlap each other in alternation, leading to the sequence $\dots 010101 \dots$ which is also seen to be a still life. These two alternatives exhaust the possibilities, and answer the question “What are the still lifes of Rule 22?”

There are three neighborhoods whose central cell evolves into one, namely 001, 010, and 100. The first of these can only fit between the third and the second, in that order. The second can only fit between the first and third, and the third must follow the second and precede the first. Thus only the sequence $\dots 100100100100 \dots$ can produce pure ones. On inspection this is reasonable—contiguous ones necessarily force a zero, while a gap of zeroes longer than two can never be filled with ones in a single generation.

Similar arguments can be given regarding the remaining lines in the table. However, something more concrete than such verbal arguments is required, and is afforded by the subdiagrams of a map outlining all the possible ways in which the neighborhoods defining the evolution can overlap.

4.3 The de Bruijn matrix

The basic diagrammatic tool is called a de Bruijn diagram[44, 98], whose k^s nodes represent the possible sequences of s symbols chosen from amongst a collection of k ; making the convenient choice of integers modulo k for symbols, the nodes become simply s -digit numbers to the base k .

Two nodes are to be linked if the first ends with the same $s-1$ digit sequence with which the second begins; in other words, if they represent overlapping sequences. Since there is no restriction on the symbols which can be placed at the free ends, each node will have k incoming links and k outgoing links. Having chosen numbers as symbols, the linkage rule is expressible in a simple arithmetic form; namely, that node i is linked to nodes $si, si + 1, \dots, si + k - 1$. Since there are $S = k^s$ nodes in all, the foregoing sums are to be taken modulo S . Indeed the most straightforward representation of a de Bruijn diagram is through the vertices of an S -gon inscribed in a circle, chords marked according to the links present. It can also be represented by a connectivity matrix whose block diagonal structure models this circle.

Since those nodes are linked which differ by shifting their label, dropping the digit on the left and inserting a new digit on the right, de Bruijn diagrams are frequently used in shift register theory. Here, the shift register is simply a moving window which can be used to scan a long chain for the neighborhoods needed by the evolutionary rule of a cellular automaton. Any chain can be reconstructed from the windows by following a path through the diagram, recording the digit associated with each link as one moves along.

Now we need to examine de Bruijn matrices in detail, but to give concrete examples we will stick to low order matrices over the binary alphabet. Let $B_{k,n}$ denote the connectivity matrix for a de Bruijn diagram for k symbols and n stages; write B_n for $B_{2,n}$.

4.3.1 B_1

Let us start with B_1 , which is a 2×2 matrix (showing a border identifying the 1-blocks which have been linked).




Figure 4.1: Single stage binary de Bruijn matrix and diagram.

The simple structure is due to the fact no matter whether 0 or 1 is dropped from a 1-block, or whether 0 or 1 is added to the block, the maneuver is feasible; thus all positions of the matrix are filled with a one. By inspection, we see that

$$B_1^2 = 2B_1,$$

which is the characteristic equation for B_1 . It has eigenvalues $\lambda = 0, 2$. Since it is symmetric it has an orthogonal eigenvector matrix, which is:

$$\frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & \Leftrightarrow 1 \end{bmatrix}.$$

The coefficient $1/2$ rather than $1/\sqrt{2}$ is chosen because we will presently use probability vectors, which are normed by sums of absolute values rather than sums of squares.

4.3.2 B_2

Let us continue with B_2 , likewise labelling its rows and columns by the 2-blocks to which they correspond. Zero matrix elements—for which there is no link in the de Bruijn diagram—will be suppressed in favor of dots.




Figure 4.2: Two stage binary de Bruijn matrix and diagram.

By a simple calculation,

$$B_2^3 = 2B_2^2,$$

which is a minimal equation rather than a characteristic equation since it is not of degree 4. Its roots are $\lambda = 2, 0, 0$. B_2 is evidently singular because of the repeated rows, but on closer inspection we find that it shows the Jordan canonical form with respect to a 2×2 block belonging to eigenvalue 0.

The matrix of principal vectors reducing B_2 to canonical form could be

$$\frac{1}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & \Leftrightarrow 1 & \Leftrightarrow 1 & 1 \\ 1 & 1 & \Leftrightarrow 1 & \Leftrightarrow 1 \\ 1 & \Leftrightarrow 1 & 1 & \Leftrightarrow 1 \end{bmatrix}.$$

The first three columns are eigenvectors, the last a principal vector.

Although characteristics of the de Bruijn matrices are already evident in B_2 , it is worth going on to examine B_3 for good measure.

4.3.3 B_3




Figure 4.3: Three stage binary de Bruijn matrix and diagram.

This time the minimal equation is

$$B_3^4 = 2B_3^3.$$

The general result is that $B_{k,n}$ is a $k^n \times k^n$ matrix whose minimal equation is

$$B_{k,n}^{n+1} = kB_{k,n}^n,$$

and which exhibits the Jordan canonical form for $k > 2$.

This result shows that it takes at most n links to get from one node to another in an n -stage de Bruijn diagram; in terms of shift register sequences this is precisely the shift required to completely replace one sequence by another. If the node labels have internal symmetry, a smaller shift and hence a shorter path through the diagram may exist.

Although the form of the minimal equation can be conjectured readily enough by inspection, an instructive proof can be given, whose elements will be useful later on. We begin with two factorizations of the de Bruijn matrix, motivated by the observation that it would be block diagonal if only its rows were rearranged.

4.3.4 B_n as product and sum

$$B_3 = \begin{bmatrix} 1 & . & . & . & . & . & . & . \\ . & . & 1 & . & . & . & . & . \\ . & . & . & . & 1 & . & . & . \\ . & . & . & . & . & . & 1 & . \\ . & 1 & . & . & . & . & . & . \\ . & . & . & 1 & . & . & . & . \\ . & . & . & . & 1 & . & . & . \\ . & . & . & . & . & . & . & 1 \end{bmatrix} \times \begin{bmatrix} 1 & 1 & . & . & . & . & . & . \\ 1 & 1 & . & . & . & . & . & . \\ . & . & 1 & 1 & . & . & . & . \\ . & . & . & 1 & 1 & . & . & . \\ . & . & . & . & 1 & 1 & . & . \\ . & . & . & . & . & 1 & 1 & . \\ . & . & . & . & . & . & 1 & 1 \\ . & . & . & . & . & . & . & 1 & 1 \end{bmatrix}$$

Expressing this equation in the form $B_n = P_n U_n$ implicitly defines the orthogonal permutation matrix P and a block diagonal U . These matrices satisfy the equations $P_n^n = I$, $U_n^2 = 2U_n$, where I is the unit matrix. Note that U_n is a Kronecker product, $U_n = I_{n-1} \otimes U_1$, and that the following two identities hold

$$B_n = P_n(I_{n-1} \otimes U_1)$$

$$B_n = (U_1 \otimes I_{n-1})P_n$$

Due to the symmetry of the matrix U , a slightly different factorization is possible, in terms of a permutation matrix which we could call Q_n .

$$B_3 = \begin{bmatrix} \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \\ 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix} \times \begin{bmatrix} 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix}$$

Q_n satisfies the equation $Q_n^{2^n} = I$, but P and Q do not commute. It is interesting to note that

$$B_n = P_n + Q_n$$

Thus there are both factorizations and sum decompositions for the de Bruijn matrix, all readily obtained in a way that indicates that many more representations are possible, although the others would not be as symmetrical as the ones shown.

If we adopted a more formal definition of P and Q we would find

$$[P_n]_{i,j} = \delta(2i, j) \bmod 2^n$$

$$[Q_n]_{i,j} = \delta(2i + 1, j) \bmod 2^n$$

from which would follow

$$\begin{aligned} B_n^2 &= (P_n + Q_n)^2 \\ &= P_n P_n + P_n Q_n + Q_n P_n + Q_n Q_n \end{aligned}$$

Writing the four terms of this sum according to their matrix elements, we find (where δ is Kronecker's delta):

$$\begin{aligned} [P_n^2]_{i,j} &= \delta(4i, j) \bmod 2^n \\ [P_n Q_n]_{i,j} &= \delta(4i + 2, j) \bmod 2^n \\ [Q_n P_n]_{i,j} &= \delta(4i + 1, j) \bmod 2^n \\ [Q_n^2]_{i,j} &= \delta(4i + 3, j) \bmod 2^n \end{aligned}$$

a representation which can readily be generalized for any power of B , and also for any number of states per cell, k . In particular, the form of both the minimal equation and the characteristic equation for the de Bruijn matrices follows. Note that the congruence in these equations is multiplicative, not additive, making 2^N congruent to 1, not 0.

The factorized form of the de Bruijn matrices can be used to obtain the determinants and inverses of their probabilistic generalization; since the strict de Bruijn matrices are singular their zero determinant can be derived in this way, but of course they have no inverses.

4.4 Periods and other properties

There are two ways that a de Bruijn diagram can be associated with the neighborhoods of a linear cellular automaton. The neighborhoods can be attributed to the nodes or to the links. The latter is the more compact representation; then the node between two links would describe the overlapping portion of the two neighborhoods. Any classification of neighborhoods immediately becomes a classification of the links. In searching for still lifes, we would discard all those links which did not stand for a neighborhood whose central cell evolved into itself.

If the remaining links formed chains, they would be chains which remained unchanged from one generation to the next. Unless they closed into loops however, they might dwindle away in each generation. If there were suitable loops, their lengths would determine the lengths of the ring automata in which they could be found. For example, chains of alternating ones and zeroes can only be found in rings of even length according to Rule 22, whereas the zero chain can exist in rings of any length.

Consequently loops often have to be extracted from larger fragments of de Bruijn diagrams, just as they have to be to get cycles from the evolutionary diagrams; in both cases the transients leading into the loops get discarded.

The formation of barriers between macrocells is one exception to the rule that only loops give interesting still lifes. The criterion for such a barrier is that the terminal nodes have a full complement of incoming or outgoing links according to the end of the chain which they occupy.

Another combination which often occurs is that a loop will be connected through a series of links to another loop, but that there is no return path. Thus some pattern may occupy the left side of an infinite chain, then undergo a transition into another pattern which occupies the right side of the chain. Such configurations were called *fuses* in Gardner's column and Wainwright's newsletter, and can exist in one dimension just as well as in two.

Nor are still lifes the only patterns which can be deduced from de Bruijn diagrams. To obtain all the patterns of period two, it is only necessary to iterate the evolutionary rule once and then look for still lifes. This would transform a (k, r) automaton into a $(k, 2r)$ automaton using a de Bruijn diagram of $4r$ stages, whose links would be classified by their behaviour after two generations of evolution. Similar considerations apply to periods of three or higher.

Any binary property at all of a neighborhood of cells can be reflected by its inclusion as a link in the de Bruijn diagram for neighborhoods of that length, and be extended to a larger region by pursuing chains throughout the diagram. Cells other than the central cell can repeat themselves, so that it is just as easy to search for shifting patterns as for periodic ones. Evolution into constant chains is likewise easy to test, and is helpful in determining the behaviour of backgrounds, particularly in automata with many states per cell.

Figure 4.4 shows several examples; in fact configurations possessing any of the attributes mentioned in Section 4.2, and many more, can be read off by inspection, the $(2, 1)$ de Bruijn diagram being quite small.

Two less obvious properties can be mentioned. One is the search for *superluminial* configurations. Quiescent automata are ones in which there is a state q




Figure 4.4: Simple properties of Rule 22 via de Bruijn diagrams.

such that $\varphi(q, q, q) = q$; when such a state exists, it is generally assigned the value zero. Actually, especially for automata which are not binary, there may be several quiescent states, in which case only one of them would be labelled zero. Quiescence simply means that a non-quiet cell cannot arise unless there is another non-quiet cell in its neighborhood, which is to say, “nearby.” In quiescent automata information can propagate through a quiet region at a maximum velocity equal to one radius per time step; this velocity is picturesquely called the *velocity of light*.

A pattern can move faster than the velocity of light if it moves through a region which is densely enough populated by non-quiet cells. Superluminal configurations can always be found for any velocity, but they may turn out to be so trivial as to be uninteresting. For example, one may simultaneously be superluminal with a smaller velocity, or even be a higher harmonic of a subluminal configuration. The argument is that one can choose an arbitrary interval consisting of one superlight-second plus half the width of the neighborhood. Say its length is ℓ . Anything beyond this interval is fixed by the combination of evolution and shifting. Of course, the shifted image which emerges does not have to be part of the original interval, but there is not much else that it can be because there are only 2^ℓ intervals and $2^{\ell-2g}$ images after g generations. So, there has to be *some* interval which generates a recurrent pattern.

Idempotent rules lead to a particularly restrictive type of behaviour, but not all rules are idempotent. Still, one of the boolean questions which can be asked of second generation evolution is whether

$$\varphi(\varphi(a, b, c), \varphi(b, c, d), \varphi(c, d, e)) = \varphi(b, c, d)$$

and consequently whether a given rule admits any configurations which have an idempotent evolution even if all of them do not.

4.5 Superluminal configurations for Rule 22

Conceptually superluminal configurations may be somewhat easier to describe than those of lower velocity. For convenience, consider a .1(2) configuration, which is supposed to have a single phase which shifts by a distance of 2 in each generation. Its evolution would have the symbolic form

$$\begin{array}{cccccccc} \cdot & \cdot & \cdot & \cdot & a & b & c & d & e & f & \cdot \\ \cdot & \cdot & a & b & c & d & e & f & \cdot & \cdot & \cdot \end{array}$$

From this it is evident that a , b , and c are arbitrary, but that necessarily

$$\begin{aligned} d &= \varphi(a, b, c) \\ e &= \varphi(b, c, d) \\ &= \varphi(b, c, \varphi(a, b, c)) \\ f &= \varphi(c, d, e) \\ &= \varphi(c, \varphi(a, b, c), \varphi(b, c, \varphi(a, b, c))) \\ &\dots \end{aligned}$$

The point is that any seed sequence a, b, c determines the entire sequence of cells from some point on, but only eight different binary sequences of three cells are possible. The original seed need not repeat itself, but some seed must repeat within at least eight cells. Thus only eight strings need to be examined to ascertain all the possible .1(2) configurations; for Rule 22 we would find:

	configuration
0 0 0 0	0^*
0 0 1 1 0 0 1 .	$(0011)^*$
0 1 0 1 0 . . .	$(01)^*$
0 1 1 0 0 1 1 .	$(0011)^*$
1 0 0 1 1 0 0 .	$(0011)^*$
1 0 1 0 1 . . .	$(01)^*$
1 1 0 0 1 1 0 .	$(0011)^*$
1 1 1 0 0 1 1 0	$(0011)^*$

These configurations are not very interesting, because they all have simpler descriptions; but they all fulfill the technical description, and are moreover the only configurations which do so.

If one wanted to find .1(3) configurations, the seed would have length 4, and in general a $.p(d)$ configuration would have a seed of length $p+d$ and the superluminal configuration would have length less than 2^{p+d} and have to be manifest within that many cycles of iteration of the above procedure.

The following superluminal configurations (single phase, shift no longer than 12) exist for Rule 22:

.1(4)	11.1	.1(8)	29.1
.1(5)	30.1	.1(9)	132.1
.1(6)	46.1	.1(10)	191.1
.1(7)	31.1	.1(11)	96.1
.1(8)	37.1	.1(12)	83.1

Subluminal configurations can be found by the same procedure, but since the displaced seed overlaps its own neighborhood, consistency requirements must be met. Thus superluminal configurations always exist even if they are degenerate or trivial, but subluminal configurations may fail the additional requirement. In any event, the de Bruijn diagram yields a more systematic procedure for all configurations, and the foregoing discussion is mainly useful as an existence theorem.

4.6 Periods for Rule 22

Whereas the diagram for the cycles of a linear cellular automaton is a collection of “trees rooted in cycles,” its periods are obtained from the de Bruijn diagram which consists of an enormous number of cycles, all interlacing. Still, the length of a ring with a given period is limited by the number of nodes in its de Bruijn diagram, which after g generations is k^{2rg} . For one-generation properties of a (2,1) automaton this number is 4. This means that any still life must be evident in a ring whose length is no longer than 4, and that if there is a still life in a longer ring, it must contain repeated segments whose length is no longer than 4.

There is only one loop of length 4 in a 2-stage binary de Bruijn diagram, formed by the sequence 0011; to require it to represent a still life makes the following demands on the transition rule:

$$\begin{array}{cccccccc} 111 & 110 & 101 & 100 & 011 & 010 & 001 & 000 \\ a & 1 & b & 0 & 1 & c & 0 & d \end{array} .$$

Thus four neighborhoods have their evolution prescribed and for the other four it is arbitrary. For the maximum variety, a , b , c , and d should be chosen opposite to the central cell, yielding Rule 105, but there are altogether sixteen rules for which the sequence $\dots 001100110011\dots$ is a still life. Of course, there may be additional still lifes besides the one which was stipulated; for example if d were chosen to be zero in this example, any number of additional zeroes could be inserted into the sequence besides the two shown, and we would have Rule 104. Then $\dots 0011000001100011\dots$ would also be a still life. At the points where additional zeroes can be inserted, there is no restriction on the exact number.

It is harder to survey the possible periods for a rule than the possible cycles because the lengths of the configurations that have to be checked grow twice as fast with increasing period as with increasing cycle length; thus if it is feasible to survey all cycles up to length 30, the corresponding limit for periods is 15. Furthermore, the diagram for cycles has only simple loops, whereas the period diagrams can have multiple loops, whose branching is somewhat harder to describe.

Shifting configurations are just as easy to determine as are the periodic configurations; let the notation $P.p(d)$ denote a configuration which has period P with p phases after which time it has shifted d cells to the left (a negative d signifies a right shift, which only need be catalogued separately for non-symmetric rules). Reflective configurations of the form $2p.p$ can be found among the $2p.2p(0)$ configurations, the calculation required to detect them being the same in either case.

Actually it is better not to specify the period, since it will depend somewhat on the length N of the ring in which it occurs. That is, if d does not divide N , the pattern may have to make several circuits of the ring before repeating itself, and

$$P = p \times \text{lcm}(d, N).$$

The smaller subluminal periods for Rule 22 have been determined; they include the following up through five generations:

1 generation :

$$\begin{array}{l} .1(0) - 0,* (01)* \\ .1(1) - 0* \end{array}$$

2 generations :

$$\begin{array}{l} .2(0) - 0,* (01)*, (0011)* \\ .2(1) - 0,* (0001101)* \\ .2(2) - 0,* (01)* \end{array}$$

3 generations :

- .3(0) - 0,* (01)*
- .3(1) - 0,* (000010011 11)*
- .3(2) - 0,* Figure 4.5
- .3(3) - 0*

4 generations :

- .4(0) - 0,* (1(0 + 00 + 000)10000(0 + 00 + 000))* and evolves
- .4(1) - 0*
- .4(2) - 0,* (01),* (0001101)*
- .4(3) - 0,* Figure 4.6
- .4(4) - 0,* (01),* (0011)*

5 generations :

- .5(0) - 0,* (01),* (100001000 0(0 + 00 + 000))* and evolves
- .5(1) - 0,* (0001101)*
- .5(2) - 0,* (0000111100 1),* (0011),* (01)*
- .5(3) - 0,* (0000011100 0010001001 1101111000 0000010)*
- .5(4) - 0,* (0001110001 0001011101 1000000010 0),* (01)*
- .5(5) - 0,* (0010011000 0101110001 0000011111)*



Figure 4.5: Configurations satisfying .3(2) for Rule 22.

The unshifted period six is interesting for a wide variety of cycles, all of whose lengths are multiples of ten, situated as vertical faces on the hexagonal prism of Figure 4.7. Of course the table continues indefinitely; shifting configurations of period six exist, not to mention longer periods. Taking into account additional shifting and longer periods tends to produce diagrams of increasing complexity, but three classes of diagram are already apparent in the results shown here.

First, there are simple cycles which evolve by shifting (including still lifes with zero shift); a single diagram serves all generations.

Next, there are diagrams, such as the .4(0) or .5(0) families, in which succeeding generations are described by distinct components of a disconnected diagram, although symmetric images of the first (and subsequent) generations may be encountered before the full period closes.

The third class, typified by the .4(3) configurations of Figure 4.6 or the .6(0) configurations of Figure 4.7, combines several different phases within a single generation, producing a much more intricate single component diagram.

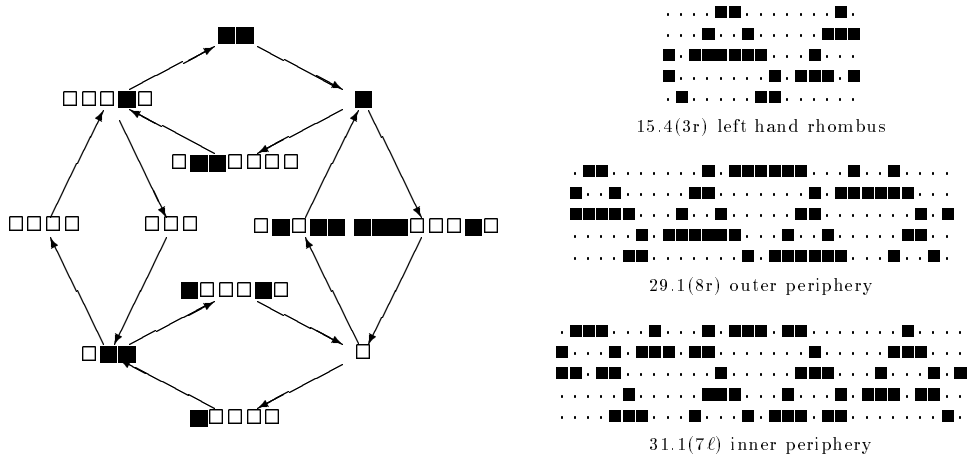


Figure 4.6: Configurations satisfying .4(3) for Rule 22.

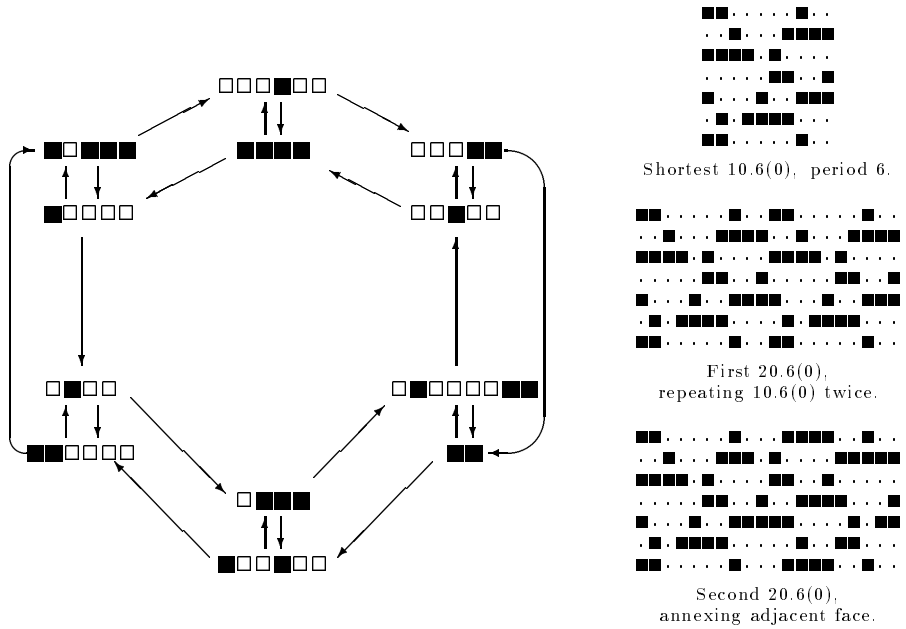


Figure 4.7: Configurations satisfying .6(0) for Rule 22.

The .5(0) configurations arise from a single pair of cells embedded in a torus; two gaps separate them of which one always has length 4, the other may have length 5, 6, or 7. Altogether the cycle length may be 11, 12, or 13; longer configurations arise from joining the triplets in sequence.

By contrast, the .4(0) configurations also feature a live cell pair, but this time the short gap can have length 1, 2, or 3 while the long gap's length can be 5, 6, or 7; this time an ennead whose members can be sequenced.

In summary, the de Bruijn diagram determines sequences of neighborhoods, and therefore cyclic or infinite configurations which meet some requirement which can be placed on its links. Cyclic configurations follow from the loops inherent in the diagram, but aperiodic configurations also exist; these can arise in two ways.

If there are two loops with a one-way connection between them, such as can be seen in Figure 4.5 describing .3(2) configurations, a transition between the two loops can be made at some point. The result in some cases has been called a "fuse," a left hand field arises which is different from the right hand field; the boundary between them may be moving or static. To make a good fuse, the field from which the boundary retreats should be quiescent.

If two loops are mutually connected, the transition back and forth can be made many times; imagine that the choice is made to correspond to the binary expansion of an irrational number to obtain an aperiodic configuration. The greater the variety of loops in the diagram, the greater the variety of configurations, but a connected pair is enough for aperiodicity.

4.7 The gap theorems

In the course of obtaining all the cycles of length up to $N = 34$ for Rule 22, rings of cells were examined in such an order that those with more leading zeroes were examined first. Moreover, rings were discarded if, in the course of evolution, a ring appeared which might have been used earlier as a starting point, or was equivalent to such a ring by some symmetry operation. In effect, that meant if a ring evolved which contained a string with more zeroes than there were leading zeroes in the original ring, the analysis was halted and the next case considered.

It was quickly observed and proved that when a single leading zero was encountered, only one cycle was going to be discovered — the still life if the length of the ring was even. Likewise, two leading zeroes would only lead to the ring of period two, and then only when the length of the ring was a multiple of four. If we say that a sequence of m zeroes (presumably bounded by ones at both ends) is a *gap* of length m , then the conditions in the following table were observed

gap	increases to	unless
1	2	(01)*
2	3	(0011)*
3	4	(0001011)*
4	5	(00100001111)*

A conjecture that the only exceptions to gap 5 were the configurations of period 4 and period 6 was made, but proved intractable to verify. Essentially the conclusion

was drawn that there would always be some phase of a cycle with a larger gap than a certain value unless the ring were of a particular form, generally of low period, in which the given gap was conserved.

These observations had practical importance because the search for cycles could be stopped when the leading gap got to be small enough. The search time doubled each time the leading gap was decreased by one, so that it was worthwhile to stop with the largest possible leading gap.

By the use of the de Bruijn diagram, the low order gap theorems can be derived; the procedure is applicable to any rule in any automaton. It is convenient to change the point of view of the theorems slightly. Rather than worry about "gaps increasing," the question could be: What strings of cells *do not* evolve into gap m ? This property is just one of many predicates which could be used to select subdiagrams from a de Bruijn diagram. For example, which strings of length three do not evolve into zero? Or, which strings of length four do not evolve into 00?

Since a small gap is contained in a larger one (gaps without worrying about boundaries), it must mean that only strictly smaller gaps can occur when a given gap is excluded from the next generation. Rings from which such gaps are excluded can only be made up from surviving loops in the de Bruijn diagram, and hence are even limited in their maximum length.

The simple exclusion of gaps from the second generation does not mean that the gaps were not present in the first generation, or that they might not appear in the third or subsequent generation. So, drop them from the first generation as well, leaving a finite set of candidates whose periods must be determined and which must be followed through their full cycle of evolution to see whether overly large gaps occur at some stage.

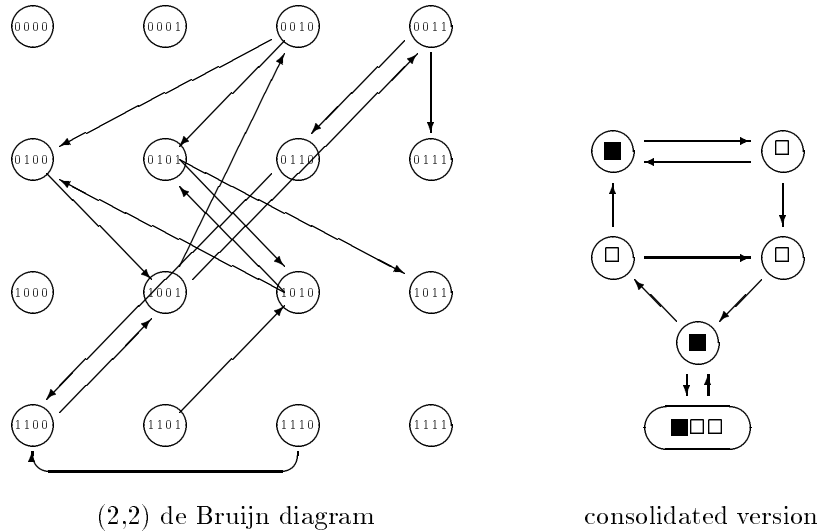


Figure 4.8: (2,2) de Bruijn diagram excluding 000 in first and second generations.

Figure 4.8 shows the process applied to Rule 22 for gap 3; using the table of Section 3.8 all links are excluded which contain three consecutive zeroes in either the first or the second generation; consolidating the results shows that either an additional generation should have been considered (resulting in a larger starting diagram) or some of the known counterimages of three zeroes should also have been dropped (a rather *ad hoc* procedure).

The elimination of a gap from a sequence is similar to the construction of Cantor's set, wherein the digit '1' is excluded from triadic decimals constructed from the digits $\{0, 1, 2\}$; but clearly additional and fancier exclusions are possible. The difference lies in extending the exclusion to several, and possibly all, generations; the latter requires skill, persistence, and luck.

Carrying a diagram through additional generations generally reduces the number of configurations which are admissible; the triangular bridge lying between the cycles $(01)^*$ and $(0011)^*$ disappears in the next generation, leaving results which are provably applicable to all future generations. The process works well because these cycles lack ancestors other than themselves; but it is another matter to arrive at gap 5 which can be found in all Rule 22's period four configurations, all of which have large numbers of ancestors with small gaps.

Useful results have been obtained for practically important cases with small gaps, but the theorems are slightly more complicated if aperiodic configurations are included: there are some interesting "black hole" states. Also to be borne in mind is the fact that for symmetrical rules such as Rule 22, symmetric images of exceptions are also exceptions, as we have tacitly assumed.

gap	nonincreasing configurations	delay
0	0^*	0
1	$(01)^*$	1
2	$(0011)^*$	3
3	$(0011)^*(0001011)^*$	5
3	$(0001011)^*(1100)^*$	5
3	$(1101000)^{3n}(1100)^{7n}0(0001011)^{3n}$	$21n$
4	$(11110000100)^*$	4
4	$(00100001111)^*$	4
4	$(1101000)^*(1100)^*11(0001011)^*$	

Chapter 5

The Garden of Eden

The de Bruijn diagram is particularly useful for the purpose of extending the properties of a neighborhood in a cellular automaton to the properties of chains of cells because they give a particularly graphical way to deal with the overlapping of neighborhoods. Links in the diagram correspond to neighborhoods, nodes to the intervals by which they overlap, and paths through the diagram either to extended neighborhoods or to the sequence of cells into which they evolve.

The de Bruijn diagram differs from the evolutionary diagram in one important respect - it consists wholly of cycles and there are as many links leading into a node as emerging from it; a number which is uniform for each node. Of course this fine balance may be upset when a subdiagram of the de Bruijn diagram is chosen, and one of the important problems may consist in locating the loops which have survived the pruning which created the subdiagram.

5.1 The subset construction

There is another way in which an imbalance may occur. Let us suppose that each link is labelled by the state into which its neighborhood evolves. A path made of such links has evolved from some state of the automaton, namely the one formed by reading off the sequence of overlapping neighborhoods rather than the sequence of evolved cells. We may invert this process. Suppose that we are given an arbitrary sequence of states. If we can locate it somewhere in the diagram, then we have found at least one extended neighborhood which is its ancestor. Sometimes we can locate the sequence along several paths; but it might also occur that such a path is nowhere to be found.

Although the number of links emanating from each node is the same, it is up to the particular rule what the mixture of labels will be. For example, Rule 0 would show a label of zero on *all* links, with a one nowhere to be found. Which is correct; it is a rule under which no ones can evolve. It is a dangerous situation if there is an imbalance at a certain node, so that some state does not label a link there and some other state labels more than one link. If one is ever obliged to use that node in forming a chain of evolved cells, the combination in which the missing cell

is the next one required will be impossible. It may or may not happen that the desired sequence can be located somewhere else in the diagram, so we need an accurate determination of just which sequences can be folded up and fitted into the diagram.

Moore's subset construction[84] fulfills this requirement very nicely. The basic idea is that if we were systematic, we would start with the first cell of our evolved sequence and look through the diagram taking note of all the links matching that cell. They would emanate from a subset of nodes, so we could make a new diagram in which that subset was a node. The nodes to which the acceptable links lead form a new subset, which we write down and connect to the first subset by a single link with the common label. Turning to the second cell of the evolved sequence, we would examine only those nodes in the second subset node to see which new subset we might form and link with the second cell as label. It would be pointless to examine more nodes in the de Bruijn diagram, because they wouldn't even make the first cell in the evolution come out right.

The final result would be a network of subsets linked and labelled with the same labels used in the de Bruijn diagram. There is no reason not to complete the diagram showing other subsets and other links, but there are exponentially many subsets and they may not all be required if the ancestor of only one chain is sought. The difference is that the nodes and links would now cover all the possibilities in a systematic fashion rather than after the alternative laborious search. In effect the search has already been made, but in an orderly way. Note that in the complete subset diagram, every node has exactly one emerging link for each of the k states of an individual cell; thus one can never reach an impasse seeking a route through the subset diagram.

What one may find instead is that one has arrived at the empty set, which is simply a polite way of saying that the search has been fruitless.

The subset diagram serves a dual purpose. By avoiding the empty set, it is possible to obtain all the possible configurations of the second generation; since the diagram is finite and sequences of cells can be arbitrarily long, the diagram necessarily contains loops, just as does the de Bruijn diagram. By seeking out the empty set, the class of sequences is found which do not have ancestors, and which therefore belong to the Garden of Eden. By examining the final approaches to the empty set, but mainly by rejecting loops, sequences may be found which spoil any other sequence in which they appear, and thus are the basic excluded words for their rule.

5.2 Excluded states for Rule 22

As an example, consider Rule 22. For brevity, label the nodes A, C, D, F. They are linked by zeroes and ones according to the following table:

node	0 leads to	1 leads to
<i>A</i>	<i>A</i>	<i>C</i>
<i>C</i>	<i>D</i>	<i>F</i>
<i>D</i>	<i>C</i>	<i>A</i>
<i>F</i>	<i>D, F</i>	<i>(none)</i>

The links are not evenly distributed between zeroes and ones; indeed there are five zeroes and three ones so there just couldn't be a balance. In this case the discrepancy shows up at node F , which will surely cause trouble.

Figure 5.1 shows the connections between the subsets of the initial subset $\{A C D F\}$, and the resulting subset diagram. Note that the penultimate row, containing the unit classes, almost reflects the de Bruijn diagram, except where the imbalance in links results in a link descending to the empty set or rising to a higher level (which can only be to the row of doublets, for a binary rule).

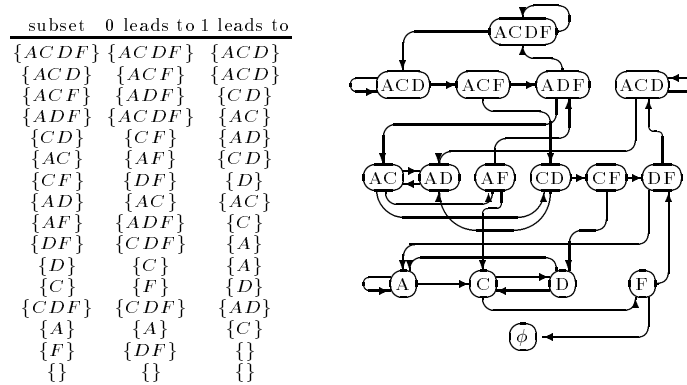


Figure 5.1: The subset diagram (exit links) for Rule 22.

It seems that all sixteen subsets have to be used to complete the analysis. It also appears that both 10101001 and its reverse, 10010101, lead directly to the null set, so they are both excluded words; the shortest ones that can be found for Rule 22.

5.3 Symbolic equations

Diagrams can be represented by connection matrices, through which many properties of the diagram can be deduced more readily than from the diagram itself. There is another representation, in terms of regular expressions, which is often useful. To see how this works, let us try to express the diagram through a set of symbolic equations. Let nodes be represented by capital letters, links by small letters. Then the two-stage de Bruijn diagram for Rule 22 can be written in the form

$$\begin{aligned}
 A &= A0 + D1 + \lambda \\
 C &= A1 + D0 + \lambda \\
 D &= C1 + F0 + \lambda \\
 F &= C0 + F0 + \lambda
 \end{aligned}$$

In such an equation, “products” represent concatenation and are usually written without any explicit operator sign, and “sums” represent alternative choices and may be read as “or.” The symbol λ represents a null string and acts like a one in

products. It is included in the equations to signify that one can get to a node by doing nothing; that is, by following a null link.

A system of equations is “solved” in much the same way that a system of algebraic equations would be solved. The distributive law holds, and any variable (the nodes) may be substituted by an equivalent expression. The process of elimination would eventually leave a single variable defined in terms of constants, following which the whole chain of substitutions could be unravelled to obtain the values of all the variables.

There is just one technical point, however, which concerns the procedure to be followed when the same variable occurs on both sides of an equation, as in

$$X = Xa + b$$

An attempt to write

$$X(\lambda \Leftrightarrow a) = b$$

and thence

$$\begin{aligned} X &= b(\lambda \Leftrightarrow a)^{-1} \\ &= b(\lambda + a + aa + aaa + \dots) \end{aligned}$$

is very suggestive but hardly justified in a system which lacks any trace of subtraction, division, infinite series, convergence, and the like.

The correct procedure is to set about a series of substitutions

$$\begin{aligned} X &= Xa + b \\ &= (Xa + b)a + b \\ &= ((Xa + b)a + b)a + b \\ &\dots \end{aligned}$$

and observe the general form of the result

$$X = Xa \dots aaa + b(\dots aaa + aa + a + \lambda).$$

It is the first term which refuses to disappear; in the context of arithmetic, multiplication, and a numerical a less than 1, it could be argued that its magnitude eventually becomes arbitrarily small and hence it can safely be neglected. Here the argument might run along the lines, that if the symbolic a is not the null chain, λ , then the first term becomes arbitrarily long and can safely be ignored if we are seeking *finite* solutions to the symbolic equation.

The case $a = \lambda$ is special, because it really doesn't define anything; any value of X whatsoever, including b , will work. On the other hand, if b is empty and a is not null, the solution must be empty because otherwise we would be requiring X to be equal to something strictly longer than itself. That, however, is a different matter.

5.4 Arden's lemma

If this point of view is adopted, then we have three rules of inference.

- two expressions equal to a third are equal to each other
- an expression may be substituted for an equal expression
- the solution to $X = Xa + b, a \neq \lambda$ is $X = ba^*$

The notation a^* stands for the continued alternative $\lambda + a + aa + aaa + \dots$, and is read "a star." Any symbolic expression constructed by the aid of concatenation, alternative selection, and the star operation, is called a *regular expression*. Regular expressions are ideally suited to describing paths through a diagram, and conversely, any regular expression has a diagrammatic representation. Conway's monograph[25] is a concise reference for regular expressions and their properties.

If we apply these principles to the de Bruijn diagram for Rule 22 above, we find the following chain of substitutions

$$\begin{aligned} A &= (D1 + \lambda)0^* \\ F &= (C0 + \lambda)0^* \\ C &= D(10^*1 + 0) + 0^*1 + \lambda \\ D &= C(1 + 00^*0) + 0^*0 + \lambda \\ C &= C(1 + 00^*0)(0 + 10^*1) + (00^* + \lambda)(10^*1 + 0) + \lambda)^*1 + \lambda \end{aligned}$$

Finally we have

$$C = ((1 + 00^*0)(0 + 10^*1)) * ((00^* + \lambda)(10^*1 + 0) + \lambda)^*1 + \lambda)$$

of which the leading term is significant and the remainder represent transients.

Although we have set up the equations for the evolution of one generation of a (2,1) automaton for the special case of Rule 22, in fact the equations have the same form for all rules. Given the general evolution table

$$\begin{array}{cccccccc} 111 & 110 & 101 & 100 & 011 & 010 & 001 & 000 \\ h & g & f & e & d & c & b & a \end{array}$$

the equations are, in symbolic form,

$$\begin{aligned} A &= Aa + Dh + \lambda \\ C &= Ab + Dd + \lambda \\ D &= Cc + Fg + \lambda \\ F &= Ce + Ff + \lambda \end{aligned}$$

with the general solution

$$\begin{aligned} A &= Dha^* + a^* \\ F &= Cef^* + f^* \\ C &= ((f^*g + \lambda)(d + ha^*b) + a^*b + \lambda)((c + ef^*g)(d + ha^*b))^* \\ D &= ((a^*b + \lambda)(c + ef^*g) + f^*g + \lambda)((d + ha^*b)(c + ef^*g))^* \end{aligned}$$

For symmetric rules we would have $b = h$ and $e = g$. If 0 were a quiescent state, we would have $a = 0$.

For infinite or cyclic chains the interesting part of these solutions would be the starred right hand expression common to the four nodes. In essence the remainder of the expression tells how to get to the designated node from any other, while the main term tells how to keep returning to that node. For purposes of reference we could tabulate the results for all the symmetric, quiescent rules for $(2, 1)$ automata. These are just Wolfram's "legal" rules.

5.5 The use of symbolic equations

Unfortunately the scheme which solves systems of symbolic linear equations tends to multiply the length of the solution by a constant factor, equal to the number of alternatives, for each node that has to be resolved by the star formula. Since the symbolic coefficients do not tend to simplify in general, symbolic solutions become impractical for all but the very smallest diagrams. We might also note that the table above describes a considerable variety of automata, and yet they are uniformly described by formulae of the same general appearance. Many of them can be further simplified, in spite of the fact that most cannot. For example the expression describing Rule 0 transforms into 000^* ; the descriptor of Rule 4, for which 11 is an excluded word, simplifies to $(100^* + 0000^*)^*$.

Some additional remarks ought to be made concerning symbolic equations. The equations which we have shown are *entrance equations*, in the sense that $C = A1 + D0 + \lambda$ says that you enter node C via link 1 from node A , via link 0 from node D or you do nothing. But it is also true that $C = 1D + 0F + \lambda$, wherein you leave node C via link 1 and continue from D , or else leave via link 0 and continue from F . The complete set of conditions would constitute a description of the de Bruijn diagram via *exit equations*, and can be equally solved by regular expressions.

Using the same example as before, the exit equations for Rule 22 would take the form

$$\begin{aligned} A &= 0A + 1C + \lambda \\ C &= 1D + 0F + \lambda \\ D &= 1A + 0C + \lambda \\ F &= 0D + 0F + \lambda \end{aligned}$$

For many purposes, the entrance equations and the exit equations of a diagram give equivalent descriptions. However, the path running between the full set and the empty set which corresponds to the excluded words for a rule ought to be defined by entrance equations rather than exit equations. After all, interest lies in how an impasse is arrived at, not in what to do once it has occurred.

We have not explained why λ has been incorporated as an alternative in each equation. Were it omitted, the equations would be shorter and thereby easier to solve. However the final equation for Rule 22 would then read $C = C(1 + 00^*0)(0 + 10^*1)$ whose solution according to Arden's lemma would be the empty set—that is,

Rule	a	b	c	d	e	f	g	h	$((c + ef^*e)(d + h0^*h))^*$
0	0	0	0	0	0	0	0	0	$((0 + 00^*0)(0 + 00^*0))^*$
4	0	0	1	0	0	0	0	0	$((1 + 00^*0)(0 + 00^*0))^*$
18	0	1	0	0	0	0	0	1	$((0 + 00^*0)(0 + 10^*1))^*$
22	0	1	1	0	0	0	0	1	$((1 + 00^*0)(0 + 10^*1))^*$
32	0	0	0	1	0	0	0	0	$((0 + 00^*0)(1 + 00^*0))^*$
36	0	0	1	1	0	0	0	0	$((1 + 00^*0)(1 + 00^*0))^*$
50	0	1	0	1	0	0	0	1	$((0 + 00^*0)(1 + 10^*1))^*$
54	0	1	1	1	0	0	0	1	$((1 + 00^*0)(1 + 10^*1))^*$
72	0	0	0	0	1	0	1	0	$((0 + 10^*1)(0 + 00^*0))^*$
76	0	0	1	0	1	0	1	0	$((1 + 10^*1)(0 + 00^*0))^*$
90	0	1	0	0	1	0	1	1	$((0 + 10^*1)(0 + 10^*1))^*$
94	0	1	1	0	1	0	1	1	$((1 + 10^*1)(0 + 10^*1))^*$
104	0	0	0	1	1	0	1	0	$((0 + 10^*1)(1 + 00^*0))^*$
108	0	0	1	1	1	0	1	0	$((1 + 10^*1)(1 + 00^*0))^*$
122	0	1	0	1	1	0	1	1	$((0 + 10^*1)(1 + 10^*1))^*$
126	0	1	1	1	1	0	1	1	$((1 + 10^*1)(1 + 10^*1))^*$
128	0	0	0	0	0	1	0	0	$((0 + 01^*0)(0 + 00^*0))^*$
132	0	0	1	0	0	1	0	0	$((1 + 01^*0)(0 + 00^*0))^*$
146	0	1	0	0	0	1	0	1	$((0 + 01^*0)(0 + 10^*1))^*$
150	0	1	1	0	0	1	0	1	$((1 + 01^*0)(0 + 10^*1))^*$
160	0	0	0	1	0	1	0	0	$((0 + 01^*0)(1 + 00^*0))^*$
164	0	0	1	1	0	1	0	0	$((1 + 01^*0)(1 + 00^*0))^*$
178	0	1	0	1	0	1	0	1	$((0 + 01^*0)(1 + 10^*1))^*$
182	0	1	1	1	0	1	0	1	$((1 + 01^*0)(1 + 10^*1))^*$
200	0	0	0	0	1	1	1	0	$((0 + 10^*1)(0 + 00^*0))^*$
204	0	0	1	0	1	1	1	0	$((1 + 10^*1)(0 + 00^*0))^*$
218	0	1	0	0	1	1	1	1	$((0 + 10^*1)(0 + 10^*1))^*$
222	0	1	1	0	1	1	1	1	$((1 + 10^*1)(0 + 10^*1))^*$
232	0	0	0	1	1	1	1	0	$((0 + 0f^*1)(1 + 00^*0))^*$
236	0	0	1	1	1	1	1	0	$((1 + 0f^*1)(1 + 00^*0))^*$
250	0	1	0	1	1	1	1	1	$((0 + 10^*1)(1 + 10^*1))^*$
254	0	1	1	1	1	1	1	1	$((1 + 10^*1)(1 + 10^*1))^*$

Table 5.1: Paths through the de Bruijn diagrams of Wolfram's "legal" Rules.

no solution. After all, the equation requires that C be equal to something strictly longer than itself. Nevertheless the “infinite recursion” $C = ((1 + 00^*0)(0 + 10^*1))^*$ is just the part of the solution we extracted to represent a cyclic or infinite state for this automaton.

There is evidently a philosophical issue involved, concerning the differences between our conceptions of finite and infinite. However, it is not a particularly difficult philosophy. Had we been more selective in the introduction of the λ 's, we would have paid more attention to the fact that they represent a starting point or a stopping point, according to the handedness of the equations. By including a λ in each of the entrance equations, we have implied that any node in the diagram could be a starting point, and thus we should be prepared to see that the cyclic part of the equation for each node is preceded by the transients leading into it from the other nodes.

If we are looking for finite paths using entrance equations, only the initial nodes should be provided with λ 's and the equations should be solved for the terminal nodes; not only will the equations be simpler, but the results will be more accurate. For exit equations, the situation is reversed; the terminal nodes get their λ 's in preparation for solving for the initial states.

If we are genuinely interested in infinite solutions, we can entertain alternatives to Arden's lemma, but we should beware that the result may depend upon the node which we have chosen for the final equation; it will only express a solution involving that node. It often happens that a diagram contains a loop which makes a transition to a second loop, but that there is no return path to the first loop. In de Bruijn diagrams for cellular automata, such a configuration corresponds to one of Wainwright's fuses. If the equations are solved for a node in the first loop, the existence of the second loop will not be evident.

The algebra of regular expressions is not as direct as one would like, and often questions involving regular expressions are most easily resolved by constructing an equivalent diagram, transforming the diagram, and converting the result back into a regular expression. For example, regular expressions form a Boolean algebra, but the definition of a regular expression only involves unions—the regular expression sum. Intersections, which represent words conforming to several regular expression descriptions, and complements, which are words which in no way conform to the description, are readily defined in terms of diagrams, but not by operations directly on the regular expression itself. For example, excluded words are obtained through the subset construction, even though they are then readily described by regular expressions.

A related problem is the one of determining the equality of two regular expressions, and is the reason that there is no canonical form for regular expressions. Equality can be determined by constructing equivalent diagrams, reducing them to their simplest form and comparing the results; however there is no unique regular expression corresponding to any given diagram.

Another result for which there is no simple calculus is to obtain the regular expression which results from the evolution through one or more generations of a given regular expression. The difficulty clearly lies in accounting for all the overlapping which has to be taken into account where sums are involved. It is

easier to form the composite rule corresponding to the desired evolution and solve the equations converting the de Bruijn diagram of the composite rule into a regular expression.

Although each generation of evolution of an automaton can be described by a regular expression, the tendency is for the expression to become more and more complicated with each increasing generation. Likewise the description of the words excluded from each generation becomes increasingly complex as the exclusions from the previous generation have to be taken into account. One may naturally wonder what kinds of limit there may be to all this increasing complexity, and indeed whether the limit may be in some cases simpler than all the confusion which preceded it.

Generally speaking, the limit of a sequence of regular expressions need not be a regular expression. One has only to recall that balanced parentheses are not described by a regular expression, although parentheses nested to any finite depth can be listed explicitly and thus described by a regular expression. Here, then, is a limit which is conceptually far simpler than its approximations, but for which there are approximations more than adequate for practical purposes.

There is something qualitatively different between keeping a huge list of acceptable parenthesis sequences to check parentheses, and a huge list of binary numbers to check the regular expression $(0+1)^*$; the latter is simply unnecessary because a simpler procedure suffices. In our present situation we might find that we have to let a cellular automaton run for a very long time to find out something that another approach might discover much more quickly. Nevertheless, it is still interesting to know what the possibilities are, or when such a situation might be occurring.

Certain transformations can be carried out on cellular automata themselves which can be used to clarify their limiting behaviour. These transformations mostly go in the direction of showing how cellular automata can be used to simulate Turing machines, from whence many known results concerning computability and decidability can be passed over to cellular automata. These results in turn affect the possible types of limiting behaviour which may be encountered.

5.6 Systems of symbolic equations

In Conway's monograph[25] it is shown that the premise of Arden's lemma, that the symbolic equation $X = aX + b$ can be solved by $X = a^*b$, is just as applicable to systems of equations if they are written in matrix notation. Thus if we define the following vectors and matrices

$$\mathbf{X} = \begin{bmatrix} P \\ Q \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} u \\ v \end{bmatrix}$$

$$\mathbf{a} = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

we can write one single matrix equation

$$\mathbf{X} = \mathbf{a}\mathbf{X} + \mathbf{b}$$

whose solution would be

$$\mathbf{X} = \mathbf{a}^*\mathbf{b}.$$

Of course, such a solution is not of much value unless there is an effective way to calculate the “star” of a matrix. Fortunately the matrix elements of the star are simple regular expressions of the elements of the matrix being starred, expressing in a concise form the results of the chain of substitutions that would otherwise be carried out explicitly each time a system of symbolic equations was to be solved. Indeed, the easiest way to establish the point is to go ahead and solve the equations: Given

$$\begin{aligned} P &= aP + bQ + u \\ Q &= cP + dQ + v \end{aligned}$$

we deduce in succession

$$\begin{aligned} P &= a^*bQ + a^*u \\ Q &= d^*cP + d^*v \\ P &= a^*bd^*cP + a^*bd^*v + a^*u \\ &= (a^*bd^*c)^*(a^*bd^*v + a^*u) \\ Q &= d^*ca^*bQ + d^*ca^*u + d^*v \\ &= (d^*ca^*b)^*(d^*ca^*u + d^*v) \end{aligned}$$

Disentangling u and v from the final equations for P and Q yields a formula for \mathbf{a}^*

$$\mathbf{a}^* = \begin{bmatrix} (a^*bd^*c)^*a^* & (a^*bd^*c)^*a^*bd^* \\ (d^*ca^*b)^*d^*ca^* & (d^*ca^*b)^*d^* \end{bmatrix}.$$

Although these results have been stated for a 2×2 matrix, the fact that the matrix elements can be matrices themselves means that the results are valid in terms of submatrices. By repeated partitioning, in principle a system of any size can be reached. In practice, each partition multiplies the complexity of the symbolic expressions involved, quickly reaching unmanageable proportions. Alternative forms for these matrix elements can be derived using some of the identities satisfied by regular expressions, but the difficulty is a fundamental one—the formulas required are intrinsically complicated.

5.7 Ancestorless states for Rule 18

Rule 18, whose statistical properties have been extensively studied by Grassberger, is very similar to Rule 22, except for the fact that single cells die out rather than remaining to affect further evolution. The consequence is that isolated cells form growth clusters whose interiors die out, just as for Rule 22. However, the fate of colliding frontiers from two different clusters depends on the parity with which they collide, and as a result the evolution of Rule 18 is largely characterized by the conflict of neighboring domains of opposite parity. Occasionally a domain will shrink and disappear, allowing its two neighbors to coalesce. Grassberger was interested in the random walk involved.

The following script shows the subset machine for Rule 18.

```
(2,1) Rule #18
```

```
Like LIFE but single cells die
```

```
(00*(0+10*1))*
```

```
Transitions from original machine
```

```
Node A: Link 0:(A ) Link 1:(C )
```

```
Node C: Link 0:(D F ) Link 1:(% )
```

```
Node D: Link 0:(C ) Link 1:(A )
```

```
Node F: Link 0:(D F ) Link 1:(% )
```

```
Transitions from subset machine
```

```
Node ADF: Link 0:(ACDF ) Link 1:(AC )
```

```
Node CDF: Link 0:(CDF ) Link 1:(A )
```

```
Node A: Link 0:(A ) Link 1:(C )
```

```
Node DF: Link 0:(CDF ) Link 1:(A )
```

```
Node : Link 0:( ) Link 1:( )
```

```
Node C: Link 0:(DF ) Link 1:( )
```

```
Node AC: Link 0:(ADF ) Link 1:(C )
```

```
Node ACDF: Link 0:(ACDF ) Link 1:(AC )
```

The next script shows the symbolic equation set of the subset machine for Rule 18. (The nodes have been relabelled according to $a = \{ACDF\}$, $b = \{AC\}$, $c = \{ADF\}$, $d = \{C\}$, $e = \{DF\}$, $f = \{CDF\}$, $g = \{A\}$, $h = \{\}$)

```
a=a.0+c.0+%
```

```
b=a.1+c.1
```

```
c=b.0
```

```
d=b.1+g.1
```

```
e=d.0
```

```
f=e.0+f.0
```

```
g=e.1+f.1+g.0
```

```
h=d.1
```

Since there are relatively few equations, we can write their symbolic coefficient matrix fairly easily. By partitioning it into submatrices certain characteristics of the equations are emphasized; namely that certain groups of nodes are self-contained. Links lead into the group but there are no links leading out again. Figure 5.2 shows the coefficient matrix (assuming the nodes to form a column vector) and the part of the subset diagram derived from the full subset.

	a	b	c	d	e	f	g	h
a	0	1
b	.	.	0	1
c	0	1
d	0	.	.	1
e	0	1	.
f	0	1	.
g	.	.	.	1	.	.	0	.
h




Figure 5.2: Subset diagram for Rule 18.

The solutions for these equations can be obtained in a way which emphasizes the triangular nature of the coefficient matrix:

$$\begin{aligned}
 b &= (0*1)(000*1+01)* \\
 d &= b.1((01+000*1)0*1)* \\
 h &= d.1 \\
 g &= d.(01+000*1)0* \\
 f &= d.000* \\
 e &= d.0 \\
 a &= b.000*+0* \\
 c &= b.0
 \end{aligned}$$

The coefficients of b and d can be simplified slightly. The new coefficient $0*1(00*1)^*$ means that (reading from left to right) after meeting the first 1, we can remain on the highest level only if 1's do not come in joined pairs. Having met the first such pair, we may now remain on the second level as long as the sequence $00*10*1$ repeats itself, which it will do unless *three* 1's come along in sequence, or else a pair of pairs enclosing a single isolated 1. Descent to level three ensues, finally leaving a sequence without an ancestor. Thus three ones in sequence is a sure mark of a Garden of Eden sequence, but the other condition is much less specific, as it does not refer to some sequence of fixed length.

5.8 Factors

Regular expressions form an algebraic system in the same way that many other structures, such as groups, semigroups, rings and fields; they probably resemble

semigroups more than any of the others. Consequently there are standard aspects of structure theory which may be brought to bear, including the study of equivalence relations, order relations, and mappings which are consistent with the regular operations of sum, concatenation and star. Conway's monograph touches on all these points, without explicitly mentioning them. His theory of "factors" is related to the decomposition of an algebraic structure into ideals, with implications for the establishment of a canonical form for regular expressions.

In the case of Garden of Eden configurations, it is noticed that the mere presence of certain strings in a sequence of cells will prevent the whole sequence from having an ancestor; for instance 10101001 is sufficient to block the formation of ancestors for Rule 22. Tracing back from the empty set in the subset diagram, it is seen that all the strings lacking ancestors have to end with 01. A question which can be asked is whether the lack of ancestors is always a local matter, for which definite strings can be assigned responsibility, or whether difficulties can build up more slowly, with the possibility of being reversed before they become inescapable. The presence of loops in the paths leading from the universal set to the empty set allow for a wide variety of ancestorless chains, included among which there is a finite collection of loopless paths, whose maximum length is (exponentially) bounded.

Chapter 6

The calculus of regular expressions

After having watched the evolution of a variety of cellular automata and developed a certain amount of theory regarding their behaviour, it is not only possible to begin to predict the behaviour of a previously unfamiliar automaton, but one also begins to feel that certain desired effects could be obtained if only certain changes were made in the rules of evolution. This is the beginning of automaton engineering, wherein automata are designed to order; at least within the limits of certain constraints. To do this job well, it is necessary to continue analyzing the structure of automata, particularly along the lines of formal structure theory, possible mappings between automata, and the possible embedment of one automaton in another.

6.1 Derivatives

Regular expressions are the natural formalism arising from an attempt to devise solutions to the symbolic equations encountered in trying to trace out all the possible paths through a network. They comprise one of three productive viewpoints; the other two being, first, the diagrammatic or graphical representation of the network, and second, a relatively numerical description of the network via its connectivity matrix.

The symbolic equations defining regular expressions are soon seen to be the simplest examples of a whole hierarchy of symbolic equations, both implicit and explicit, with various restrictions on the precise form which the system is permitted to assume. Just because they lie at the bottom of the hierarchy, it is worthwhile to make an exhaustive study of their properties, which will surely be encountered again and again in any more comprehensive theory.

Having obtained regular expressions as the description of the paths through a diagram, the converse question arises as to whether there is a diagram for an arbitrarily given regular expression. If such there were, it would be a diagram relating how one regular expression differed from another by its initial letter, in the case the regular expression were solving exit equations, or by its terminal letter, in

the case of solutions to entrance equations. Since the two cases are symmetric, it suffices to study one of them—which might as well be the case of prefixes.

One of the charming aspects of this theory is its formal resemblance to the differential calculus of analysis. J. Brzozowski developed this point of view, explicitly incorporating the notation of partial derivatives, in an article[17] published in the *Journal of the Association for Computing Machinery* in 1961. The basic concept is the following: Let $\Sigma = \{a, b, c, \dots\}$ be a set of letters, Σ^* be the set of all finite sequences of letters from Σ , and Ω be some other subset of Σ . Perhaps it could be the set of words represented by a given regular expression, but this is not necessary for the definition.

The derivative of Ω with respect to a is defined to be the set

$$\frac{\partial \Omega}{\partial a} = \{\omega \in \Sigma^* \mid a\omega \in \Omega\}$$

In other words, it is the set of tails of the words in Ω which begin with the letter a .

Sets are the derivatives of sets; since regular expressions define sets it remains only to verify that their derivatives are also described by regular expressions, to be able to talk directly about the derivatives of regular expressions. The easiest way to attack the problem is to resort to the axiomatic definition of regular expressions, thereby obtaining a list of rules defining the derivatives of composites in terms of the derivatives of their constituents. It is at this point that the formal resemblance to differential calculus becomes apparent, evidently being thereby responsible for the terminology.

A functional notation is convenient; for example $f(a, b, c, \dots)$ might be a regular (or other) expression built up explicitly from variables a, b, c, \dots , all belonging to Σ , using specific operations or further functions similarly defined. At one point it is required to know whether the null word λ could be represented by f ; the easiest way to find out is to evaluate f with zero (empty set, not null word; they're different) arguments. Let us designate this particular value of f by $[f]$:

$$[f(a, b, c, \dots)] = f(\phi, \phi, \phi, \dots)$$

Then, given Σ as above, λ null, and ϕ zero, we have:

- definition of a regular expression
 1. λ is a regular expression
 2. ϕ is a regular expression
 3. $a \in \Sigma$ is a regular expression
 4. if x and y are regular expressions, so is xy
 5. if x and y are regular expressions, so is $x + y$
 6. if x is a regular expression, so is x^*
- regular expression whose arguments are zero
 1. $[\lambda] = \lambda$

2. $[\phi] = \phi$
3. $[a] = \phi$
4. $[xy] = [x][y]$
5. $[x + y] = [x] + [y]$
6. $[x^*] = \lambda$

- derivative of a regular expression

1. $\frac{\partial \lambda}{\partial a} = \phi$
2. $\frac{\partial \phi}{\partial a} = \phi$
3. $\frac{\partial a}{\partial a} = \lambda$
4. $\frac{\partial b}{\partial a} = \phi \quad (a \neq b)$
5. $\frac{\partial xy}{\partial a} = \frac{\partial x}{\partial a}y + [x]\frac{\partial y}{\partial a}$
6. $\frac{\partial (x+y)}{\partial a} = \frac{\partial x}{\partial a} + \frac{\partial y}{\partial a}$
7. $\frac{\partial x^*}{\partial a} = \frac{\partial x}{\partial a}x^*$

Note that $[x]$ is required in the “product” rule because we must devote our attention to the second “factor” if the first is the null word; otherwise the first “factor” is long enough to yield a derivative without regard to what follows it.

- canonical form

$$f(a, b, c, \dots) = [f] + a\frac{\partial f}{\partial a} + b\frac{\partial f}{\partial b} + c\frac{\partial f}{\partial c} + \dots$$

The evident way to get a diagram corresponding to a regular expression f is to start with a node labelled f , and to connect it to a node labelled $\frac{\partial f}{\partial a}$ via an outgoing link labelled a , to a node labelled $\frac{\partial f}{\partial b}$ via an outgoing link labelled b , and so on. The canonical form guarantees that we have the correct set of exit equations; the term $[f]$ determines whether we have a terminal node or not.

To incorporate additional nodes into the diagram, it is necessary to take further derivatives, which can be done recursively by applying the definitions once again. Our only real doubt is whether or not there is a finite total number of derivatives, so that the diagram will be finite. Higher derivatives correspond to words which begin with a specified sequence of letters rather than just a single letter; sometimes this is written as a single derivative with respect to the longer sequence.

6.2 Ideals and factors

One of the goals of any theory of the structure of an algebraic system is to decompose the system into standard combinations of standard components; in the case of vector spaces the result is to express any vector as a linear combination of a certain

number of basis vectors. However, a vector space represents the ultimate in uniformity; the typical decomposition in other systems may be neither as complete nor as uniform, but there are usually some vestiges of both a basis and of coefficients.

When the system has both a multiplication and an addition, these terms bearing their usual connotations with respect to things like the commutative and the distributive laws, the principal object of interest is the behaviour of products. If there is a subset, such as Ω , typically closed with respect to addition, and an entire set Σ , we may classify products according to which of these two sets contains them, particularly with respect to the source of the factors. The most notorious example of this relationship is the fact that the product of any factor with zero results in zero; the idea is to generalize the properties of zero.

The following table expresses the possible relationships:

left factor	right factor	product	name
$a \in \Sigma$	$w \in \Omega$	$aw \in \Omega$	left ideal
$w \in \Omega$	$a \in \Sigma$	$wa \in \Omega$	right ideal
$x \in \Omega$	$y \in \Omega$	$xy \in \Omega$	subwhatever

Any subset which is simultaneously a left ideal and a right ideal is simply called an ideal; the set consisting entirely of zero is the prototypical ideal, but so also is the entire set Σ .

In the case of regular expressions, Conway has explored these concepts directly, without making any reference either to ideals or to the general theory of semi-groups from which the results might possibly be taken. In doing so he exploits a natural order which exists for regular expressions, or any system having the regular expression operators; namely

$$x \leq y. \equiv \exists z \ni x + z = y$$

Zero is clearly a minimum element for this order; the goal is to factorize any regular expression E as a product of other regular expressions, but this may not be possible; thus one must work up to the concept in stages. The first is to define a *trial factorization* through the equation

$$F.G \dots H \dots J.K \leq E$$

and then define successively

- another trial factorization $F'.G' \dots H' \dots J'.K'$ is *dominant* if $F \leq F', G \leq G', \dots, H \leq H', \dots, J \leq J', K \leq K'$
- any term is *maximal* if it cannot be increased while keeping the product less than E
- a *factorization* is a trial factorization in which every term is maximal
- any term of a factorization is a *factor*, further characterized as *left* or *right* according to whether it occupies that position in some factorization

Amongst others, the canonical form of a regular expression provides trial factorizations whose terms are letters and derivatives. By taking further derivatives, trial factorizations may be developed with arbitrary words as leading terms. The point of starting with trial factorizations is the expectation that they can be promoted to factorizations by systematically increasing each of their terms.

In doing so we depend upon the fact that sums provide the route to upper bounds. Thus if $P \leq E$ and $Q \leq E$, $P + Q \leq E$, follows readily from the definition of inequality; as long as we do not insist that the result necessarily be a regular expression, the same conclusion also holds for arbitrary sums. Consequently the maximal element less than E is simply the sum of *all* the elements less than E .

A product is linear in each of its terms individually, so replacing G by summing the set

$$G' = \{\sigma \in \Sigma^* | F.\sigma \dots H \dots J.K \leq E\}$$

gives a new trial factorization dominating the first. Moreover the replacement does not spoil the maximality of any of the other terms; if a maximal term were increased, the distributive law shows that the expression which previously offended the inequality would still be present. Consequently trials can be refined one term at a time until they are finally factorizations.

Without further restriction, a factorization of E need not be unique; the trial factorizations even less so. One such restriction is that all terms but one be maximal; there is only one way the remaining term can be refined, and that must produce the term itself if it were already maximal. The concept of factor, defined as a maximal term, depends upon the environment in which it is encountered, but the environment can be reduced to a product of two terms for either a left or a right factor, or three terms for an interior factor, by using the associative law to consolidate the remaining terms.

Finally it is possible to conclude that, whatever the number of left factors, they are paired with unique right factors; thus they can be given a common index set reflecting this correspondence; let us call them L_i and R_i ; they all satisfy $L_i R_i \leq E$, each of which is dominant. It is a temptation to combine them with each other. Thus we further define $E_{i,j}$ as the maximal solution to the trial factorization

$$L_i E_{i,j} R_j \leq E$$

dominating $L_i.\phi.R_j$

Several special cases come to mind. Since $\lambda.E = E$, there must be an index ℓ for which $\lambda \leq L_\ell$ and $R_\ell = E$; likewise an index r for which $\lambda \leq R_r$ and $L_r = E$. Finally, since $L_\ell.E.R_r = E$ we conclude $E_{\ell,r} = E$.

As for the L 's and R 's, note that by the definition of inequality, $L_i R_i + X = E$, so that $L_\ell L_i R_i + L_\ell X = L_\ell E = E$, making $L_\ell L_i R_i$ a trial factorization with maximal terms—a factorization. Consequently $L_i = E_{\ell,i}$, and similarly $R_i = E_{i,r}$. Thus the $E_{i,j}$'s can be arranged into a square matrix, wherein the left factors form a certain row, the right factors a certain column, and E lies at their intersection. The $E_{i,j}$'s in general behave like elementary matrices, satisfying the following rules:

- $\lambda \leq E_{i,i}$ [non-zero diagonal]
- $E_{i,j} E_{j,k} \leq E_{i,k}$ [matrix basis]

- $AB \leq E_{i,k} \Leftrightarrow A \leq E_{i,j} \& B \leq E_{j,k}$
- $F_1 F_2 \dots F_n \leq E \Leftrightarrow F_1 \leq E_{\ell,i} \& c$

One consequence of these rules is that by extending inequality elementwise to matrices, $E \times E \leq E$, which implies that

$$E^* = E.$$

Interesting as it may be, this discussion has so far been purely theoretical, depending on maximizing through the formation of arbitrary sums. Turning to the canonical representation and symbolic derivatives it is possible to have more concrete results. Note that a right factor of E is a maximal R satisfying $LR \leq E$; in terms of words belonging to L this means

$$R = \bigcap_{w \in L} \frac{\partial E}{\partial w};$$

The pairing between L and R is quite evident from the structure of this formula. By building up the parallel theory of right derivatives, a similar relationship can be created to express L in terms of R . If it is found inconvenient to work with intersections, de Morgan's rules for regular expressions can be invoked to obtain

$$\bar{R} = \sum_{w \in L} \frac{\partial \bar{E}}{\partial w},$$

using the notation that \bar{X} is the complement of X , which may itself require some effort to obtain.

6.3 Rule 18

The generation evolving from the regular expression $(0 + 1)^*$ according to Rule 18 is $((0 + 00^*0)(0 + 10^*1))^*$, which can be simplified slightly to $(00^*(0 + 10^*1))^*$. If we calculate derivatives strictly according to the rules, simplifying only to remove λ 's ($\%$ in the following script) and terms beginning with ϕ 's ($\#$'s), we obtain for the first few derivatives

```
(00*(0+10*1))*/0= (0*(0+10*1))(00*(0+10*1))*
(00*(0+10*1))*/1= #
(00*(0+10*1))*/00= (0*(0+10*1)+%) (00*(0+10*1))*
(00*(0+10*1))*/01= (0*1)(00*(0+10*1))*
(00*(0+10*1))*/10= #
(00*(0+10*1))*/11= #
(00*(0+10*1))*/000= (0*(0+10*1)+%) (00*(0+10*1))*+(0*(0+10*1))(00*(0+10*1))*
(00*(0+10*1))*/001= (0*1)(00*(0+10*1))*
(00*(0+10*1))*/010= (0*1)(00*(0+10*1))*
(00*(0+10*1))*/011= (00*(0+10*1))*
(00*(0+10*1))*/100= #
(00*(0+10*1))*/101= #
(00*(0+10*1))*/110= #
(00*(0+10*1))*/111= #
```

If we introduce the definitions

$$\begin{aligned} e &= (00^*(0 + 10^*1))^* \\ a &= 0^*(0 + 10^*1) \\ b &= 0^*1 \end{aligned}$$

we may draw the conclusion that there are only five distinct derivatives, one of which is zero, and one of which is a sum. Thus the following table is sufficient to deduce all the remaining derivatives.

$$\begin{aligned} \frac{\partial e}{\partial 0} &= ae \\ \frac{\partial e}{\partial 1} &= \phi \\ \frac{\partial ae}{\partial 0} &= (a + \lambda)e \\ \frac{\partial ae}{\partial 1} &= be \\ \frac{\partial be}{\partial 0} &= be \\ \frac{\partial be}{\partial 1} &= e \end{aligned}$$

The five distinct derivatives are also the nodes in a diagram describing the second generation of Rule 18; it is an exit diagram, whose links are shown in Figure 6.1. Comparison with Figure 5.2 shows them to be equivalent, although slightly different.

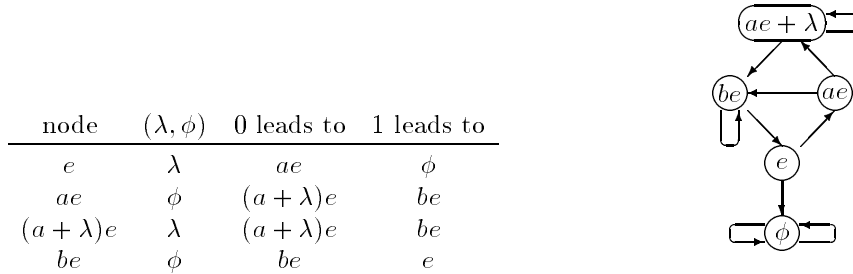


Figure 6.1: Exit diagram for Rule 18.

The right factors will be formed from intersections of these left derivatives.

It is possible to work out the right derivatives of this same regular expression; One might think that the results should be the same because Rule 18 is a symmetrical rule. However, the regular expression describing it is not, and consequently the remainder of the analysis must proceed accordingly. We can use the same program to calculate the derivatives, however, just by reversing the order of the factors:

$$((0+10^*1)0^*0)^*/0 = (0^*0)((0+10^*1)0^*0)^*$$

```

((0+10*1)0*0)*/1= ((0*1)0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/00= (0*0+%)((0+10*1)0*0)*
((0+10*1)0*0)*/01= #
((0+10*1)0*0)*/10= ((0*1)0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/11= (0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/000= (0*0+%)((0+10*1)0*0)*+(0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/001= ((0*1)0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/010= #
((0+10*1)0*0)*/011= #
((0+10*1)0*0)*/100= ((0*1)0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/101= (0*0)((0+10*1)0*0)*
((0+10*1)0*0)*/110= (0*0+%)((0+10*1)0*0)*
((0+10*1)0*0)*/111= #

```

This time we define

$$\begin{aligned}
 e &= (00^*(0+10^*1))^* \\
 f &= 00^* \\
 g &= 00^*(10^*),
 \end{aligned}$$

obtain the derivatives

$$\begin{aligned}
 \left(\frac{\partial e}{\partial 0}\right)_r &= ef \\
 \left(\frac{\partial e}{\partial 1}\right)_r &= eg \\
 \left(\frac{\partial ef}{\partial 0}\right)_r &= e(f+\lambda) \\
 \left(\frac{\partial ef}{\partial 1}\right)_r &= \phi \\
 \left(\frac{\partial eg}{\partial 0}\right)_r &= eg \\
 \left(\frac{\partial eg}{\partial 1}\right)_r &= ef
 \end{aligned}$$

and the entrance diagram

node	(λ, ϕ)	0 comes from	1 comes from
e	λ	ef	eg
ef	ϕ	$e(f+\lambda)$	ϕ
$e(f+\lambda)$	λ	$e(f+\lambda)$	eg
eg	ϕ	eg	ef

While different from the exit diagram, it evidently serves the same purpose. To use an entrance diagram one builds up an expression from right to left reading *backwards* along the arrows until an initial node is reached.

6.4 Factors for Rule 18

Now that we have all the derivatives for the second generation of Rule 18, it is possible to calculate its factors using the representation that they are intersections of the derivatives. The easiest way to calculate an intersection is to represent each intersectand via its diagram, collect corresponding nodes into pairs, and work out the regular expression representing the terminal nodes in this new diagram. A terminal pair is one which is terminal for both of its components, signified by the presence of λ 's in their own diagrams.

Accordingly the exit diagram for the intersections of the left derivatives can be gotten from a table of links, which can be written in a compact form if certain preparations are made. Define $x = a + \lambda$; then deduce the links for e and its derivatives by writing them in canonical form, and finally display all the pairwise intersections in a matrix.

$$\begin{array}{c|cccc}
 & e & ae & xe & be \\
 \hline
 e & \begin{pmatrix} \lambda + 0ae + 1\phi \\ \lambda + 0ae + 1\phi \end{pmatrix} & \begin{pmatrix} \lambda + 0ae + 1\phi \\ \phi + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \lambda + 0ae + 1\phi \\ \lambda + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \lambda + 0ae + 1\phi \\ \phi + 0be + 1e \end{pmatrix} \\
 ae & \begin{pmatrix} \phi + 0xe + 1be \\ \lambda + 0ae + 1\phi \end{pmatrix} & \begin{pmatrix} \phi + 0xe + 1be \\ \phi + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \phi + 0xe + 1be \\ \lambda + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \phi + 0xe + 1be \\ \phi + 0be + 1e \end{pmatrix} \\
 xe & \begin{pmatrix} \lambda + 0xe + 1be \\ \lambda + 0ae + 1\phi \end{pmatrix} & \begin{pmatrix} \lambda + 0xe + 1be \\ \phi + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \lambda + 0xe + 1be \\ \lambda + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \lambda + 0xe + 1be \\ \phi + 0be + 1e \end{pmatrix} \\
 be & \begin{pmatrix} \phi + 0be + 1e \\ \lambda + 0ae + 1\phi \end{pmatrix} & \begin{pmatrix} \phi + 0be + 1e \\ \phi + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \phi + 0be + 1e \\ \lambda + 0xe + 1be \end{pmatrix} & \begin{pmatrix} \phi + 0be + 1e \\ \phi + 0be + 1e \end{pmatrix}
 \end{array} .$$

To read this array, suppose that the intersection $ae \wedge be$ is to be calculated. Then from the row ae , column be we see from (ϕ, ϕ) that the node is not terminal, that 0 leads from (ae, be) to (xe, be) and that 1 leads from (xe, be) to (be, e) . If further links are traced onward from these, a diagram with eight nodes is obtained, but none of them is a terminal state. Thus this intersection is empty, $ae \wedge be = \phi$, a conclusion which is confirmed by observing that ae always contains an even number of 1's, while the number of 1's in be is always odd. Since the node (be, e) occurs in the same diagram with (ae, e) we can also conclude that that intersection is empty.

It is evident from the definition of xe that $e \leq xe$. The diagram for (ae, e) reads

$$\begin{aligned}
 (ae, e) &= (\phi, \lambda) + 0.(xe, ae) + 1.(be, \phi) \\
 (xe, ae) &= (\lambda, \lambda) + 0.(xe, xe) + 1.(be, be) \\
 (xe, xe) &= (\lambda, \lambda) + 0.(xe, xe) + 1.(be, be) \\
 (be, be) &= (\phi, \phi) + 0.(be, be) + 1.(e, e) \\
 (e, e) &= (\lambda, \lambda) + \dots
 \end{aligned}$$

These equations can be solved by inspection to obtain

$$\begin{aligned}
 ae \wedge e &= 0 + 00xe + 01be \quad (= 0ae), \\
 ae \wedge xe &= 0xe + 1be \quad (= ae).
 \end{aligned}$$

The only triple intersection not deducible from the foregoing results is

$$e \wedge ae \wedge xe = 00xe + 01be \quad (= 0ae),$$

finally giving us all the possible right factors. By examining canonical forms for the derivatives these results can be somewhat simplified to the forms shown in parentheses. As was to be expected, they all involve incomplete segments of the basic unit in e . All told, including the null intersection Σ^* , they are seven in number— ϕ , Σ^* , ae , be , $xe (= ae + \lambda)$, $0ae$, and $e (= 0ae + \lambda)$.

By repeating essentially the same calculation with $y = e(f + \lambda)$, the left factors are found to be ϕ , Σ^* , ef , eg , $ey (= ef + \lambda)$, $ef0$, and $e (= ef0 + \lambda)$.

In general, the n -fold intersections can be deduced from a table of n -tuples, and in each case a diagram can be made up and solved to determine the regular expression corresponding to each intersection. The only slight doubt that might remain is whether the diagram could be reduced, due to the unsuspected equality of some of its nodes. Sometimes this is evident from inspection; as when two different nodes solve the same equation, or when the expression is fairly short with a clear interpretation. Otherwise it might be necessary to resort to testing the diagram for equivalence classes.

Having developed this theory to a fair degree of elaboration, the question naturally arises as to its applicability. For studying infinite or cyclic systems formed from symmetric rules, it does not seem likely that there will be a useful distinction between left factors and right factors; they seem more related to starting up or ending a finite expression, as the analysis which we have made for Rule 18 shows. They are also likely to be of importance for rules which admit fuses or gliders.

On the other hand, it cannot be denied that the combination of Arden's lemma and the calculus of regular expressions yields a very concise and elegant way to describe diagrams symbolically.

6.5 A geometrical representation

Regular expressions represent sequences of symbols—letters taken from some alphabet. When the sequence is relatively short, the sequence is readily perceived and its properties may be examined. The notation is especially nice for summarizing cycles which can be repeated an arbitrary number of times, since a single star signifies this infinite collection of alternatives. When the regular expression is complicated, and particularly when it subsumes several distinct cycles at the same time, it rapidly becomes more complicated, both to write down completely and to comprehend once it has been written.

Although difficult of typography, the diagram representing a regular expression is preferable for expressions of moderate complexity, although eventually the many crossing lines and increased density of nodes take their toll, and the figure can no longer be comprehended.

Another representation has some advantages, although it is typographically even more difficult to deal with. To begin with, suppose that the symbols are chosen to be consecutive numbers, just as we are accustomed to assuming when dealing with automata. If a decimal point were to be placed in front of a finite sequence, it could be interpreted as a number with a positional representation in base k if there were k symbols running from 0 to $k \Leftrightarrow 1$. Finite sequences would be finite decimals, infinite sequences could be interpreted after the fashion of real numbers.

The principal difference is that in base 10 $.10000\dots$ and $.09999\dots$ are ordinarily considered to be the same real number; but have to be understood to be distinct regular expressions— 10^* in the first case and 09^* in the second.

The advantage of this representation is that every regular expression, no matter how complicated, is a point in the interval $0 \leq x \leq 1$, and so is represented in a finite region of space. The counterbalancing penalty is that this region can be filled more densely than is convenient, often too densely for a clear presentation, since the individual points are always fairly bulky and can never be presented visually with the mathematical precision of occupying zero area.

We are interested in regular expressions describing configurations of cellular automata, but these are sequences infinite in both directions. To capture both directions in such a string, break it at some point, then write the left hand side backwards. Inserting a figurative decimal point as before now produces two intervals of unit length; the natural thing is to treat them as two coordinates generating a unit square. The vertical, or y -coordinate represents the left half of the sequence, the horizontal or x -coordinate represents the right half.

Shifting, or changing the decision as to where to break the sequence, will naturally change the point which represents the sequence; in any event it is a geometric operation whose general nature can be described. Indeed this is a fundamental operation on sequences: a right shift splits the unit square into k horizontal strips and reassembles them into k vertical strips after stretching them in one direction and compressing them in the other.

The numerical value associated with each sequence defines a distance, and consequently the distance between two sequences is just the sequence of differences. At this point the difference between this representation and that of the real numbers manifests itself; although different terms are included in the sum in the normal way, there are no carries from one term to another and so there is no confusion that a number may have two different representations as a repeating decimal.

In mathematical terms, suppose that there are two sequences

$$a_{-n}, a_{-n+1} \dots, a_{-1}, a_0, a_1, \dots, a_{n-1}, a_n$$

$$b_{-n}, b_{-n+1} \dots, b_{-1}, b_0, b_1, \dots, b_{n-1}, b_n$$

Their distance is defined as


$$\|\mathbf{a} \leftrightarrow \mathbf{b}\| = \sum_{i=-n}^n \frac{|a_i \leftrightarrow b_i|}{k^{|i|}}.$$

The distance is small, or the sequences are close, when their central regions agree; the closer they are, the longer the matching core must be.

To have an idea of how this representation works, let us consider Rule 18 for a moment. We have seen that the sequence 111, amongst others, is excluded from the second generation. If we divide the interval $0 \leq x \leq 1$ into eighths, any sequence $.xyz\dots$ beginning with $.111$ will occupy some position in the last eighth of this interval. This interval will include sequences of the form $.1111\dots$, but sequences of the form $.0111\dots$ will occupy the last eighth of the interval $0 \leq x \leq \frac{1}{2}$. Recursively, the last eighth of any binary subdivision of the unit interval is excluded. Similar

comments apply to sequences extending to the left, so that as the unit square represents various doubly infinite sequences the excluded sequences will come to resemble a Scotch plaid.

Additionally it is necessary to consider sequences whose central point lies amongst the three excluded 1's. For example 1.11 will refer to points within the upper half of the square, but in the rightmost quarter; 11.1 to the upper quarter but right half. These regions have to be added to the plaid, as well as the further subdivisions implied by further decimals. Thus the same basic pattern is repeated recursively in each quadrant, 16-ant, and so on of the unit square.



Sequences of three 1's are not the only configuration excluded from the second generation, since it is also impossible to encounter a single 1 in a sequence bounded by pairs of 1's. Taking into account all of the excluded sequences simply gives an even finer structure to the plaid. In general the plaid is a two dimensional version of the classical Cantor set, which is the residue remaining after the middle thirds of the unit interval and surviving subintervals have been removed. In arithmetic terms, the Cantor set is the set of all decimals in the unit interval which contain no 1's when they have been displayed with respect to the base 3; which is one of the simplest possible examples of an excluded sequence.

The plaid can be used as the starting point for further discussions. For example the third generation must occupy a subset of the permitted portion of the second generation plaid, and so on as the automaton evolves through further generations; it is even meaningful to think of a limiting set which consists of just those points, if any, which remain after an infinite number of generations. The limiting set will include all the points representing cycles of whatever period, as well as gliders, fuses, and a variety of structures which may have no finite or cyclical representation.

Metρίζing regular expressions, and Σ^* in general, has the advantage of creating a topological space, to which the whole machinery of analysis may be applied. As a compact topological space, even stronger conclusions, such as the necessary existence of certain limits, can often be drawn. An extensive literature concerns the application of these ideas to probability and symbolical dynamical systems; but there is no reason that it cannot be broadened to include regular expressions as well. In particular, the interaction between probability measures and linear cellular automata can be studied through this model, to the mutual enlightenment of both disciplines.

Chapter 7

Probabilistic de Bruijn matrix

The reduced evolutionary matrix and the de Bruijn diagram have probabilistic versions, in which the zeroes and ones which enable links in the diagram are replaced by probabilities that the links are to be used. This does more than express the likelihood that one thing or another will occur; it allows some quantitative comparisons to be made.

Since the reduced evolution matrix enumerates the numbers of $(n + 2r)$ -block ancestors of n -blocks, the probabilistic evolution matrix can be used estimate the likelihood that the ancestors actually occur, and thus to develop self-consistent estimates for their probabilities. Probabilistic de Bruijn matrices are useful for studying correlations between cells or strings of cells situated at a distance from one another because its matrix elements could describe the probability that one n -block will overlap the next; powers of the matrix would relate blocks through a chain of overlaps.

7.1 Block probabilities

The rules of probability are widely known and frequently employed in a wide variety of contexts; nevertheless it is a probability which assumes such things as independent events and the lack of correlation. Under those conditions probabilities of alternatives are summed and probabilities of coincidences are multiplied. Sometimes it seems reasonable that the justifying hypotheses are valid, other times it is simply assumed so with the hope that the calculations will still be reliable.

When the evolution of cellular automata is treated probabilistically, it does not seem reasonable that the evolution of neighboring cells is independent, inasmuch as they share a common portion of their neighborhood and evolve according to deterministic rules. Nevertheless the rules are complicated, and the overlapping segment of the neighborhoods enters differently into the evolution of each cell, so the coupling between the two evolutions might not be all that strong.

Going ahead with naive probabilistic calculations for cellular automata produces results which are neither very good nor very bad. Tendencies are evident and the general conclusions seem to be reliable, but the results are clearly defective as soon as any accuracy is required.

One way to get better results is to work out the rules for more complicated combinations of probabilities, as is routinely done to get the probabilities for the values of all kinds of mathematical functions in terms of the probability distributions of their arguments and the fundamental arithmetic operations. This approach does not work well for cellular automata because of the complicated way that the rule of evolution would have to be expressed as an arithmetic function of a real variable.

A more tractable approach seems to be to work with probabilities for sets of cells rather than with individual cells; in particular, with the linear sequences of cells from which linear cellular automata are already composed. Thus correlations between cells are seen in the differences in the probabilities of different pairs of cells, which do not necessarily have to be the products of probabilities for the individual cells. Further correlations can be deduced by studying the probabilities of triples of cells, quadruples, and so on. The extreme logical limit to this approach would be to assign probabilities to each and every possible sequence, no matter how long.

Sequences of cells will be called *blocks*, to emphasize their linear arrangement; the term n -block will be used when its exact length is required. Some blocks form parts of others; indeed the best way to describe $(n + 1)$ -blocks is by systematically extending n -blocks in all possible ways; since they are linear, the extension can be made to the left or to the right. Insertions could be made in the middle, but that will be regarded as a compound extension.

The assignment of probabilities to blocks ought to respect the process of extension. Since there are more extensions than bases, each of them should have a smaller probability than its base; moreover the probability of the base could be expected to be the sum of the probabilities of all its extensions. In this way a constant amount of probability will be subdivided more and more finely as the blocks grow longer.

When formalized, these restraints are called the Kolmogorov consistency conditions. For linear sequences there are two such conditions, according to the handedness of the extension. Thus, if abc were a 3-block, Σ , the collection of cell states, the requirements are

$$\begin{aligned} p(abc) &= \sum_{x \in \Sigma} p(xabc) \\ p(abc) &= \sum_{x \in \Sigma} p(abcx) \end{aligned}$$

By repeated application of the consistency conditions the probabilities of all m -blocks for $m \leq n$ are determined once all the n -block probabilities are known. But because there are *two* consistency conditions, the choice of n -block probabilities is not completely arbitrary. For example, in a binary alphabet, 1 extends to the left to become 01 or 11, to the right to become 10 or 11. This implies that $p(01) + p(11) = p(10) + p(11)$ or in short that $p(01) = p(10)$, which is a restraint on the assignments of probabilities to 2-blocks.

7.2 Kolmogorov conditions in matrix form

Block probabilities for linear chains can be defined and their properties described by an extremely elegant matrix formulation. Binary sequences already illustrate the principles quite clearly, as we can see by writing the equations defining block probabilities up to length two. Generally it is understood that $p(\lambda) = 1$.

$$\begin{aligned}
 p(\lambda) &= p(0) + p(1) \\
 p(0) &= p(00) + p(01) \\
 p(1) &= p(10) + p(11) \\
 p(0) &= p(00) + p(01) \\
 p(00) &= p(000) + p(001) \\
 p(01) &= p(010) + p(011) \\
 p(10) &= p(100) + p(101) \\
 p(11) &= p(110) + p(111) \\
 &\dots
 \end{aligned}$$

Concentrating on the last group of equations, four 2-block probabilities are written as linear combinations of eight 3-block probabilities:

$$\begin{bmatrix} p(00) \\ p(01) \\ p(10) \\ p(11) \end{bmatrix} = \begin{bmatrix} 1 & 1 & \dots & \dots & \dots & \dots \\ \dots & \dots & 1 & 1 & \dots & \dots \\ \dots & \dots & \dots & \dots & 1 & 1 & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & 1 & 1 \end{bmatrix} \begin{bmatrix} p(000) \\ p(001) \\ p(010) \\ p(011) \\ p(001) \\ p(001) \\ p(001) \\ p(001) \end{bmatrix}.$$

If $U = (1, 1)$ is a row vector, I a unit matrix, the matrix in the above equation can be written as a tensor product, $I \otimes U$.

However, there is another way that the same system of equations can be written, which is:

$$\begin{bmatrix} p(00) \\ p(01) \\ p(10) \\ p(11) \end{bmatrix} = \begin{bmatrix} p(000) & p(001) & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & p(010) & p(011) & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & p(100) & p(101) & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & p(110) & p(111) & \dots & \dots \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}.$$

Next, we partition the rectangular matrix of probabilities into two square matrices A and B , and partition the column of eight 1's correspondingly into two identical columns. Then, in terms of submatrices, the right hand side of this equation takes the form $AV + BV = (A + B)V$.

Finally the equations relating 2-blocks to 3-blocks become

$$\begin{bmatrix} p(00) \\ p(01) \\ p(10) \\ p(11) \end{bmatrix} = \begin{bmatrix} p(000) & p(001) & \cdot & \cdot \\ \cdot & \cdot & p(010) & p(011) \\ p(100) & p(101) & \cdot & \cdot \\ \cdot & \cdot & p(110) & p(111) \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}.$$

in which the matrix $(A + B)$ will be recognized as having the form of a three-stage de Bruijn matrix, but with probabilities as non-zero elements.

The interesting point is that this entire derivation can be repeated, with slight variations, for left extension rather than right extension, to obtain the 2-block probabilities as sums of 3-block probabilities. The same “de Bruijn” matrix will make its appearance if we write the 2-block probabilities as a row vector rather than a column vector:

$$[p(00) \ p(01) \ p(10) \ p(11)] = [1 \ 1 \ 1 \ 1] \begin{bmatrix} p(000) & p(001) & \cdot & \cdot \\ \cdot & \cdot & p(010) & p(011) \\ p(100) & p(101) & \cdot & \cdot \\ \cdot & \cdot & p(110) & p(111) \end{bmatrix}.$$

In this form, the Kolmogorov consistency conditions require that the row sums of the “de Bruijn” matrix match the column sums. The reason that the adjective ‘de Bruijn’ has been enclosed in quotes is that this is not quite the matrix which we intend to honor by such a description; the term is reserved for a slightly different matrix for which the consistency conditions will be met.

7.3 Probabilistic de Bruijn matrix

A de Bruijn diagram is a basic structure which can be applied in various ways. It describes all the possible ways that two sequences of symbols, or windows, can overlap if they are displaced with respect to each other; just the relationship between the neighborhoods of two successive cells in a linear cellular automaton. Selecting a subset of the diagram by retaining or discarding selected links discloses states of a given period, gliders, and other information. An intermediate approach would assign probabilities to the links rather than selecting them. As an extreme case, a probability of zero would exclude a link, a probability of one would surely include it, giving other values in the same interval intermediate interpretations.

The connectivity matrix of a graph is readily adapted to show the probabilities of the links; one has only to define

$$M_{i,j} = \begin{cases} p(i \rightarrow j) & i \rightarrow j \\ 0 & \text{otherwise} \end{cases}$$

The rules of matrix multiplication coincide with the rules for compound probabilities if it is desired to calculate the probability of forming chains of several links. Even the de Bruijn matrix itself is readily converted into a probability matrix by

replacing its non-zero elements by the value $1/k$, making the choice of all paths equally likely.

Just as probabilities can be associated with the links of a de Bruijn diagram, they also can be associated with the nodes; to conserve the probabilistic interpretation of matrix multiplication the nodal probabilities might be used to form a vector, with matrix equations describing the linkages between nodes. A vector of probabilities would have non-negative real components—with a unit sum, if it happens that all the different components comprise a complete set of alternatives.

Any matrix intended to preserve each and every such column vector must necessarily have columns of unit sum, as can be seen by testing the matrix on unit vectors. Such a matrix is called a *stochastic* matrix. Likewise the preservation of unit row sums would require the matrix to have unit row sums. Sometimes the conditions for rows and columns must be met simultaneously, for which the adjective *doubly stochastic* would be used.

In terms of the de Bruijn diagram, if the incoming probabilities at a node were required to have a unit sum, the probabilistic de Bruijn matrix would have to be column-stochastic. This property would refer to the rows if the outgoing probabilities were constrained. Expressing either constraint in matrix form guarantees a unit eigenvalue, with a probability vector as the equilibrium eigenvector of opposite handedness. These and certain other conclusions characterize stochastic matrices, all of which information can be found in any one of several recent textbooks[41, 13, 103] which can be consulted for descriptions and proofs of the results.

If the links in an n -stage de Bruijn matrix are replaced by their corresponding n -block probabilities, the row or column probabilities do not sum to 1, but rather to the $(n \Leftrightarrow 1)$ -block probabilities; this suggests dividing each column by its column sum to get a stochastic matrix; alternatively the rows could be divided by the row sums. Generally these are two distinct choices and result in two different de Bruijn matrices. For right extension, we define

$$\begin{aligned} p(ab \rightarrow bc) &= \frac{p(abc)}{\sum_x p(xbc)} \\ &= \frac{p(abc)}{p(bc)}; \end{aligned}$$

for left extension,

$$\begin{aligned} p(ab \leftarrow bc) &= \frac{p(abc)}{\sum_x p(abx)} \\ &= \frac{p(abc)}{p(ab)}; \end{aligned}$$

Among other things, these denominators compensate for the overlapping between blocks when de Bruijn matrices are multiplied.

A stochastic de Bruijn matrix (of either handedness) can be used to generate a set of n -block probabilities which will satisfy the Kolmogorov consistency conditions. Suppose that the rows are indexed by the $(n \Leftrightarrow 1)$ -block sequences, so that

writing one of the non-zero matrix elements as $m(axb)$ identifies it as belonging to the ax th row and the xb th column. Suppose further that the matrix is column stochastic. Then the row $(1, 1, \dots, 1)$ is an eigenvector belonging to unit eigenvalue, and we may call $p(xb)$ the components of the corresponding column eigenvector. By rewriting

$$p(ax) = \sum_b m(axb)p(xb)$$

in the form

$$p(ax) = \sum_b (m(axb)p(xb)) \times 1,$$

we see that the elements $p(ax)$ are row sums of a new matrix whose elements are defined implicitly by the equation

$$m(axb) = \frac{p(axb)}{p(xb)}.$$

Since we have only renamed the elements $m(axb)$, the matrix M is still column stochastic; since $p(xb)$ has the same denominator throughout a given column we have

$$1 = \sum_a \frac{p(axb)}{p(xb)};$$

when rewritten in the form

$$p(xb) = \sum_a p(axb)$$

we have the other half of the consistency condition.

7.4 Some properties of n -block probabilities

The stochastic de Bruijn matrices suggest one way to obtain a set of n -block probabilities satisfying the Kolmogorov consistency conditions, at least for $(n \Leftrightarrow 1)$ -block probabilities in terms of n -block probabilities. All that is required is to fill in the matrix elements of the de Bruijn matrix with positive numbers, taking care not to exceed, say, the unit column sum. The choice is relatively arbitrary until the selection of the final element of the column, which must be chosen to complete the sum. For k symbols, this means that only k^{n-1} of the k^n non-zero elements lack any freedom of choice; the remaining $(k \Leftrightarrow 1)k^{n-1}$ elements are then parameters to be chosen.

Once this is done, the equilibrium eigenvector is to be determined and each of the columns of the stochastic matrix is to be multiplied by the corresponding component, to produce the n -block probabilities. The components of the eigenvector are the $(n \Leftrightarrow 1)$ -block probabilities; the probabilities for smaller blocks follow from summing longer blocks.

It is interesting to observe that once consistency has been established for $(n \Leftrightarrow 1)$ -block probabilities relative to n -block probabilities, the consistency of all smaller

blocks follows. Starting with n -blocks, the inductive step consists of verifying consistency for $(n \Leftrightarrow 2)$ -blocks:

$$\begin{aligned}
 p(a) &= \sum_x p(ax) \\
 &= \sum_x \sum_y p(axy) \\
 &= \sum_x \sum_y p(yax) \\
 &= \sum_y \sum_x p(yax) \\
 &= \sum_y p(ya),
 \end{aligned}$$

and to arrive at successively shorter blocks. In similar fashion it follows that the probability of any block up to length n is simply the sum of the probabilities of all the possible extensions which can be derived by the same route; for example,

$$p(0) = \sum_x \sum_y p(x0y).$$

Given the interrelation between n -blocks and all the shorter blocks, some consideration can be given to using the probabilities of short blocks as parameters. It is not possible to use *only* the probabilities of short blocks as parameters; were it otherwise there would be no need to build up a theory of block probabilities. Nevertheless it is possible to attempt the maximum usage of the probabilities of short blocks. For example, $p(1), p(2), \dots, p(k \Leftrightarrow 1)$ can all be chosen before the value of $p(0)$ is fixed in order to make the complete sum equal to 1—a total of $k \Leftrightarrow 1$ parameters. Or, of a total of k parameters, one— $p(0)$ —has been excluded.

Then, among the 2-blocks, $p(x1), p(x2), \dots, p(x(k \Leftrightarrow 1))$ can be chosen before $p(i0)$ is needed to complete a sum of $p(x)$. Alternatively, we can consider that $p(x0)$ is excluded as a parameter, as is $p(0x)$ when the left extensions are considered. Going on to 3-blocks, we see that all the probabilities $p(0xy)$ or $p(xy0)$ are excluded. Continuing through n -blocks, the total number of parameters comprising probabilities in which 0 is neither an initial nor a terminal symbol is

$$(k \Leftrightarrow 1) + (k \Leftrightarrow 1)(k \Leftrightarrow 1) + (k \Leftrightarrow 1)k(k \Leftrightarrow 1) + \dots + (k \Leftrightarrow 1)k^{n-2}(k \Leftrightarrow 1)$$

which can be summed to give the value $(k \Leftrightarrow 1)k^{n-1}$; just the same as the number of parameters required to make a stochastic de Bruijn matrix.

Even when full use is made of the short blocks to define parameters the fraction of parameters which are n -block probabilities is

$$\frac{(k \Leftrightarrow 1)k^{n-2}(k \Leftrightarrow 1)}{(k \Leftrightarrow 1)k^{n-1}},$$

which reduces to $(k \Leftrightarrow 1)/k$; $1/2$ for a binary sequence, larger and approaching 100% for others.

7.5 Some simple examples

The generic form of a 2×2 matrix with unit column sum is

$$\begin{bmatrix} a & b \\ 1 \Leftrightarrow a & 1 \Leftrightarrow b \end{bmatrix},$$

whose eigenvalues are

$$\lambda = 1, a \Leftrightarrow b;$$

whose matrix of left eigenvectors is

$$\begin{bmatrix} 1 & 1 \\ 1 \Leftrightarrow a & b \end{bmatrix},$$

and whose matrix of right eigenvectors is

$$\begin{bmatrix} b/(1 + b \Leftrightarrow a) & 1 \\ (1 \Leftrightarrow a)/(1 + b \Leftrightarrow a) & \Leftrightarrow 1 \end{bmatrix}.$$

The eigenvalue 1, whose eigenvector has components proportional to the off diagonal elements of the matrix, represents an equilibrium probability. The second eigenvalue lies in the range $\Leftrightarrow 1 \leq \lambda \leq 1$, with an eigenvector suitable for probability differences. Indeed it is the factor by which the disequilibrium decreases in each generation; if it is zero equilibrium is reached in one step, otherwise there will be an exponential approach to equilibrium. Depending on the sign, there may be oscillations about equilibrium, or a uniform approach to equilibrium. However, if the second eigenvalue is as large as one, there will either be a degeneracy by which any probability is in equilibrium, or the probabilities for zeroes and ones will be exchanged with each other and alternate forever after.

A doubly stochastic one-stage de Bruijn diagram would have a probability matrix

$$\mathcal{P}_1 = \begin{bmatrix} p & q \\ q & p \end{bmatrix}$$

with $q = 1 \Leftrightarrow p$, whose eigenvalues would be $\lambda = 1, p \Leftrightarrow q$; being symmetric, the matrix of both left and right eigenvectors would be

$$\frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & \Leftrightarrow 1 \end{bmatrix}.$$

The probability matrix of a two stage diagram is a little more complicated.

$$\mathcal{P}_2 = \begin{bmatrix} p_1 & q_2 & \cdot & \cdot \\ \cdot & \cdot & p_3 & q_4 \\ q_1 & p_2 & \cdot & \cdot \\ \cdot & \cdot & q_3 & p_4 \end{bmatrix}$$

with $q_i = 1 \Leftrightarrow p_i$, and the characteristic polynomial

$$(\lambda^4 - \lambda^3(p_1 + p_4) - \lambda^2(p_1 p_4 - p_2 p_3) + \lambda(p_2 d_{34} + p_3 d_{12}) - d_{12} d_{34}).$$

with two determinants defined by

$$\begin{aligned} d_{12} &= (p_1 p_2 \Leftrightarrow q_1 q_2) \\ d_{34} &= (p_3 p_4 \Leftrightarrow q_3 q_4) \end{aligned}$$

The determinant of the de Bruijn matrix is the product of two smaller determinants corresponding to the evident blocks in the de Bruijn matrix; in this we have a special case of a quite general result. The p 's and q 's were defined as they are with the thought of equating the subscript pairs (1,2) and (3,4) to get a doubly stochastic matrix, but the best interpretation of the vanishing of the small determinants is that the biases of the probabilities in their submatrices are equal.

The constant term in the characteristic equation will be zero, giving a single zero root, if either determinant is zero. The coefficient of λ will also be zero if both are zero, producing a double zero root. In order to get three equal roots and thus open up the possibility of the Jordan normal form, requires some similarity between the two submatrices in addition.

7.6 Determinant and inverse

The representations of the de Bruijn matrix as sums and products need very little change to apply to the probabilistic versions as well. The factorizations give immediate formulas for determinants; when they do not vanish, the block diagonal form of one of the factors, together with the fact that the other is a permutation matrix quickly reveals a form for the inverse matrix. Taking as an example the matrix of the last section, we have

$$\mathcal{P}_2^{-1} = \begin{bmatrix} \frac{p_2}{d_{12}} & \cdot & \frac{-q_2}{d_{12}} & \cdot \\ \frac{-q_1}{d_{12}} & \cdot & \frac{p_1}{d_{12}} & \cdot \\ \cdot & \frac{p_4}{d_{34}} & \cdot & \frac{-q_4}{d_{34}} \\ \cdot & \frac{-q_3}{d_{34}} & \cdot & \frac{p_3}{d_{34}} \end{bmatrix}.$$

The Gerschgorin limit of this matrix requires its maximum eigenvalue to be less than the greatest reciprocal determinant $1/d_{12}$ or $1/d_{34}$ (all with respect to absolute values). Generalizing, this means that *every* eigenvalue of \mathcal{P}_n must be greater than the lesser of these determinants. Of course, if one of them is zero, we know that the lower bound is reached and \mathcal{P}_n is singular.

In the other direction, unless some p_i or q_i is zero, every matrix element of \mathcal{P}_n^n will be strictly positive, although in the probabilistic case no two of them have to be equal. This ensures that the maximum eigenvalue, $\lambda = 1$, will be unique. Conversely, if some probability or coprobability vanishes, the possibility of degeneracy exists (but is not obligatory).

7.7 Characteristic equation

Unfortunately information about the other eigenvalues, particularly the second largest of them, is not so readily available—either directly or from the factored forms. However, formulas for the coefficients of the characteristic equation can be derived from the symbolic form of the de Bruijn matrix elements; their expression is even more elegant if a new rule of composition is introduced for regular expressions, particularly for sequences of symbols.

7.7.1 merged product

Let us suppose a and b are letters, and that x is a sequence, possibly null and possibly empty. We define the merged product of the sequences s and t , denoted $s \vee t$, by

$$s \vee t = \begin{cases} axb & s = ax \text{ and } t = xb \\ \phi & \text{otherwise} \end{cases}$$

Likewise, define the overlap, $s \wedge t$, by

$$s \wedge t = \begin{cases} x & s = ax \text{ and } t = xb \\ \phi & \text{otherwise} \end{cases}$$

These are just the combinations of indices required to work with the de Bruijn matrices. For example, the $2^n \times 2^n$ column-stochastic matrix is defined in terms of n - and $(n + 1)$ -block probabilities by the formula

$$[\mathcal{C}_n]_{ax,yb} = \begin{cases} \frac{p(ax \vee yb)}{p(ax)} & p(ax) \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

Correspondingly the row-stochastic matrix is defined by

$$[\mathcal{R}_n]_{ax,yb} = \begin{cases} \frac{p(ax \vee yb)}{p(xb)} & p(xb) \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

Unless stated otherwise, we will assume that a probabilistic de Bruijn matrix is column-stochastic. Since it is not always convenient to show the matrix elements as a quotient, but it is essential to know which of them intrinsically vanish, let us write

$$[\mathcal{C}_n]_{ax,yb} = m(ax, yb)$$

and note that it vanishes unless $x = y$, or alternatively when $ax \vee yb = \phi$.

7.7.2 trace

The trace of \mathcal{C}_n , which is the coefficient of λ^{n-1} in the characteristic equation, satisfies

$$Tr(\mathcal{C}) = \sum_{ax} m(ax, ax)$$

$$\begin{aligned}
&= \sum_{i|ax=xi} m(ax, xi) \\
&= \sum_i m(i^*, i^*)
\end{aligned}$$

The essential element in all of these calculations is to factor the column index in such a way that it can match a row index and yet produce a non-zero merged product.

7.7.3 second coefficient

$$\begin{aligned}
a_2 &= \sum_{ax \leq by} (m(ax, ax)m(by, by) \Leftrightarrow m(ax, by)m(by, ax)) \\
&= \sum_{i|ax=xi} \sum_{j|by=yj} m(ax, xi)m(by, yj) \Leftrightarrow \sum_{i|by=xi} \sum_{j|ax=yj} m(ax, xi)m(by, yj) \\
&= \sum_{i \leq j} m(i^*, i^*)m(j^*, j^*) \Leftrightarrow \sum_{i \leq j} m((ij)^*)m((ji)^*)
\end{aligned}$$

In essence, this sum runs over all cycles of length 2 in the de Bruijn diagram. The solution to the requirements that $by = xi$ and $ax = yj$ is that $y = (ab)^*$ insofar as this notation can be respected on account of the fact that y might not have even length.

7.7.4 principal minors

Striking out rows and columns intersecting on a diagonal will always leave some surviving elementary determinants, until so many have been stricken out that it is no longer possible to form a complete determinant. The coefficient of the elementary determinants will depend on how much of the remaining elementary determinant can be salvaged.

In general, the determinant of the de Bruijn $|m(ix, yj)|$ factorizes, so that

$$|m(ix, yj)| = \prod_{x,x} |m(ix, xj)|,$$

where the factors are $k \times k$ determinants obtained by varying the indices i and j with fixed x . To get the minor of the hz, hz diagonal element, we need to consider two cases. If this pivot does not belong to any subdeterminant, it can be ignored. If it does, we drop the z, z term from the product, replacing it by the internal minor from within the omitted factor. Thus

$$a_1 = \sum_{hz=wh} \text{cofactor}(m(hz, hz)) \prod_{x \neq z} |m(ix, xj)|.$$

7.7.5 determinant

Although we already have the determinant of the probabilistic de Bruijn matrices from their factored forms, it is worth observing that it can be calculated directly by following the symbolic procedure.

$$|m(ix, yj)| = \prod_{x, x} |m(ix, xj)|,$$

7.8 Correlations

One application of the probabilistic $(n + 1)$ -stage de Bruijn matrix is to finding the probability of encountering one n -block in the vicinity of another; depending upon the direction, the basic probabilities are obtained by multiplying a vector of n -block probabilities by the matrix of the corresponding handedness to get the probabilities of a shifted block. Powers of the matrix would then yield blocks shifted to greater and greater distances, giving some importance to knowing about the behaviour of such powers. The degeneracy of the largest eigenvalue and the size of the second largest eigenvalue are the quantities which determine the limiting behaviour of powers and how rapidly their limits are reached.

If the largest eigenvalue is non-degenerate, which will be assured if none of the links has zero probability, then both the dominant eigenvalue and its normalized eigenvector will be unique. The second largest eigenvalue governs the exponential approach to equilibrium, which will be the more rapid the smaller the eigenvalue. In general a slow approach to equilibrium is associated with a large bias. Unfortunately the commonest estimators of eigenvalues tend to estimate either the largest eigenvalue, or else the smallest by first inverting the matrix, but not any of the others.

There is a certain informal expectation that if a block is long enough, probabilities will not be much affected by dropping the final cell, especially if comparative probabilities are considered. Illustrating this concept with a fairly short sequence of three cells, we might expect to find

$$\frac{p(abc)}{p(ab)} = \frac{p(bc)}{p(b)}.$$

There is no real justification for such an assumption, but it gives an alternative viewpoint to the equivalent equation

$$p(abc) = \frac{p(ab)p(bc)}{p(b)},$$

which says that we should multiply the probabilities for the two sequences ab and bc because the presence of both is required to form the sequence abc , but that the probability of b should be divided out because it is common to both sequences and thus counted twice when they are joined.

As yet, the source of the block probabilities which enter into the formation of an actual probabilistic de Bruijn matrix has not been discussed; only the fact that that

they ought to satisfy the Kolmogorov consistency conditions. One way to obtain them would be to look for probabilities which were self consistent with respect to evolution; the context in which to do this is the probabilistic version of the reduced evolution matrix.

Chapter 8

Probabilistic evolution matrix

Empirically it has been noted that automata go through three phases of evolution. If a ring is started out with an arbitrary initial configuration the randomness decreases fairly rapidly, followed by a longer period with a fairly constant density of states, and a final period which will be reached after a time depending on the length of the ring, in which the evolution is completely periodic.

Paying closer attention to the intermediate phase, it is seen that the density of states is not constant but that it fluctuates, and that there may be short range correlations between the states. Nevertheless, these densities, correlations, and even standard deviations can be calculated and they seem to agree fairly well with experience.

8.1 Regularities and anomalies

There are both temporal and spatial correlations; we have seen how the latter can be estimated by the probabilistic version of the de Bruijn diagram. It is interesting that the most serious logical discrepancy in probabilistic estimates arises from the excluded states which can be determined from the subset construction. Exclusion amounts to assigning a probability of zero, but zero probabilities cannot arise in probabilistic calculations, which involve only sums and products of positive quantities. Thus calculations based on the probabilities of individual cells will not suffice, and strings of cells must be taken into account from the outset.

Since there are arbitrarily long strings which are excluded for the first time (which is to say that none of their shorter segments is excluded, but that they are) it would seem that no theory based on the extrapolation of probabilities from finite strings could be mathematically exact. Nevertheless it can be hoped that there is a degree of approximation which is sufficient for practical purposes, but still not so complicated as to be beyond reasonable access.

For example, 10101001 is an excluded word for Rule 22, but it would have a probability of $1/256 = 0.004$ if zeroes and ones were considered equally probable.

This is not the worst distortion which Rule 22 suffers, because the distribution of frequencies for other short sequences of is by no means uniform, much less a simple function of the number of zeroes and ones they contain. For example, segments of the form $(01)^*$ are generally hard to come by, inasmuch as that is a sequence which it is its only ancestor. Most rules have their own excluded words, some of them even shorter than the eight letters which form the shortest excluded words for Rule 22.

In its probabilistic version, the de Bruijn diagram serves to predict the probabilities of the possible sequences comprising a shifting window in a long chain of states, and thus the correlations between such a sequence and a similar one occurring somewhere else in the chain. The conclusion afforded by the Frobenius-Perron theory is that there will always be an equilibrium vector of probabilities, that under certain circumstances will be unique. The particular form of the de Bruijn matrix shows that the lack of any extreme bias is sufficient to make the equilibrium unique; also that any zero biases will result in zero eigenvalues. We even know that the equilibrium probabilities are the n -block probabilities if the de Bruijn matrix is derived from $(n + 1)$ -block probabilities.

Although the disequilibrium eigenvalues and eigenvectors do depend in detail upon the transition probabilities of each particular diagram, the general conclusion is that the greater the bias the slower the approach to equilibrium, the node with the least bias establishing an upper limit to the speed. Zero bias results in an immediate equilibrium which may be accompanied by degenerate eigenvectors, while extreme bias offers the only possibility of alternate equilibria.

Whatever may be the spatial correlations of the moment, it is time evolution which drives the statistics of a linear automaton, continually challenging any correlations or lack thereof which may be found in the spatial distribution of the cells. Consequently we should look for probabilities, either of single cells or correlated probabilities of sequences of cells, which are consistent with the rule of evolution of the automaton.

8.2 Mean field theory

There are two approaches to this computation in the literature. One is called *mean field theory* and begins by assigning probabilities to each of the k states of the automaton, and then calculates the probabilities in the next generation on the basis of the usual combinatorial rules of probability, assuming that the probabilities for each of the cells in a neighborhood are independent. It is then possible to solve for a set of self consistent probabilities for each state. The results of such calculations are generally plausible but do deviate significantly from empirical observations. The suspect element in the calculation is the assumption of independence.

To see how this works, let us once again recall the transitions defining Rule 22:

$$\begin{array}{cccccccc} 111 & 110 & 101 & 100 & 011 & 010 & 001 & 000 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{array} .$$

Five neighborhoods evolve into zeroes, three into ones; thus one might predict 37.5% ones would be found each generation on the basis of the number of ancestral

neighborhoods. This is a better estimate than saying that 50% of the cells ought to be ones because there are only two different values they can have, but we have no reason to believe that all neighborhoods are equally likely either.

Taking the probability of finding a one as p , its coprobability as q , we could estimate $p' = 3pq^2$ for the probability of finding a one in the following generation, based on the makeup of the three neighborhoods that evolve to one. Mean field theory takes the fixed point of this estimate as the equilibrium density of ones for this rule. The self-consistent values for p are 0 and $1 \Leftrightarrow 1/\sqrt{3}$, or approximately 42%.

A slightly more detailed approach to the same information would be to set up an evolution matrix, in which the probabilities of each of the cell values are components of a vector, while the elements of the matrix describe the probabilities that one value of the cell evolves into another.

$$\begin{bmatrix} q \\ p \end{bmatrix}' = \begin{bmatrix} p^2 + q^2 & p^2 + 2pq \\ 2pq & q^2 \end{bmatrix} \begin{bmatrix} q \\ p \end{bmatrix}.$$

The eigenvalues of this column stochastic matrix are $\lambda = 1, q(2q \Leftrightarrow 1)$. A matrix of column eigenvectors is

$$U = \begin{bmatrix} p^2 + 2pq & 1 \\ 2pq & \Leftrightarrow 1 \end{bmatrix}$$

Self-consistency is judged as before, with the same equation for $p' = p$; but now additional information about the rate of decay of disequilibrium is available. Disequilibrium is very long lived for the self-consistent value $p = 0$, and vanishes for $p = 50\%$.

8.3 More refined theories

Dresden and Wong[35] showed how to write the rule of evolution of Conway's Life in an algebraic form to which the rules for combining probability distribution functions could be applied. Schulman and Seiden[102] proceeded to obtain an explicit form for the evolution of probability, finding the cummulants hard to deal with, but nevertheless worked out an approximation and applied it to Life.

Wilbur, Lipman, and Shamma[116] decided instead to work with the probabilities of chains or blocks of cells rather than individual cells, including a survey of the self-consistent probabilities of triples of cells according to Wolfram's thirty two "legal" $(2, 1)$ automata in their article. They obtain self-consistency from estimating the probabilities of the different ancestors in the reduced evolution matrix. They use the the same chains they are studying to estimate the probabilities of the extensions required as ancestors, treating the extensions as though they were a Markov process.

That is, starting with n -block probabilities, they use the quotient

$$p_\ell(a, x) = \frac{p(ix)}{\sum_i p(ix)}$$

as the probability that the $(n \Leftrightarrow 1)$ -block chain x can be extended to the n -block chain ax ; in terms of the notation we have already introduced, they assume $p_\ell(ax) =$

$p(ax)/p(x)$. A similar result, $p_r(xa) = p(xa)/p(x)$ is supposed to hold for right extensions. Their working hypothesis is that the same relation serves to extend an n -block probability to an $(n+1)$ -block probability; making two extensions, one on each side, to get the ancestor of a given block, their equations for self-consistent block probabilities read

$$p(z) = \sum_{\varphi(abxcd)=z} \frac{p(abx)p(xcd)}{p(bx)p(xc)} p(bxc).$$

All terms refer to n -block probabilities or their sums.

8.4 Local structure theory

Perhaps the most extensive pursuit of this theme is to be found in Gutowitz, Victor, and Knight's[51] *local structure theory*. It differs from the approach of Wilbur *et.al.* by their use of the theory of probability measures to justify the derivation of their equations for the probabilities of the blocks, even though the final equations are the same. Both groups of authors survey several classes of rules, comparing the theoretical results with empirically observed frequencies. There is general agreement that the probability of a block in one generation should equal the probability of its ancestor in the previous generation; differences arise both from the philosophies and the actual techniques used to obtain the required probability.

To derive their equations, Gutowitz *et.al.* require $2r+1$ consecutive overlapping n -blocks to build up the $n+2r$ -block ancestor of a given n -block from other n -blocks. The choice of a notation in which to express the equations is important; we shall explore certain alternatives. To begin with, suppose that B and X are words, the remaining symbols single letters. The equations for self consistent probabilities then read

$$p(ABC) = \sum_{VWXYZ \in \varphi^{-1}(ABC)} \frac{p(VWX)p(WXY)p(XYZ)}{p(WX)p(XY)}$$

As Gutowitz *et.al.* demonstrated, these equations reduce to the mean field equations for 1-block probabilities, making their local field theory a plausible extension of mean field theory. Since the 1-block, or mean field theory, denominators reduce to the constant value $p(\lambda) = 1$, the mean field theory equations are polynomial equations, whereas the general local structure theory equations involve rational fractions. As a result it is much harder to establish the existence, uniqueness, or stability of their solutions.

Nevertheless, the local structure theory approach makes contact with a very general framework within probability theory, and also presents the basic equations in a symmetrical form admitting a greater variety of interpretations. In particular we can recover the point of view of Wilbur *et.al.* , interpreting the equations as a set of linear equations whose matrix of coefficients depends on the same solutions which it defines. The nonlinearity of the equations is thus of a very special form, reminiscent of the Hartree equations or the Hartree-Fock equations of quantum mechanics, which are generally solved by an iterative process.

Numerically, the equations of Gutowitz *et.al.* can be solved exclusively by iteration. Diagonalizing the coefficient matrix might accelerate convergence, but at the cost of the time expended on diagonalization. It is of greater importance that the form of the equations reveals particular properties of the solution and its convergence that are not otherwise evident.

The ostensible variables are not manifest in the denominators of the equations, but they are readily obtained by using the Kolmogorov consistency conditions. Since the terms of the numerators are associated with links in the de Bruijn diagram, and the denominators with nodes, the whimsically minded might fancy that they see a resemblance to the propagators and Feynman diagrams of field theory. In any event we are going to take up a slightly different interpretation.

8.5 Hartree-Fock approach

The Hartree-Fock-like assumption is to define

$$\mu_V = \frac{p(VWX)}{p(WX)}$$

$$\nu_Z = \frac{p(XYZ)}{p(XY)}$$

both these quantities are positive numbers less than 1; if any numerator is zero, the entire fraction is taken to be zero, even in those cases where the denominator might also be zero.

In fact μ and ν are probabilities which can be used in various ways, such as constructing the Markov matrix for spatial correlations; in an ℓ -stage de Bruijn diagram they give the relative probabilities for the different entering or emerging symbol during a shift.

The self-consistency equations can now be written

$$p(ABC) = \sum_{WXY} \sum_{V,Z} \delta(VW XYZ, \varphi^{-1}(ABC)) \mu_V \nu_Z p(WXY)$$

which presents the appearance of being a system of linear equations if one overlooks the fact that the μ 's and ν 's are not constants, but depend upon the very same unknown probabilities for which one is solving. The function δ is a set-theoretic Kronecker delta, 1 when its arguments coincide, zero when they do not.

The inner sum defines a matrix (writing all indices as arguments, not subscripts):

$$\mathcal{H}(ABC, WXY) = \sum_{V,Z} \delta(VW XYZ, \varphi^{-1}(ABC)) \mu(V) \nu(Z),$$

mapping one vector of probabilities for the ℓ -blocks into another.

For symmetry we have linearized the equations of Gutowitz *et.al.* by taking $p(WXY)$ as the vector component, but we could just as easily have taken one of the other numerator terms; the Kolmogorov conditions are flexible enough to permit μ and ν to retain the same form in either case. Should the occasion arise to do

so, we could distinguish the three different definitions of \mathcal{H} by the adjectives left, central, or right.

To implement the theory of Gutowitz *et.al.*, recall the reduced evolution matrix E_n . For n -chains of cells i and j , it is defined by

$$[E_n]_{i,j} = \begin{cases} 1 & i\epsilon\varphi^{-1}(j) \\ 0 & \text{otherwise} \end{cases}.$$

The probabilistic version of E is just \mathcal{H} defined above, at least for one particular way of estimating the probabilities of ancestors. In any event, many essential properties of \mathcal{H} are determined by E , since both matrices have the same block diagonal structure, no matter whether the nonzero matrix elements are determined self-consistently from the nonlinear local structure theory or otherwise. E_n , which is always fairly crowded for low values of n , becomes sparser and sparser as n increases. One hopes that it tends toward a stable form which could be described analytically; and which might also describe the full matrix \mathcal{H}_n .

8.6 Kolmogorov consistency conditions

The motivation for working with block probabilities rather than cell probabilities is the hope that a better agreement with empirical observations can be obtained, supposing that up to some point longer blocks can better account for correlations between cells than short blocks can. Nevertheless, the empirical quantity which is usually calculated is the probability of individual cells. Sometimes variances or pair probabilities may also be calculated, but it is the density of cells which is the primary concern.

$$\begin{bmatrix} p(00) \\ p(01) \\ p(10) \\ p(11) \end{bmatrix} = \begin{bmatrix} 1000001100010111 \\ 0100000000101000 \\ 0001010010000000 \\ 0010100001000000 \end{bmatrix} \begin{bmatrix} p(0000) \\ p(0001) \\ p(0010) \\ p(0011) \\ p(0100) \\ p(0101) \\ p(0110) \\ p(0111) \\ p(1000) \\ p(1001) \\ p(1010) \\ p(1011) \\ p(1100) \\ p(1101) \\ p(1110) \\ p(1111) \\ p(1100) \\ p(1101) \\ p(1110) \\ p(1111) \end{bmatrix}$$

Figure 8.1: Two-block probabilities summed from four-block probabilities.

Consolidating these two equations produces a 4×4 matrix of coefficients

$$\mathcal{H}_2 = \left[\begin{array}{c|cc|cc} & 00 & 01 & 10 & 11 \\ \hline 00 & \frac{p(00)}{p(0)} \frac{p(00)}{p(0)} & \frac{p(10)}{p(0)} \frac{p(11)}{p(1)} & \frac{p(11)}{p(1)} \frac{p(01)}{p(0)} & 1 \\ 01 & \frac{p(00)}{p(0)} \frac{p(01)}{p(0)} & \frac{p(10)}{p(0)} \frac{p(10)}{p(1)} & \frac{p(11)}{p(1)} \frac{p(00)}{p(0)} & 0 \\ \hline 10 & \frac{p(10)}{p(0)} \frac{p(00)}{p(0)} & \frac{p(00)}{p(0)} \frac{p(11)}{p(1)} & \frac{p(01)}{p(1)} \frac{p(01)}{p(0)} & 0 \\ 11 & \frac{p(10)}{p(0)} \frac{p(01)}{p(0)} & \frac{p(00)}{p(0)} \frac{p(10)}{p(1)} & \frac{p(01)}{p(1)} \frac{p(00)}{p(0)} & 0 \end{array} \right],$$

in which it can be verified by inspection that $p(0)$ is consistently defined by either $p(0) = p(00) + p(01)$ or by $p(0) = p(00) + p(10)$; or that equivalently $p(10) = p(01)$ on the right hand side implies the same relation for the new values on the left hand side of the equation.

It is easier to give the general proof in symbolic form, having first written the field equations in terms of the merged product:

$$p(Q) = \sum_{\varphi(X \vee Y \vee Z) = Q} \frac{p(X)p(Y)p(Z)}{p(X \wedge Y)p(Y \wedge Z)}$$

If it is intended to write $Q = AB$ so as to obtain

$$p(B) = \sum_A p(AB)$$

we need to know that

$$\varphi^{-1}(AB) = \varphi^{-1}(A) \vee \varphi^{-1}(B),$$

If $\varphi^{-1}(A)$ is incorporated into X we introduce W

$$X = \varphi^{-1}(A) \vee W$$

and note that

$$\sum_A \varphi^{-1}(A) \vee W = X \wedge Y,$$

so that

$$p(B) = \sum_{\varphi(Y \vee Z) = B} \frac{p(Y)p(Z)}{p(Y \wedge Z)}.$$

Since an entirely symmetrical expression results when Q is factored into $Q = BC$, comparing the two establishes the consistency of $p(B)$ after iteration, supposing that it was consistent before. The essential points in the proof are, first, that running through all initial letters in the sum for $p(B)$ guarantees a traversal of all the initial letters in the ancestors, and second, that the denominators in the probabilities are cancelled by the sums in the numerators as a consequence of the consistency hypothesis.

8.7 The vector subset diagram

If $\alpha(x)$ is the number of ancestors of x , we have

$$0 \leq \alpha(xy) \leq \alpha(x)\alpha(y),$$

because the $2r$ cells at which the ancestors of x overlap the ancestors of y when x and y are joined must coincide. Moreover, there can be no ancestor of xy which does not begin with an ancestor of x , nor which does not terminate with an ancestor of y .

The overlapping of neighborhoods which must be taken into account in calculating ancestors can be summarized very concisely by the relation

$$\varphi^{-1}(ax) = \varphi^{-1}(a) \vee \varphi^{-1}(x)$$

when a is a letter and x is a word (or the reverse). If both are words a similar relation holds, but the merged product was not defined with quite the generality required to restrict the overlap involved.

Since the determination of the ancestors of a given word plays an important role in all aspects of the calculation of evolutions, it is fortunate that the subset diagram used to locate excluded words can be generalized slightly to a form which will provide both the number of counterimages of a word and the counterimages themselves.

The subset diagram was derived from the $2r$ -stage de Bruijn diagram whose links were labelled for the neighborhoods of a (k, r) automaton because the links also correspond to the cells that evolve from the neighborhoods which the links represent. Thus following out a path according to the evolved cells automatically yields the ancestor involved; the subset construction was invoked because of a need to obtain *all* the possible paths, and moreover to obtain them systematically.

Links were the only information recorded in the subset diagram because interest in the diagram was limited to knowing which combinations, or subsets, of nodes could be linked to get a given evolution. However, if each node is assigned a 2^{2r} -ple (quadruple for a $(k, 1)$ automaton), the exact linkage of neighborhoods could be shown, not just the mere indication that such a linkage exists.

If we use Rule 126 as an example, we need the table of transitions

$$\begin{array}{cccccccc} 111 & 110 & 101 & 100 & 011 & 010 & 001 & 000 \\ 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \end{array} .$$

In other words all neighborhoods evolve into ones with the exception of 111 and the quiescent neighborhood 000.

The de Bruijn diagram for this rule is

node	0 leads to	1 leads to
<i>A</i>	<i>A</i>	<i>C</i>
<i>C</i>	ϕ	<i>D, F</i>
<i>D</i>	ϕ	<i>A, C</i>
<i>F</i>	<i>F</i>	ϕ

while the vector subset diagram will have the form

subset	0 leads to	1 leads to
$\{ACDF\}$	(A, ϕ, ϕ, F)	$(D, A + D, C + F, C)$
$\{AF\}$	(A, ϕ, ϕ, F)	(ϕ, A, F, ϕ)
$\{CD\}$	(ϕ, ϕ, ϕ, ϕ)	(D, D, C, C)
$\{\}$	(ϕ, ϕ, ϕ, ϕ)	(ϕ, ϕ, ϕ, ϕ)

The shortest excluded word for this rule is 010; generally any sequence ending in 010 is excluded. To determine the subset to which a quadruple leads we have only to determine its non- ϕ components. To obtain the number of ancestors for any sequence, say 1101, we begin with the quadruple (1,1,1,1), and apply the formula within the quadruple belonging to each digit in turn. Finally we sum the components.

node	count	
$ACDF$	(1, 1, 1, 1)	(00, 01, 10, 11)
1 $ACDF$	(1, 2, 2, 1)	(100, 001 + 101, 010 + 110, 011)
1 $ACDF$	(2, 3, 3, 2)	(0100 + 1100, 1001 + 0101 + 1101, 0010 + 1010 + 0110, 0011 + 1011)
0 AF	(2, 0, 0, 2)	(01000 + 11000, $\phi, \phi, 00111 + 10111$)
1 CD	(0, 2, 2, 0)	($\phi, 010001 + 110001, 001110 + 101110, \phi$)

Thus we conclude that there are 4 ancestors of 1101 according to Rule 126. If we wanted the actual ancestors, we should have performed the arithmetic symbolically, as shown in the last column of the table.

8.8 Estimating the number of ancestors

Instead of working with formula bearing quadruples, we could have represented the links in the vector subset diagram by square matrices; most of them would be singular, but calculations involving them could be carried out entirely in terms of matrix algebra. Moreover, it would be possible to ask selective questions about ancestors by examining individual matrix elements of the product; a multiple product is already indexed by the $2r$ letters with which the ancestor commences and the $2r$ letters with which it ends.

Even more useful is to construct the connection matrix of the vector subset diagram; links to the empty set can be omitted because of the triangular form of the full matrix. For the example of Rule 126, this matrix has the form

. 1 . . .	1
. . 1 1
1 1
. . 1 1
<hr/>		
.	1 1 . . .
.
.
.
<hr/>		
.
. . 1 1
1 1
.

This matrix contains four distinct non-zero submatrices

$$\alpha = \begin{bmatrix} \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 \\ 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \end{bmatrix}, \beta = \begin{bmatrix} 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \end{bmatrix}, \gamma = \begin{bmatrix} \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \end{bmatrix}, \delta = \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 \\ 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}.$$

which are the matrix representations of the quadruples of the vector subset diagram. It is convenient to regard them as hypercomplex numbers formed from a basis consisting of the following 8 4×4 matrices

$$\beta = \begin{bmatrix} 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \end{bmatrix}, \gamma = \begin{bmatrix} \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \end{bmatrix}, \epsilon = \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \\ \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}, \zeta = \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \\ 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix},$$

$$\eta = \begin{bmatrix} \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot \end{bmatrix}, \theta = \begin{bmatrix} \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot & \cdot \end{bmatrix}, \iota = \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \end{bmatrix}, \kappa = \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix},$$

according to the definitions

$$\begin{aligned} \alpha &= \gamma + \zeta + \epsilon \\ \delta &= \zeta + \epsilon. \end{aligned}$$

These basis elements obey the multiplication table

	β	θ	ι	ϵ	γ	η	ζ	κ
β	β	θ	ϕ	ϕ	γ	η	ϕ	ϕ
θ	θ	β	ϕ	ϕ	η	γ	ϕ	ϕ
ι	ϕ	ϕ	ι	ϵ	ϕ	ϕ	ζ	κ
ϵ	ϕ	ϕ	ϵ	ι	ϕ	ϕ	κ	ζ
γ	ϕ	ϕ	γ	η	ϕ	ϕ	θ	β
η	ϕ	ϕ	η	γ	ϕ	ϕ	β	θ
ζ	ζ	κ	ϕ	ϕ	ϵ	ι	ϕ	ϕ
κ	κ	ζ	ϕ	ϕ	ι	ϵ	ϕ	ϕ

In general it is probably easier to work with the connection matrix and its powers than to deal with the algebra of hypercomplex numbers. Sometimes special properties of the algebra manifest themselves, as they do here. Every basis matrix has an even element sum, and likewise an even trace. This shows that there must always be an even number of ancestors of whatsoever open chain, and also of every ring. Thus there can be no orphan periods among the rings, in which no transients lead into the ring.

The following table shows the numbers of ancestors, of both open sequences, and rings, of three cells. To count open sequences all the elements of the final product matrix are to be summed; for rings only the trace is involved because the ring cannot close unless the loose ends coincide.

cells	000	001	010	011	100	101	110	111
matrix product	$\beta\beta\beta$	$\beta\beta\gamma$	$\beta\gamma\phi$	$\beta\gamma\delta$	$\alpha\beta\beta$	$\alpha\beta\gamma$	$\alpha\alpha\beta$	$\alpha\alpha\alpha$
total	2	2	0	4	2	2	4	16
ring	2	0	0	0	0	0	0	6

8.9 Trivial solutions

It would be surprising if the equations for self consistency had unique solutions. In fact, even for one-block equations, the quiescent rules will always have a solution for the quiescent state. Generally there is another solution (except for Wolfram's classes i and ii, which practically exclude additional solutions by definition.) If there is more than one solution it will be found that some of them are stable while others are unstable; in fact there is a whole theory surrounding the existence and properties of the fixed points of nonlinear equations.

Whatever may be the nature of the solutions of the local field theory equations in general, there are some solutions which can be foreseen. If the length of the block corresponds to the length of a cycle for a given automaton, then the members of the cycle can be assigned equal probabilities and other blocks can be assigned probability zero. We have two cases to consider—the blocks obtained by symmetry but belonging to the same phase of evolution, and the blocks belonging to different phases. All the phases of the same block will form a cyclic submatrix of the local structure matrix, while different symmetry images will generate additional diagonal blocks. In any event, a given block has exactly one counterimage within the sequence of evolution, and everything has been assigned equal probabilities.

There are some patterns of cyclic evolution which have no other ancestors than their immediate predecessors in the cycle of evolution—the still life of alternating zeroes and ones in Rule 22 for example, as well as the cycles of period seven and eleven. Presumably they represent unstable fixed points in the parameter space of blocks. Other cycles may be the endpoint of various transients, and thus be somewhat more stable. Generally the self-consistent probabilities for blocks will not correspond to the actual densities of cells for any particular pattern of evolution, but rather will be formed from a composite of all of them.

It is also not excluded that there will be probabilities which are not self-consistent, but rather which alternate between a finite number of values. This behaviour tends to occur for rules which do not have quiescent states, so that alternation between two or more backgrounds can occur.

Chapter 9

Positive matrices

Starting with the evolution and de Bruijn matrices, and continuing with their probabilistic versions, we find that we are dealing with a specialized class of matrices all of whose elements are positive—or more accurately, non-negative. The first two of these classes of matrices is more restricted, inasmuch as their elements must be integers, but it is not especially easy to obtain much advantage from that particular characteristic. However, positive matrices—including those with integer elements—enjoy two properties which can be exploited to considerable advantage. Avoiding some limiting cases arising from unfavorable groupings of zero matrix elements, these properties are:

- there is a unique maximum eigenvalue, whose value is bracketed by the row sums of the matrix (as well as the column sums), and
- whose eigenvector can be normalized so that all its components are strictly positive. That is, all have the same sign and none are zero.

9.1 Gerschgorin's disks

The fundamental tool for estimating the size of eigenvalues is the eigenvalue equation itself, written out in terms of components of the eigenvectors. Thus, from the matrix equation $MX = \lambda X$ there follows

$$\lambda x_i = \sum_{j=1}^n m_{ij} x_j.$$

Since the zero vector is never considered to be an eigenvector, there is at least one component x_i which is not zero; let us select the largest among the non-zero components, move the corresponding diagonal element to the left hand side of the equation, and apply the triangle inequality. If there is any ambiguity in selecting the largest component, any one of them will do. Then

$$|\lambda - m_{ii}| |x_i| \leq \sum_{j \neq i} |m_{ij}| |x_j|.$$

The purpose of selecting the largest component and ensuring that it was non-zero was to allow division of this equation by $|x_i|$ to obtain the factors $|x_j|/|x_i|$ which are all less than, or equal to, 1 (even when $i = j$). Thus

$$|\lambda \leftrightarrow m_{ii}| \leq \sum_{j \neq i} |m_{ij}| \frac{|x_j|}{|x_i|}.$$

Since the inequality can only be enhanced by making the terms on the right hand side larger by dropping their small multipliers, we finally obtain the basic result

$$|\lambda \leftrightarrow m_{ii}| \leq \sum_{j \neq i} |m_{ij}|,$$

which applies even for complex matrices or for real matrices with complex eigenvalues. The complex setting allows a picturesque description of the results with respect to the geometry of the plane. Regarding the column sum excluding the diagonal element as a radius and the diagonal element as a center, we have found a circle which surely contains the eigenvalue.

In general, without having actually found an eigenvector, which would be its largest component would be rarely be evident; but if all the different disks were joined together, no eigenvalue could escape the collection. There is even some choice of the precise disks to be used, according to whether the diagonal element is taken as part of the radius or as part of the center; the procedure shown yields the smallest disks but, at times a common center might be preferable. Of course, the smallest disks are centered on the diagonal elements.

This estimate is due to Gerschgorin, whose name is associated with the disks. Generally the disks all intersect, and there is no correspondence between eigenvalues and disks because of the uncertainty as to which would be the largest component of any given eigenvector. Sometimes there will be disjoint clusters of disks for which continuity arguments give each cluster its quota of eigenvalues. That is, if the diagonal elements are retained and the remainder multiplied by a small parameter, continuity of the eigenvalues with respect to the coefficients of the matrix will locate the eigenvalues within small disks surrounding the diagonal elements. As the parameter is increased, eigenvalues can only wander amongst disks which overlap, but the number within a given cluster must remain constant.

Using the origin as a common center for the Gerschgorin disks leads to expressing the bounds directly in terms of row sums:

$$|\lambda| \leq \sum_{j=1}^n |m_{ij}|;$$

since the inequality must hold for at least one row, we could always select the worse case with full confidence that it expresses a valid bound.

9.2 Eigenvalues on the boundary

The approximations used to derive the Gerschgorin limits would seem to be fairly generous. If any one of the eigenvalues actually lies on the circumference of a

Gerschgorin disk, the implied equality for that one equation could presumably be used to draw additional conclusions; even close proximity to the boundary might possibly be a source of further information.

Consider the point in the derivation at which the ratios $|x_j|/|x_i|$ were replaced by 1, and recall that i was the index of the largest component of the eigenvector; this substitution would surely alter the sum unless the matrix element m_{ij} multiplying it were zero. If *all* the remaining offdiagonal matrix elements were zero, we could stop and think about reducing the matrix; otherwise we have forced a number of components of the eigenvector to be equal in absolute value. But this means that there are other equations in which this same largest component occurs, but with other indices.

The same argument forces equality for any further components related through non-zero matrix elements, and we arrive at the eventual conclusion that all the components of the eigenvector have the same absolute value, and that *all* the Gerschgorin disks intersect at least in the eigenvalue in question. The only exception would arise from a matrix whose diagram were not connected, so that either the class of matrices involved must be restricted, or the conclusion must be confined to the components belonging to connected parts of the matrix.

Bearing this restriction in mind, we conclude that the Gerschgorin bounds can only be realized for matrices having uniform row sums, and thus any nonuniformity is certain evidence that the limits cannot be reached. Nevertheless, we have only derived a necessary—not a sufficient—condition, and it may well happen that the eigenvalues of matrices with uniform row sums lie well within the Gerschgorin disks. Examining simple 2×2 matrices with elements of mixed signs will quickly confirm this statement. The lack of sufficiency results from not yet having taken into account the liberalizing influence of the triangle inequality on Gerschgorin's inequalities; but if all matrix elements were positive, the uniformity of the row sums would immediately establish the vector with unit components as an eigenvector and the row sum as its eigenvalue.

Given that our arguments are insensitive to scalar multiplication of the whole matrix by a complex factor, we cannot postulate that we are working with a positive matrix. However, if we note that a collection of vectors of fixed lengths cannot reach their maximum sum unless they are all parallel, we see that the only way that the natural eigenvalue equation and the one resulting from inserting absolute values can hold simultaneously is for all the terms to have a common phase.

We can begin by removing phase factors from the components of the eigenvector by performing a similarity transformation via diagonal matrices bearing the phase factors. The rows of the resulting matrix must then have constant phase, which is necessarily the phase of the eigenvalue. If the eigenvalue were not zero, this would be the phase of the whole matrix, and could be discarded by treating it as a constant factor.

Thus non-uniform row sums guarantee that the largest row sum is not an eigenvalue, but even uniformity will only make the common row sum an eigenvalue when the matrix is effectively a positive matrix.

Even though we did not assume that a possible boundary eigenvalue was the largest eigenvalue, that conclusion seems inescapable, inasmuch as we have seen

that a connected matrix possessing a boundary eigenvalue is essentially positive, for which summing row elements with anything other than equal weights will necessarily produce a smaller sum and with it a smaller eigenvalue. Likewise, we seem to be led to the uniqueness of this eigenvalue since only one eigenvector can be associated with it.

9.3 Minimax principle

Using Dirac's notation for row and column vectors, the equations

$$\lambda \leq \frac{\langle x|M|y\rangle}{\langle x|y\rangle},$$

used in deriving the Gerschgorin disks would arise directly from the so-called Rayleigh quotient, were it not for the fact that an inequality can only be stated for real quantities. However the necessary absolute values entering the equations could be avoided, given an assurance that all the terms were real, as perhaps might be anticipated for the largest eigenvalue of a real matrix.

It is generally understood that eigenvalues and eigenvectors satisfy a variational principle; for symmetric or hermitean matrices this leads to bounds involving the Rayleigh coefficient or to the Courant minimax principle. If we suppose that $\langle x+\delta x|$ and $|y+\delta y\rangle$ are vectors deviating slightly from $\langle x|$ and $|y\rangle$ respectively, and ignoring quantities of second order we know that

$$\begin{aligned} \frac{\langle x+\delta x|M|y+\delta y\rangle}{\langle x+\delta x|y+\delta y\rangle} &= \frac{\langle x|M|y\rangle}{\langle x|y\rangle} + \\ &\quad \langle \delta x| \left(M|y\rangle \Leftrightarrow \frac{\langle x|M|y\rangle}{\langle x|y\rangle} |y\rangle \right) + \\ &\quad \left(\langle x|M \Leftrightarrow \langle x| \frac{\langle x|M|y\rangle}{\langle x|y\rangle} \right) |\delta y\rangle + \\ &\quad \dots \end{aligned}$$

Requiring the vanishing of the first order terms independently of $\langle \delta x|$ and $|\delta y\rangle$ establishes the Rayleigh quotient as an eigenvalue, providing that $\langle x|$ and $|y\rangle$ are its corresponding eigenvectors. However it would be preferable to work directly from explicit inequalities and not have to extrapolate from extremals to bounds.

Consequently, without yet specifying P and Q , define

$$\begin{aligned} r &= \max_{y \in Q} \min_{x \in P} \frac{\langle x|M|y\rangle}{\langle x|y\rangle}, \\ s &= \min_{y \in Q} \max_{x \in P} \frac{\langle x|M|y\rangle}{\langle x|y\rangle}. \end{aligned}$$

If P were the set of coordinate vectors and Q were the set of vectors with positive components, the Rayleigh quotient would correspond to the right hand side of the Gerschgorin equations. Interchanging the roles of P and Q in these same

equations would produce results for column sums equivalent to those which we would otherwise obtain for row sums.

With these choices of P and Q , it is easy to follow the rationale behind the definition of r ; vector by vector, we decide which component suffers the least magnification, even if it is very, very small. At least we know that the vector as a whole will receive at least that much magnification. No one can blame us for finally choosing the vector with the greatest overall magnification, the value of which is r .

Nevertheless, the conservatism involved in singling out the component with the least magnification means that the entire vector will always be magnified slightly more than the worst case estimate unless in fact each and every component is magnified by the same amount—in which case we have an eigenvector on our hands. Likewise, supposing the magnification is not uniform for all components, we could construct another vector whose minimum magnification was slightly larger by simply reducing the size of the minimally amplified component relative to the others. However, this latter procedure would definitely fail for the vector of maximum minimal amplification, so it would just have to be an eigenvector.

Similar considerations apply to the definition of s , since we would be estimating the maximum magnification that M could possibly impart to a transformed vector, and examining the vector whose estimate was the least extravagant, only to discover that the estimate was quite valid throughout the whole vector, and so to conclude once again that we had an eigenvector and that s was its eigenvalue.

To be certain of our conclusions we need to be assured that zero matrix elements do not occur in inconvenient places, and also that the limits involved in assuming the existence of maxima or minima over sets of vectors exist; the reason that the Rayleigh quotients contain the denominators which they do is to allow us to confine our attention to vectors of fixed norm, and thus to compact sets. The question of the possible configurations of zeroes does indeed reward more careful consideration, because of the possibility of finding disconnected portions or cyclic structures within M . Nevertheless, the central role is played by those matrices whose elements are strictly positive.

9.4 Largest eigenvalue

Whereas it is clear that the definition of r in the last section yields the largest eigenvalue with a positive eigenvector, we would like to know that it is the largest possible eigenvalue, and that it is unique. This information can be obtained somewhat indirectly by noticing that any eigenvector lacking zero components defines an invertible diagonal matrix R , containing those same components on its diagonal. Let X be the eigenvector, U the vector with unit elements. Then $X = RU$, so that the eigenvalue equation

$$MX = \lambda X$$

is equivalent to

$$R^{-1}MRU = \lambda U$$

which means that λ is the common row sum of the matrix $R^{-1}MR$. Since the properties of boundary eigenvalues of Gerschgorin disks establish the desired result

for the equivalent matrix $R^{-1}MR$ they establish it for M as well.

Since M does not necessarily have uniform row sums, we are now interested in relating λ to the row sums, the largest of which we know to be an upper bound to λ , yet unequal to it. Here the variational definition of r is useful; let us write $r(M)$ to acknowledge the dependence of r on M , and increase M in any way whatsoever, even by increasing one single matrix element. Then each quantity $\langle x|M|y \rangle$ must strictly increase, and with it both the minima and maxima, leading to a strictly increased maximum eigenvalue.

At this point we could alter some of the matrix elements of M to obtain a new matrix A , whose row sums were uniformly the minimal row sum of M . It is important to do this exclusively by reducing matrix elements in M . Conversely, we could selectively increase the elements of M to obtain a matrix B whose row sums were uniformly the row sums of M , but we already have the result which this would imply.

Applying the theorem on boundary eigenvalues to the matrices A , M , and B , we obtain both upper and lower bounds to λ , defined as r above:

$$\min_i \sum_{j=1}^n m_{ij} \leq \lambda \leq \max_i \sum_{j=1}^n m_{ij}.$$

However both inequalities are strict unless the bounds coincide. Although an estimate of the gaps would be useful, often the assurance that they exist is sufficient.

9.5 Second largest eigenvalue

The largest eigenvalue of a positive matrix is unique and belongs to eigenvectors with positive components, both the row eigenvector and the column eigenvector. Thus any other eigenvectors must have mixed signs both in their real and possible imaginary parts. Knowledge of the relative magnitude of the second eigenvalue of a matrix is often convenient, for example in judging the rate of convergence of most procedures for diagonalizing the matrix. The traditional way to obtain this information is to subtract the contribution of the largest eigenvalue from the matrix, followed by an estimate of the largest eigenvalue of the remaining matrix. There will often be a pair of second largest eigenvalues, complex conjugates on account of belonging to a real matrix; if so they must be considered together.

Removal of the largest eigenvalue requires an explicit knowledge of both the largest eigenvalue and its eigenvectors. It is convenient to avoid this explicit knowledge by supposing the matrix to have been previously transformed to stochastic form; if it is already stochastic, then no preparation is required. The transformation to a column stochastic matrix is accomplished by incorporating the elements of its row eigenvector into the matrix itself by column divisions and row multiplications, finally dividing the entire matrix by its eigenvalue.

Even so, we would still need to know the components of the column eigenvector to achieve the traditional reduction. By substituting a more readily available vector for the unknown column eigenvector we can still obtain useful estimates. If the matrix were stochastic to start with, no additional knowledge would be required.

A plausible choice, which would still leave matrix elements of uniform sign, would be to form a vector by taking the least element in each row, or alternatively, the greatest element in each row.

To recapitulate, suppose that M is a column stochastic matrix so that 1 is its largest eigenvalue, with U^T , the row of ones, as its eigenvector. We are seeking an upper bound for λ , which is another eigenvalue with column eigenvector X , and we have formed the column C by taking the least element from each row of M . Since

$$U^T X = 0$$

we know that both

$$MX = \lambda X$$

and

$$(M \Leftrightarrow CU^T)X = \lambda X.$$

Either of these two matrices can be used to obtain bounds for λ since it is an eigenvalue of each. Taking Gerschgorin disks from the column sums, we see that the column sums of M are unity and those of C are all equal and equal to the sum of the minimal elements of each row. Thus

$$|\lambda| \leq 1 \Leftrightarrow \sum_{i=1}^n \min_{j=1}^n m_{ij}.$$

If maximal elements were utilized, there would be a reversal of sign and

$$|\lambda| \leq \left(\sum_{i=1}^n \max_{j=1}^n m_{ij} \right) \Leftrightarrow 1.$$

Obviously one would choose the more restrictive of the two bounds.

If there is a considerable discrepancy in the sizes and an unfavorable distribution of matrix elements of M , the bounds provided by these two inequalities may not be very useful; for example the probabilistic de Bruijn matrices generally contain zeroes in every row and column so that the minimal inequality is vacuous. Likewise, the concentration of nonzero elements makes them so large that the maximal inequalities give worse limits than the knowledge that the matrices are stochastic. Nevertheless, powers of the de Bruijn are better behaved, and are more suitable for estimates.

9.6 Averaging and convergence

The fact that the largest eigenvalue of a positive matrix is isolated means that powers of the matrix should converge to a multiple of the idempotent matrix formed from the corresponding eigenvectors. By normalizing the eigenvalue and one of the eigenvectors, as has already been done for stochastic matrices, the limiting form of high powers is even further simplified. The rate of convergence depends on the ratio between the largest and second largest eigenvalue, or directly on the second if the matrix is stochastic. Conversely, one way of estimating the second eigenvalue

is to study the behaviour of matrix powers, preferably already having transformed the matrix to stochastic form.

This can be done by commencing with a positive vector $X^{(0)}$ and defining recursively

$$\begin{aligned} X^{(k+1)} &= MX^{(k)} \\ \lambda_k^+ &= \max_i \frac{x_i^{(k+1)}}{x_i^{(k)}} \\ \lambda_k^- &= \min_i \frac{x_i^{(k+1)}}{x_i^{(k)}}. \end{aligned}$$

Then it will be found that the largest eigenvalue satisfies

$$\lambda_0^- \leq \lambda_1^- \leq \dots \leq \lambda \leq \dots \leq \lambda_1^+ \leq \lambda_0^+,$$

which confines the eigenvalue to a succession of ever smaller intervals. The outermost, widest, interval follows from the minimax characterization of the largest eigenvalue. So do each of the others, by specializing the choice of the column vector in the same characterization, but the important part of this result is that the intervals are nested.

Nesting is a consequence of multiplying successive vectors by M , but it is sufficient to see how multiplication by M^2 relates to multiplication by M . Moreover it will be more understandable to write the definition of λ_k^+ in terms of

$$MX \leq \lambda_k^+ X,$$

the vector inequality being understood as holding for each component, and λ_k^+ being understood as the greatest constant for which the inequality holds. Then we would have

$$\lambda MX \leq M(MX) \leq \lambda_k^+ MX \leq (\lambda_k^+)^2 X,$$

the first inequality holds since λ itself serves for a general vector such as MX . Given that $\lambda \leq \lambda_k^+$, we conclude that $\lambda_{k+1}^+ \leq \lambda_k^+$, since it is the factor which applies to this final equation.

There is a corresponding convergence of vectors which is interesting, especially since the arguments apply equally well to any product of stochastic matrices and not just to powers of a single matrix. For simplicity suppose that the family of matrices is column stochastic, which means that U^T , the row of 1's, is a uniform eigenvector for the family. Then the dominant part of such a matrix would be the vector product of some column X with the row U^T . The action of U^T its action in multiplying any column would be to "average" the components of the column, just as its action on a matrix would be to form a new row of column averages. Strictly, an average would also include division by the number of terms summed.

One of the outstanding properties of averaging is its tendency to reduce variance. In the present context this would mean that the product of a large number of positive stochastic matrices would tend to a matrix with uniform rows, independently of whether the same matrix or different matrices were being multiplied.

In the particular case of powers of a single matrix, including an additional factor would not change the limit, which would have to be a matrix of eigencolumns.

In verifying the details of the averaging process it is found that certain estimates are required for the size of the matrix elements. In particular, permitting very small or zero elements would be nice for certain applications, but would either complicate or invalidate the proofs. Here we only present one of the preliminary theorems, which describes the reduction in range of the elements of a vector with real components when multiplied by a stochastic matrix.

Suppose that M is the matrix, that X is a real vector whose algebraically largest and smallest components are A and a respectively, and that the corresponding extreme elements of MX are B and b , respectively. We also require the smallest element of M , which we might call ε . For the purpose of argumentation we introduce another vector Y constructed from X by replacing all of its elements save one of the smallest by A . Starting from the evident inequality

$$X \leq Y$$

we obtain

$$MX \leq MY$$

and a requirement to estimate MY . Each component will have the form $\alpha a + (1 - \alpha)A$, which can be rewritten $A \Leftrightarrow \alpha(A \Leftrightarrow a)$. The value of α will vary from component to component, but it is always greater than ε , so

$$B \leq A \Leftrightarrow \varepsilon(A \Leftrightarrow a).$$

The easiest way to get an estimate for b is to repeat the same argument for $\Leftrightarrow X$, which reverses all comparisons. Combining the result,

$$\Leftrightarrow b \leq \Leftrightarrow a \Leftrightarrow \varepsilon(\Leftrightarrow a + A).$$

with the bound for B , we find

$$B \Leftrightarrow b \leq (1 - \varepsilon)(A \Leftrightarrow a).$$

By its nature, we have $\varepsilon \leq 1/2$, so the range of variation of the components of MX is necessarily less than the corresponding range for X .

If we could rely on a fixed lower limit to ε for an entire family of matrices, it is clear that X would gradually be reduced to a multiple of U if it were multiplied by increasingly larger numbers of matrices of the family.

9.7 Non-negative matrices

The presence of zeroes in an otherwise positive matrix complicates the derivation of its properties. The sporadic presence of zeroes does not greatly complicate the derivations nor change the results; on the contrary certain systematic patterns of zeroes can alter the conclusions considerably. Fortunately it is relatively easy to predict the influence of zeroes, and to state the correct conclusions. Displaying the

graph associated with the matrix leads to one of the most concise descriptions that seems to be possible.

The essential point is that zero elements are tolerable in a matrix if they disappear from a power of the matrix. Since powers of a matrix have the same eigenvectors as the original, and powers of the original eigenvalues, the necessary details of proofs can be worked out for the matrix power and then attributed to the matrix itself. Such circumvention could fail in two evident ways: either the zeroes persist through all the powers, or they simply keep shifting around. Both circumstances are fairly evident from an inspection of the graph of the matrix.

In the former case, there are nodes in the graph which are not accessible from one another, so that the basic requirement is for a connected graph. If the nodes are grouped according to their accessibility, inaccessible nodes are reflected in a triangular structure of the matrix which can be handled by partitioning the matrix and applying the relevant theorems to the resulting submatrices. Thus the solution to the problem is to treat disconnected portions of the graph separately, just as it is well to separate out transient parts of the graph and treat them separately from the others.

Even if the graph is connected, it may result that links of only certain lengths may exist; consider for example a graph with two subsets of nodes such that the only direct links occurred between nodes of different sets, none between two nodes of the same set. This structure signifies a partitioned matrix whose diagonal submatrices are zero. Its square would have diagonal submatrices but zero antidiagonal submatrices, while its cube would once again have the original form. But no power at all would be free of blocks of zeroes. Fortunately this alternative has a particularly elegant solution.

For example, the example cited refers to a matrix of the form

$$M = \begin{bmatrix} O & A \\ B & O \end{bmatrix},$$

but there exists another matrix

$$J = \begin{bmatrix} I & O \\ O & \Leftrightarrow I \end{bmatrix},$$

for which $MJ = \Leftrightarrow JM$. If

$$X = \begin{bmatrix} U \\ V \end{bmatrix}$$

is an eigenvector of M belonging to λ , then JX is another belonging to the eigenvalue $\Leftrightarrow\lambda$, as the use of the anticommutation relation readily shows. Thus all nonzero eigenvalues occur in negative pairs, with eigenvectors partitioned so that corresponding components differ at most in sign, in appropriate places.

The squared matrix M^2 is block diagonal, of the form

$$M^2 = \begin{bmatrix} AB & O \\ O & BA \end{bmatrix},$$

but the diagonal components AB and $BA = A^{-1}(AB)A$ are equivalent as long as A is square and non-singular. If it is not, a subspace belonging to eigenvalue 0 can be partitioned off, and the conclusions applied to the complementary space.

Generally M would contain a longer superdiagonal of submatrices, higher roots of unity would occupy the diagonal of J , and the canonical form would require k equivalent copies of a matrix of more tractable form. All eigenvalues would occur in cycles, related by k th roots of unity.

Chapter 10

Zeta function

The evolution and de Bruijn matrices are sparse, both in their normal and in their probabilistic versions; that is a consequence of their origins as connectivity matrices of diagrams of one type or another. It is therefore natural to seek a description of such matrices in terms of the visual properties of the same diagrams, one of the most evident of which is the network of cycles or loops which they contain. It is even more satisfying if the description will accommodate the accumulation of a variety of matrices into families. A recent article of Cvitanovic [31] reveals the process, as applied to an analysis of strange attractors.

10.1 Counting loops

It is convenient to begin with matrices with integer elements, and even more precisely with the connection matrices for diagrams, such as those which form the evolution matrices or the de Bruijn matrices. There is a mutual relationship between matrices and diagrams. To begin with, a diagram can be described in various ways by matrices, one of the most natural being that in which the rows and columns of the matrix are indexed by the nodes of the diagram. The matrix elements are then zero or one, according to whether the node indexing the row is linked to the node indexing the column. Such a matrix need not be symmetric, inasmuch as the links may be directed and not run in both directions.

It is even possible to accommodate multiple links between nodes, by allowing the matrix elements to be positive integers rather than confining them to just the two values zero and one. This interpretation is useful because it allows powers of the connection matrix to be interpreted in terms of the number of composite paths joining two nodes, the power of the matrix determining the number of links involved. Whether or not one would want to suppose that negative matrix elements denote a link in the reverse direction depends upon whether one expects to cancel a path by retracing it or not, a complication which we shall avoid.

Conversely it is possible to represent a square matrix by a diagram, by setting up nodes to correspond to the dimensions of the matrix, and then introducing a link wherever there are non-zero matrix elements. Of course, if none of the

matrix elements is zero, every node will then be linked to every other. Although the diagrammatic representation would be most interesting for disentangling the relationships between the elements of sparse matrices, complete mutual accessibility is nevertheless an important special case. In particular, given a finite diagram in which every node can be reached from any other node through a long enough chain of links, there will be a corresponding power of the connection matrix in which no zeroes remain.

Diagonal elements of the connectivity matrix represent nodes linked to themselves. It might seem at first sight that a node would naturally be linked to itself, but it should be borne in mind that links and nodes are two different things, and that even null links have to be shown explicitly. Diagonal elements of powers of the connection matrix are probably easier to understand; aside from possible null links, they would represent loops in which a node was connected to a succession of others, the final link nevertheless returning to the original node.

Generally, the ij th element of the k th power of the connectivity matrix M counts the number of paths with exactly k links leading from node i to node j . The trace, $Tr\{M^k\}$, gives the total number of paths leading back to their point of origin, but under favorable circumstances each path is counted once for each of the k nodes which it contains. This is the case for loops which do not retrace a part of themselves, but not otherwise. Thus

$$N(k) = \frac{1}{k} Tr\{M^k\}$$

is a function which would represent the number of *loops* of length k in the diagram whenever the non-redundancy criterion was met.

Sometimes useful information can be gained by constructing a power series whose coefficients count some quantity or other, particularly if the count can be related to a convolution of the coefficients in some other series. We might begin by forming the series

$$m(t) = \sum_{k=0}^{\infty} M^k t^k.$$

By either ignoring convergence, or hoping for the best, the series is seen to be equivalent to

$$m(t) = (I \Leftrightarrow Mt)^{-1},$$

which is a variant form of the resolvent of M . It can be used as a generating function for the number of random walks through the diagram of M , but a more interesting generating function results from counting loops rather than random walks. Thus consider the series $z(t)$ defined by

$$z(t) = \sum_{k=1}^{\infty} \frac{1}{k} Tr\{M^k\} t^k.$$

whose coefficients count loops. Since the trace is a linear operator, we can apply it after summing the matrix series, which this time is seen to correspond to a logarithm—the indefinite integral of the series for $m(t)$. Thus we have

$$z(t) = Tr\{\Leftrightarrow \ln(I \Leftrightarrow Mt)\}.$$

Since the logarithm is an awkward function to compute for matrices, it is fortunate that there is an identity relating exponentials of traces to determinants of exponentials, namely

$$\exp(\text{Tr}\{X\}) = \text{Det}\{\exp(X)\},$$

which can be used to define the *zeta function*,

$$\zeta(t) = \exp z(t) = \text{Det}\{m(t)\}.$$

Because the loop counts are embedded in an exponential they must be evaluated by forming logarithmic derivatives of the zeta function at the origin, but this is a small price to pay for working with such a convenient quantity as the determinant of the “resolvent” (note that t is not in its accustomed place).

Certain properties of the zeta function can be deduced from its representation as a determinant. Although the determinant of a product is a product of determinants, the involvement of the variable t in the resolvent precludes resolvents having the same property. However a matrix will sometimes have a block diagonal structure allowing the determinant of the matrix to be related to the determinants of the blocks; an example would be the connection matrix for two independent diagrams.

10.2 Traces, $\chi(t)$, $\zeta(t)$

Although the zeta function was not always given that name, there is a classical relationship between it, the characteristic polynomial, and the traces of the powers of a matrix, which are simply restatements of certain properties of polynomials. By convention, the characteristic polynomial of the matrix M is defined to be the determinant

$$\chi(t) = |M \Leftrightarrow It|$$

which is equivalent to

$$\chi(t) = \prod_{i=1}^n (\lambda_i \Leftrightarrow t),$$

for eigenvalues λ_i . Written as a polynomial in powers of t ,

$$\chi(t) = \sum_{i=0}^n (\Leftrightarrow 1)^i a_i t^{n-i},$$

wherein the coefficients a_i can be expressed either as homogeneous products of the roots λ_i or as sums of principal minors, both according to well known formulas.

For the transformations which we intend to make, it is more convenient to use $1/t$ as a variable or, to avoid confusion, to introduce the function

$$\varphi(t) = |I \Leftrightarrow Mt|$$

whose factored form is

$$\varphi(t) = \prod_{i=1}^n (1 \Leftrightarrow \lambda_i t).$$

As a polynomial it differs from χ only in attaching the coefficients a_i to ascending rather than descending powers of t :

$$\varphi(t) = \sum_{i=0}^n (\Leftrightarrow 1)^i a_i t^i.$$

Both χ and φ are finite polynomials, whose reciprocals could be represented as power series, even treating them formally and ignoring questions of convergence. We are interested in ζ as the reciprocal of φ ,

$$\zeta(t) = \prod_{i=1}^n (1 \Leftrightarrow \lambda_i t)^{-1},$$

which can be represented as a single power series by representing each factor in the denominator as a geometric series, multiplying them term by term, and collecting coefficients.

In that case

$$\zeta(t) = \sum_{i=0}^{\infty} h_i t^i,$$

where h_i is the sum of all possible homogeneous products of λ_i ; when $n = 2$, for instance, $h_2 = \lambda_1^2 + \lambda_1 \lambda_2 + \lambda_2^2$.

Of more interest are the coefficients of $z(t) = \ln \zeta(t)$

$$z(t) = \sum_{i=0}^{\infty} m_i t^i,$$

for which

$$m_i = \sum_{j=1}^n \lambda_j^i;$$

in other words, the traces of the powers of M .

The practical situation is that the coefficients m_i are always fairly easy to obtain, while the quantities that are really desired are either the coefficients a_i or h_i . Comparison of the series involved in each of the definitions yields equations of convolution type relating the coefficients, typically in the form of single determinants. Recovering the eigenvalues from the coefficients is much harder since that is just the problem of finding the roots of a polynomial; all the same, the task is hard to avoid since it is likely that some or all of the eigenvalues will be sought for.

Newton's identities relate the power sums m_i to the coefficients a_i of the characteristic polynomial; in the present context they are obtained by differentiating the formula $\varphi(t) = \exp(\Leftrightarrow z(t))$ to obtain $\varphi'(t) = \Leftrightarrow \varphi(t) z'(t)$. Comparing coefficients yields the convolution (with $a_0 = 1$)

$$i a_i = \sum_{j=1}^i (\Leftrightarrow 1)^{j-1} m_j a_{i-j},$$

which can be solved for either a or m in terms of the other.

For example, the system of equations

$$\begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ \cdot \\ \cdot \end{bmatrix} \begin{bmatrix} 1 & \cdot & \cdot & \cdot & \cdot \\ m_1 & \Leftrightarrow 2 & \cdot & \cdot & \cdot \\ m_2 & \Leftrightarrow m_1 & 3 & \cdot & \cdot \\ m_3 & \Leftrightarrow m_2 & m_1 & \Leftrightarrow 4 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ \cdot \\ \cdot \end{bmatrix}$$

is readily solved for the a 's, to obtain

$$a_i = \frac{1}{i!} \begin{vmatrix} m_1 & 1 & \cdot & \cdot & \dots & \cdot & \cdot \\ m_2 & m_1 & 2 & \cdot & \dots & \cdot & \cdot \\ m_3 & m_2 & m_1 & 3 & \dots & \cdot & \cdot \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ m_{i-1} & m_{i-2} & m_{i-3} & m_{i-4} & \dots & m_1 & i \Leftrightarrow 1 \\ m_i & m_{i-1} & m_{i-2} & m_{i-3} & \dots & m_2 & m_1 \end{vmatrix}.$$

For example, $a_2 = \frac{1}{2}(m_1^2 \Leftrightarrow m_2)$, $a_3 = \frac{1}{6}(m_1^3 + 2m_3 \Leftrightarrow 3m_1m_2)$.

10.3 Infinite de Bruijn matrix

It is rare that we will find a matrix which meets the non-redundancy criterion for *all* of its cycles, so consideration should be given to the contrary case. The situation is nicely illustrated for the de Bruijn matrices, which have cycles of all possible varieties while still retaining a certain amount of systematic order. The shift-register interpretation of a de Bruijn matrix makes it clear that an n -stage matrix will have 2^n distinct cycles; but those which possess translational symmetry will be redundant. For example, a cycle of 5 zeroes will show up only once in the fifth power of a two-stage matrix, whereas the string 11010 will show up five times and be counted as one single loop in evaluating $\frac{1}{5}Tr\{M^5\}$.

There would appear to be two alternatives: either insist that the name “zeta function” refer exclusively to the generating function which counts loops although it might not incorporate the resolvent, or retain the elegance of the resolvent even though it may not always count loops correctly. It would appear that the latter alternative is the one that is generally chosen; moreover that such a definition was implicit in classical treatments of the theory of equations even though such a name was never used explicitly.

Since we know the characteristic equation of a finite de Bruijn matrix, it is a straightforward matter to see that they all have the same zeta function

$$\zeta(t) = (1 \Leftrightarrow 2t)^{-1}$$

from which we obtain

$$z_k = \frac{2^k}{k}$$

as the “number of cycles” of length k , although we know that this is in no way a natural representation of the true state of affairs. At least for de Bruijn matrices,

we can get some idea of the discrepancy; using the probabilistic de Bruijn matrix produces an even more general result.

Beginning with the two stage matrix whose elements incorporate both a probability and the parameter t ($t_{ij} = p_{ij}t$),

$$\begin{bmatrix} t_{00} & t_{01} \\ t_{10} & t_{11} \end{bmatrix},$$

we have

$$\begin{aligned} Tr\{B\} &= t_{00} + t_{11} \\ Tr\{B^2\} &= t_{00}^2 + t_{11}^2 + 2t_{01}t_{10} \\ Tr\{B^3\} &= t_{00}^3 + t_{11}^3 + 3(t_{00} + t_{11})t_{01}t_{10} \\ Tr\{B^4\} &= t_{00}^4 + t_{11}^4 + 2(t_{01}t_{10})^2 + \\ &\quad 4(t_{00}^2 + t_{00}t_{11} + t_{11}^2)t_{01}t_{10}. \end{aligned}$$

On the other hand if we work with the three stage matrix,

$$\begin{bmatrix} t_{000} & t_{001} & \cdot & \cdot \\ \cdot & \cdot & t_{010} & t_{011} \\ t_{100} & t_{101} & \cdot & \cdot \\ \cdot & \cdot & t_{110} & t_{111} \end{bmatrix},$$

we have

$$\begin{aligned} Tr\{B\} &= t_{000} + t_{111} \\ Tr\{B^2\} &= t_{000}^2 + t_{111}^2 + 2t_{010}t_{101} \\ Tr\{B^3\} &= t_{000}^3 + t_{111}^3 + 3t_{001}t_{010}t_{100} + 3t_{011}t_{110}t_{101} \\ Tr\{B^4\} &= t_{000}^4 + t_{111}^4 + 2(t_{010}t_{101})^2 + \\ &\quad 4t_{000}t_{001}t_{010}t_{100} + 4t_{001}t_{011}t_{110}t_{100} + \\ &\quad 4t_{011}t_{111}t_{110}t_{101}. \end{aligned}$$

The first thing to be observed is that the structure of these matrices is quite similar, thanks to the fact that the de Bruijn matrices form a very regular family. In fact the same cycles will reappear in all the trace formulas, the only difference being the length of the overlappable fragment represented in each subscript. For example, $Tr\{B\}$ will always be the sum of two contributions, one for the string 0^* and the other for the string 1^* ; likewise a contribution from these two strings and another from $(01)^*$ will always appear in $Tr\{B^2\}$.

Another observation, although it is just barely evident from the small number of terms that have been written, is that the series can be rearranged to group the terms belonging to one particular cycle into its own logarithmic series. It is clear that the terms t_{00} and t_{11} in the first example can be so arranged. There is an indication that the term $t_{01}t_{10}$ would participate in another, since the factor 2 cancels opportunely wherever it occurs.

Recalling the shift register interpretation of de Bruijn diagrams, we see that every cycle which has no internal translational symmetry will show up k times in

$Tr\{B^k\}$ only to be divided by k in the term $\frac{1}{k}Tr\{B^k\}$. If there is an internal symmetry present, the unit cell will have length ℓ ; its cycle will occur to the power k/ℓ , but will only be ℓ such terms in the trace. Note how $(t_{01}t_{10})^2$ has coefficient 2, not 4, in $Tr\{B^4\}$, or that t_{00}^4 only has coefficient 1.

The series for $z(t)$ belonging to the two-stage de Bruijn matrix could then be rearranged to take the form

$$\begin{aligned} z(t) &= t_{00} + \frac{1}{2}t_{00}^2 + \frac{1}{3}t_{00}^3 + \dots \\ &+ t_{11} + \frac{1}{2}t_{11}^2 + \frac{1}{3}t_{11}^3 + \dots \\ &+ t_{01}t_{10} + \frac{1}{2}(t_{01}t_{10})^2 + \frac{1}{3}t_{00}^3 + \dots \\ &\Leftrightarrow \ln(1 \Leftrightarrow t_{00}t_{01}t_{10}) \Leftrightarrow \\ &\Leftrightarrow \ln(1 \Leftrightarrow t_{01}t_{11}t_{10}) \Leftrightarrow \\ &\Leftrightarrow \dots \end{aligned}$$

In turn this produces a zeta function

$$\begin{aligned} \zeta(t)^{-1} &= (1 \Leftrightarrow t_{00})(1 \Leftrightarrow t_{11})(1 \Leftrightarrow t_{01}t_{10}) \times \\ &\quad (1 \Leftrightarrow t_{00}t_{01}t_{10})(1 \Leftrightarrow t_{01}t_{11}t_{10}) \times \\ &\quad \dots \end{aligned}$$

This is a very beautiful and elegant formula, and also blatantly at odds with the fact that B has only two eigenvalues, its characteristic polynomial can only vanish for two values of t , and that its zeta function should have at most two poles. Clearly it is a formula which does not meet the requirements for a convergent infinite product.

10.4 Cluster expansion

Before feeling unduly discouraged about the discrepancy between the apparent number of poles of the zeta function and the dimension of its matrix, we should examine the zeta function more carefully. Since we know the exact zeta function for the normal de Bruijn matrix B , all of whose non-zero matrix elements have the form $t_{ij} = t$, we would want to examine

$$\begin{aligned} \zeta(t)^{-1} &= (1 \Leftrightarrow t)(1 \Leftrightarrow t)(1 \Leftrightarrow t^2)(1 \Leftrightarrow t^3)(1 \Leftrightarrow t^3) \times \dots \\ &= (1 \Leftrightarrow 2t + t^2)(1 \Leftrightarrow t^2)(1 \Leftrightarrow t^3)(1 \Leftrightarrow t^3) \times \dots \\ &= (1 \Leftrightarrow 2t + 2t^3 \Leftrightarrow t^4)(1 \Leftrightarrow t^3)(1 \Leftrightarrow t^3) \times \dots \\ &= (1 \Leftrightarrow 2t + 3t^4 \Leftrightarrow 3t^6 + 2t^9 \Leftrightarrow t^{10}) \times \dots \\ &= \dots \\ &= (1 \Leftrightarrow 2t). \end{aligned}$$

The limit expression for the zeta function is exactly the one that was expected, although truncating the product representation or otherwise taking it too literally

leads to the appearance of spurious poles. Although the algebra is slightly more complicated for the probabilistic matrix, multiplying out the factors of the zeta function reproduces the characteristic polynomial for the general case just as well as it does for the normal case. In principle it must be possible to reconstruct the factored form of the characteristic polynomial, but to do so directly is little more than an exercise in going around in circles. Instead we should utilize the fact that the zeta function has been described in terms of cycles of matrix elements.

The difference between primitive and composite cycles is an important one. A cycle is primitive when its translational symmetry is trivial. Composites are aggregates of primitive cycles, not necessarily connected to one another. Thus, for a two stage diagram $(t_{00})(t_{01}t_{10})$, joins into one long cycle with repeated nodes; but in a three stage diagram the analogous $(t_{000})(t_{010}t_{101})$ forms two disconnected cycles. Because their primitive cycles are selfcontained and independent, composites are commutative, but they never contain repeated factors.

Let c stand for a product of matrix elements forming a cycle, such as $t_{00}t_{01}t_{10}$ and let C be a composite cycle, such as $(t_{00})(t_{01}t_{10})$. Let $|C|$ denote the total length of the composite, but let 1 be the composite cycle of length 0. Finally, let $\mu(C)$ be the parity of the composite cycle, which is required to expand the product of differences.

If a three-valued definition of parity is used, making it vanish for composites with repeated factors; and taking it to be unity for the null composite, μ is the classical Möbius function redefined for cycles:

$$\mu(C) = \begin{cases} \Leftrightarrow 1 & \text{odd parity} \\ 0 & \text{repeated factor} \\ 1 & \text{even parity} \end{cases} .$$

Then, according to the distributive law, the product

$$\zeta(t) = \prod_{\text{cycles}} (1 \Leftrightarrow c)$$

can be replaced by a sum

$$\zeta(t) = \sum_{\text{composites}} \mu(C)C.$$

If the composites are grouped by length, the sum presents once again the form of a polynomial:

$$\zeta(t) = \sum_{i=0}^{\infty} \sum_{|C|=i} \mu(C)C.$$

If the dimension of B were finite, we would expect all the terms of high degree in this sum to vanish. For example, the terms with $|C| = 3$ belonging to the two-stage de Bruijn matrix are

$$(t_{00})(t_{01}t_{10}) + (t_{11})(t_{01}t_{10}) \Leftrightarrow (t_{00}t_{01}t_{10}) \Leftrightarrow (t_{11}t_{01}t_{10}),$$

which evidently vanishes. The corresponding term for the three-stage matrix is

$$(t_{000})(t_{010}t_{101}) + (t_{111})(t_{010}t_{101}) \Leftrightarrow (t_{001}t_{010}t_{100}) \Leftrightarrow (t_{110}t_{101}t_{110}),$$

which need not.

Thus we have a scheme which is capable of displaying the structure of the characteristic polynomial of a matrix explicitly in terms of loops formed from the elements of the matrix. Nevertheless, the same results would also have followed from a careful analysis of the characteristic determinant, although the assignment of parity to the terms is phrased in a somewhat different form.

Using subscripts to identify the cycles by their lexicographically ordered repeating unit, the zeta function for binary de Bruijn matrices takes a form independent of the number of stages, namely

$$\begin{aligned} \zeta(t)^{-1} &= 1 \Leftrightarrow (c_0 + c_1) \Leftrightarrow (c_{01} \Leftrightarrow c_{01}c_1) \Leftrightarrow \\ &\Leftrightarrow (c_{001} + c_{011} \Leftrightarrow c_0c_{01} \Leftrightarrow c_{01}c_1) \Leftrightarrow \\ &\Leftrightarrow (c_{0001} + c_{0011} + c_{0111} \Leftrightarrow c_0c_{001} \Leftrightarrow c_{011}c_1 \Leftrightarrow c_0c_{01}c_1) \Leftrightarrow \\ &\Leftrightarrow \dots \end{aligned}$$

When it is no longer possible to find long cycles which are not composites of shorter cycles, it may be concluded that the series represents a matrix of a finite number of stages; otherwise there is no intrinsic limitation on the series.

The parameters upon which this series depends are no longer the probabilities of individual links in the diagram, but rather the collective probabilities belonging to the primitive cycles. There is also a tendency to disguise the contribution of the block determinants to the zeta function, but this simply means that we have a complementary approach to its interpretation.

10.5 Reduced evolution matrix

The zeta function approach is useful for determining the characteristic equation of a matrix, or even of a family of matrices with a similar structure. Its best application lies in contexts where the parameters associated with all cycles beyond a certain length are known to be negligible, or at least thought to be so. Information about eigenvalues and quantities derived from them is available, but eigenvectors are beyond the scope of the zeta function. Since reciprocal eigenvalues determine the poles of the zeta function, it would seem to be best suited for determining the first few of the largest eigenvalues of a family of matrices.

The reduced evolution matrix has a constant column sum, so its largest eigenvalue is a foregone conclusion. The equilibrium eigenvector is an object of interest, but cannot be encountered via the zeta function. The second largest eigenvalue determines the stability of equilibrium, making it a worthwhile quantity to determine.

The non-zero elements e_{xy} of the reduced evolution matrix are those for which there are cells i and j for which $\varphi(ixj) = y$. Diagonal elements must therefore satisfy

$$\varphi(ixj) = x$$

and therefore are given by the number of ways that x is surrounded by acceptable neighbors; but in any event are determined by the still lifes.

Diagonal elements of E^2 are those for which

$$\varphi(ixj) = y, \varphi(kyl) = x,$$

so that if there is a cycle of length two, there will be a corresponding diagonal element. Continuing the analysis, it is clear that there will be diagonal elements in the powers of E for all the periods whose cycle is equal to the block length. The fact that the least pole of the zeta function is known will exercise a certain restraint on the possible values which the cycle probabilities can assume.

It is in the nature of a reduced evolution matrix that the actual cycles will not just be counted, but overcounted to the extent that they can be broken open and embedded in a larger neighborhood whose central segment follows the same cycle; such is the influence of the dangling cells in the formulas above. Nevertheless this discrepancy is expected to diminish as longer blocks are used, or if the reduced evolution matrix for iterated evolution were used.

Let us work out some of the reduced evolution matrices for Wolfram's (2,1) Rule #22: Its symbolic de Bruijn matrix is

$$\begin{bmatrix} 0 & 1 & . & . \\ . & . & 1 & 0 \\ 1 & 0 & . & . \\ . & . & 0 & 0 \end{bmatrix}.$$

giving the two matrices

$$\alpha = \begin{bmatrix} 1 & . & . & . \\ . & . & . & 1 \\ . & 1 & . & . \\ . & . & 1 & 1 \end{bmatrix}, \beta = \begin{bmatrix} . & 1 & . & . \\ . & . & 1 & . \\ 1 & . & . & . \\ . & . & . & . \end{bmatrix}.$$

from whose products the number of counterimages of any given configuration can be determined. These products give directly the row sums of the reduced evolution matrix, but the individual matrix elements have to be worked out on their own account. Evolution matrices describe evolution differently, sorting counterimages by their central strings, rather than by the cells with which they begin and end.

The reduced evolution matrix for 1-blocks is

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix}.$$

The zeta function is

$$\begin{aligned} \zeta(t) &= \frac{1}{(1 \Leftrightarrow 4t)(1+t)} \\ &= 1 + 3t + 17t^2 + \dots \end{aligned}$$

For 2-blocks it is

$$\begin{bmatrix} 1 & 1 & 1 & 4 \\ 1 & 1 & 1 & . \\ 1 & 1 & 1 & . \\ 1 & 1 & 1 & . \end{bmatrix}.$$

with the same zeta function, in which the presence of two zero eigenvalues is not visible

$$\begin{aligned}\zeta(t) &= \frac{1}{(1 \Leftrightarrow 4t)(1+t)} \\ &= 1 + 3t + 17t^2 + 63t^3 + 257t^4 + \dots\end{aligned}$$

For 3-blocks it is

$$\begin{bmatrix} 1 & . & . & 2 & 1 & 1 & 2 & 4 \\ 1 & . & . & . & . & . & 2 & . \\ . & 2 & 1 & . & . & . & . & . \\ . & 2 & 1 & . & 1 & 1 & . & . \\ 1 & . & . & 2 & . & 1 & . & . \\ 1 & . & . & . & . & 1 & . & . \\ . & . & 1 & . & 1 & . & . & . \\ . & . & 1 & . & 1 & . & . & . \end{bmatrix}.$$

Appendix A

Cycles for Rule 22

A.1 Summary

One of the fundamental conclusions of automata theory is that finite automata or those which are effectively finite, such as the cyclic configurations of an infinite automaton, eventually fall into periodic behavior. Generally the only way to discover these periods is by exhaustive enumeration, according to which the evolution of every possible initial configuration is followed out until the first repetition of a previous configuration is observed.

The notation $\pi.\phi$ means that the configuration repeats itself after π generations, but that there are only ϕ different phases, variations of which can occur through rotation or reflection. Accordingly the representative configuration shown is the one with a gap of maximal length, which is always placed at the extreme right. If a shift to the right of δ cells is involved, the notation $\pi.\phi(\delta)$ is used; a reflective symmetry is often associated with $2\pi.\pi$.

On general principles a ring of N cells cannot have a cycle longer than 2^N , but in practice the lengths of the cycles grow at a far slower rate with increasing N . Some of the longest cycles are a consequence of a shift, whose least common multiple with N often results in a far longer period than the time required to simply repeat the pattern.

Likewise, the periods are determined by the lengths of possible loops in a de Bruijn diagram of 2^{2^P} nodes, which limits the maximum length of a primitive period to this value. In practice this is also found to be a generous upper limit.

The results can be summarized in a table of periods versus cycle lengths. The columns correspond to torii of fixed circumference, the rows to fixed periods, the values of both of which are stated in the margins. The first twenty periods are shown below.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34			
1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	.	1	
2	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	.	.	1	
3
4	1	1	2	1	1	1	1	3	3	5	3	3	1	2	1	4	6	12	13	18	13	13	7	4	.	.		
5	1	1	1	1	1	2	1	1	1	1		
6	1	1	1	.	.	.	2	.	.	.	1	5	.	1	.	.	.		
7	.	.	.	1	1	1	1		
8	1	2	1	1	
9	
10	1	1	
11	1	1	1	2	
12	1	2	1	2	1	1	1	.	.	.	2	1	3	.	.	.		
13	
14	1	1	.	.	1	.	1	1		
15	1	.	
16	
17	
18	1	
19	
20	1	1	

There is only a single still life, a single configuration of period two, and none at all of period three.

From the table above and the continuation below it will be seen that there are numerous configurations of periods 4, 12, and 28. They belong to an extensive family which is formed by the expansion of isolated live cells, and which contains many more members besides. Isolated cells, in Rule 22 especially, follow a growth pattern which resembles a binary counter, or a fractal as the image would be classified nowadays. The boundary of a region filled with live cells expands into the gap separating such regions, but periodically the interior of the expanding region dies out.

Thus if a pair of cells are allowed to expand around a ring, a complementary configuration may be created when the expanding frontiers have nearly met but the expanding region has just become vacant. If the symmetry is just right the complement will repeat the performance to recreate the original configuration, setting the stage for a repetitive cycle.

The final fourteen periods are shown in the table below; as before columns correspond to the circumference of the ring, rows to a fixed period.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34		
21	
22	1	
23	
24	1	.	1	1	1	
25
26	1	1	.	1	1	
27
28	1	1	1	3	1	3	.	.	.	3	3	6	
29
30
31
32
33
34

Longer rings might have been attempted were it not for the fact that the computing time required to analyze the ring doubles with every increment in length, leaving $N = 34$ as something of a limit of endurance with present equipment and techniques. A few additional increments could be gained by using faster equipment, and by proving some “gap theorems” which would render unnecessary the analysis of a certain fraction of cases.

Nevertheless, in practice the interval $20 \leq N \leq 34$ provided enough special cases that it was possible to discover several general families of configurations, finally leading to the analysis of periods in terms of de Bruijn diagrams.

A.2 $N = 1$

· zero

A.3 $N = 2$

· · zero
 · ■ 1.1

A.4 $N = 3$

· · · zero

A.5 $N = 4$

· · · · zero
 · · ■ ■ 2.1
 · ■ · ■ 1.1

A.6 $N = 5$

· · · · · zero

A.7 $N = 6$

· · · · · · zero
 · ■ · ■ · ■ 1.1

A.8 $N = 7$

· · · · · · · zero
 · · · ■ · ■ ■ 7.1(3)

A.9 $N = 8$

· · · · · · · · zero
 · · · · · ■ ■ 6.3
 · · · · · ■ · ■ 4.2
 · · ■ ■ · · ■ ■ 2.1
 · ■ · ■ · ■ · ■ 1.1

A.10 $N = 9$

.....	zero
.....■.■	4.4

A.11 $N = 10$

.....	zero
.....■.■	4.4
.....■.■	4.2
.....■.■■	6.3
.■.■.■.■.■	1.1

A.12 $N = 11$

.....	zero
.....■.■	4.4
.....■■■■	5.5
.....■.■■■■	11.1(4)

A.13 $N = 12$

.....	zero
.....■■■■	5.5
.....■.■	4.2
.■■.■■.■■.■■	2.1
.■.■.■.■.■.■	1.1

A.14 $N = 13$

.....	zero
.....■■■■	5.5

A.15 $N = 14$

.....	zero
.....■.■■■■	12.6
.....■.■■	7.1
.■.■.■.■.■.■	1.1

A.16 N = 15

..... zero
■.....■■ 20.4(3)

A.17 N = 16

..... zero
■■ 14.7
■.■ 12.6
■.....■ 12.12
■■.....■■ 6.3
■.■.....■.■ 4.2
■■.....■■.....■■.....■■ 2.1
■.■.■.■.■.■.■.■.■.■ 1.1

A.18 N = 17

..... zero
■.■ 12.12
■■.....■ 26.13
■.■.....■.■ 4.4


A.19 N = 18

..... zero
■.■ 12.12
■.■ 12.6
■■.■.....■■ 18.9
■.■.....■.■ 4.4
■.■.....■.■ 4.4
■.■.....■.■ 4.2
■.■.■.■.■.■.■.■.■.■ 1.1


A.20 N = 19

..... zero
■.■ 12.12
■.....■■■■.■ 28.14
■.■.....■.■ 4.4
■.■.....■.■ 4.4
■.■.....■.■ 4.4


A.21 $N = 20$




A.22 $N = 21$




A.23 $N = 22$




A.24 N = 23




A.25 N = 24




A.26 N = 25




A.29 $N = 28$




A.30 $N = 29$




A.32 $N = 31$




A.33 $N = 32$



A.34 $N = 33$



A.35 N = 34



in the tables. However, the results obtained are not nearly as comprehensive as for individual cells. Even the table shown requires further analysis to deduce its general structure and ensure that it applies equally well for all the cases not determined empirically.

Bibliography

- [1] Milton Abramowitz and Irene A. Stegun (Editors), *Handbook of Mathematical Functions*, U. S. Government Printing Office, Washington D. C., 1964.
- [2] Roy Adler and Leopold Flatto, "Geodesic Flows, Interval Maps, and Symbolic Dynamics," *Bulletin of the American Mathematical Society* **25** 229-334 (1991).
- [3] Jürgen Albert and Karel Culik II, "A Simple Universal Cellular Automaton and its One-Way and Totalistic Version," *Complex Systems* **1** 1-16 (1987).
- [4] S. Amoroso and Y. N. Patt, "Decision Procedures for Surjectivity and Injectivity of Parallel Maps for Tessellation Structures," *Journal of Computer and System Sciences* **6** 448-464 (1972).
- [5] R. C. Backhouse and B. A. Carré, "Regular Algebra Applied to Path-finding Problems," *Journal of the Institute for Mathematics and its Applications* **15** 161-186 (1975).
- [6] Carter Bays, "The game of three dimensional life" (11/20/86) unpublished (available as a supplement to A. K. Dewdney's February 1987 column).
- [7] Carter Bays, "Candidates for the Game of Life in Three Dimensions," *Complex Systems* **1** 373-400 (1987).
- [8] Carter Bays, "Patterns for Simple Cellular Automata in a Universe of Dense-Packed Spheres," *Complex Systems* **1** 853-875 (1987).
- [9] Carter Bays, "Classification of Semitotalistic Cellular Automata in Three Dimensions," *Complex Systems* **1** 373-400 (1987).
- [10] Richard Bellman, *Introduction to Matrix Analysis*, McGraw-Hill Book Company, New York, 1960.
- [11] I. S. Berezin and N. P. Zhidkov, *Computing Methods*, volume 1, Pergamon Press, Oxford, 1965.
- [12] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy, *Winning Ways for your Mathematical Plays*, Academic Press, 1982 (ISBN 0-12-091152-3) vol. 2, chapter 25.

- [13] Abraham Berman and Robert J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Academic Press, New York, 1979 (ISBN 0-12-092250-9).
- [14] R. Bidaux, N. Boccara, and H. Chaté, "Order of the transition versus space dimension in a family of cellular automata," *Physical Review A* **39** 3094-3105 (1989).
- [15] George D. Birkhoff, *Dynamical Systems*, American Mathematical Society, Providence, Rhode Island, 1927.
- [16] David B. Brown, "Competition of cellular automata rules," *Complex Systems* **1** 169-180 (1987).
- [17] Janusz A. Brzozowski, "Derivatives of regular expressions," *Journal of the Association for Computing Machinery* **11** 481-494 (1964).
- [18] David J. Buckingham, "Some facts of life," *Byte* vol. 3, no. 12 (December, 1978), pp. 54-67.
- [19] Arthur W. Burks (editor), *Essays on Cellular Automata*, University of Illinois Press, Urbana, 1970.
- [20] Noam Chomsky, "Three models for the description of language," *IRE Transactions on Information Theory and IT-2* 113-124 (1956).
- [21] Noam Chomsky and George A. Miller, "Finite state languages," *Information and Control* **1** 91-112 (1958).
- [22] Noam Chomsky, "On certain formal properties of grammars," *Information and Control* **1** 137-167 (1958).
- [23] Noam Chomsky, "Formal properties of grammars," in *Handbook of Mathematical Psychology, volume 2*, John Wiley and Sons, New York, 1963, pp. 323-418.
- [24] Gerardo Cisneros and Harold V. McIntosh, *Notas sobre los lenguajes REC y Convert* (Spanish), Universidad Autonoma de Puebla, 1986.
- [25] E. F. Codd, *Cellular Automata*, Academic Press, New York, 1968.
- [26] J. H. Conway, *Regular Algebra and Finite Machines*, Chapman and Hall, Ltd., London, 1971 (ISBN 412-10620-5).
- [27] P.-J. Courtois and P. Semal, "Bounds for the Positive Eigenvectors of Non-negative Matrices and for their Approximations by Decomposition," *Journal of the Association for Computing Machinery* **31** 804-825 (1984).
- [28] Karel Culick II, "On Invertible Cellular Automata," *Complex Systems* **1** 1036-1044 (1987).
- [29] Karel Culick II and Sheng Yu, "Undecidability of CA Classification Schemes," *Complex Systems* **2** 177-190 (1988).

- [30] Karel Culick II, Jan Pachl, and Sheng Yu, "On the Limit Sets of Cellular Automata," *SIAM Journal on Computing* **18** 831-842 (1989).
- [31] Predrag Cvitanovic, "Invariant measurement of strange sets in terms of cycles," *Physical Review Letters* **24** 2729-2732 (1988).
- [32] A. K. Dewdney, "Computer Recreations - Building computers in one dimension sheds light on irreducibly complicated phenomena," *Scientific American*, May 1985, pp. 10-16.
- [33] A. K. Dewdney, "Computer Recreations - The game Life acquires some successors in three dimensions," *Scientific American*, February 1987, pp 8-13.
- [34] A. K. Dewdney, *The Armchair Universe*, W. H. Freeman and Company, New York, 1988. (ISBN 0-7167-1939-8 pbk)
- [35] M. Dresden and D. Wong, "Life Games and Statistical Models," *Proceedings of the National Academy of Sciences (U.S.A.)* **72** 956-960 (1975).
- [36] Michel Dubois-Violette and Alain Rouet, "A Mathematical Classification of the One-Dimensional Deterministic Cellular Automata," *Communications in Mathematical Physics* **112** 627-631 (1987), "Addendum," **118** 529 (1988).
- [37] Manfred Eigen, "Abstract: The Hypercycle: A Principle of Natural Self-Organization," *International Journal of Quantum Chemistry, Quantum Biology Symposium* **5** 219 (1978).
- [38] Samuel Eilenberg, *Automata, Languages, and Machines, Volume A*, Academic Press, New York, 1974. (ISBN 0-12-234001-9 (pt. A.))
- [39] Samuel Eilenberg, *Automata, Languages, and Machines, Volume B*, Academic Press, New York, 1976. (ISBN 0-12-234002-7 (pt. B.))
- [40] William Feller, *An Introduction to Probability Theory and its Applications*, volume 2 (second edition), John Wiley and Sons, New York.
- [41] F. R. Gantmacher, *The Theory of Matrices*, volume 2, Chelsea Publishing Company, New York, 1959.
- [42] Martin Gardner, "Mathematical Games - The fantastic combinations of John Conway's new solitaire game *Life*," *Scientific American*, October 1970, pp. 120-123.
- [43] Martin Gardner, *Wheels, Life, and Other Mathematical Amusements*, W. H. Freeman and Company, New York, 1983. (ISBN 0-7167-1589-9 pbk)
- [44] Solomon W. Golomb, *Shift Register Sequences*, Holden-Day, Inc., San Francisco, 1967.
- [45] Walter Helbig Gottschalk and Gustav Arnold Hedlund, *Topological Dynamics*, American Mathematical Society, Providence, Rhode Island, 1955.

- [46] J. M. Greenberg and S. P. Hastings, "Spatial patterns for discrete models of diffusion in excitable media," *SIAM Journal on Applied Mathematics* **34** 515-523 (1978).
- [47] J. M. Greenberg, B. D. Hassard, and S. P. Hastings, "Pattern formation and periodic structures in systems modelled by reaction-diffusion equations," *Bulletin of the American Mathematical Society* **84** 1296-1327 (1978).
- [48] J. M. Greenberg, C. Greene, and S. Hastings, "A combinatorial problem arising in the study of reaction-diffusion equations," *SIAM Journal of Algebra and Discrete Mathematics* **1** 34-42 (1980).
- [49] P. Guan and Y. He, "Upper bound on the number of cycles in border-decisive cellular automata," *Complex Systems* **1** 181-186 (1987).
- [50] Howard A. Gutowitz and Jonathan D. Victor, "Local structure theory in more than one dimension," *Complex Systems* **1** 57-68 (1987).
- [51] Howard A. Gutowitz, Jonathan D. Victor, and Bruce W. Knight, "Local structure theory for cellular automata," *Physica* **28D** 18-48 (1987).
- [52] Howard A. Gutowitz and Jonathan D. Victor, "Local structure theory: calculation on hexagonal arrays, and interaction of rule and lattice," *Journal of Statistical Physics* **54** 495-514 (1989).
- [53] Fred Hapgood, "Let there be Life," *Omni*, vol. 9, no. 7 (April 1987) pp. 40-46, 116-117.
- [54] Brian Hayes, "Computer Recreations - The cellular automaton offers a model of the world and a world unto itself," *Scientific American*, March 1984, pp. 10-16.
- [55] Tom Head, "One-Dimensional Cellular Automata: Injectivity from Unambiguity," *Complex Systems* **3** 343-348 (1989).
- [56] G. A. Hedlund, "Endomorphisms and automorphisms of the shift dynamical system," *Mathematical Systems Theory* **3** 320-375 (1969).
- [57] J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation*, Addison-Wesley, Reading, Massachusetts, 1979.
- [58] Lyman P. Hurd, "Formal Language Characterizations of Cellular Automaton Limit Sets," *Complex Systems* **1** 69-80 (1987).
- [59] Lyman P. Hurd, "The Non-wandering Set of a CA Map," *Complex Systems* **2** 549-554 (1988).
- [60] Oscar H. Ibarra and Tao Jiang, "On one-way cellular arrays," *SIAM Journal on Computing* **16** 1135-1154 (1987).
- [61] Hiroyuki Ito, "Intriguing properties of global structure in some classes of finite cellular automata," *Physica D* **31** 318-338 (1988).

- [62] Erica Jen, "Scaling of preimages in cellular automata," *Complex Systems* **1** 1045-1062 (1987).
- [63] Erica Jen, "Cylindrical Cellular Automata," *Communications in Mathematical Physics* **118** 569-590 (1988).
- [64] Kunehiko Kaneko, "Attractors, basin structures and information processing in cellular automata" in *Theory and Applications of Cellular Automata* (Stephen Wolfram, ed.), World Scientific, Singapore, 1986. pp. 367-399.
- [65] Youichi Kobuchi and Hidenosuki Nishio, "Some regular state sets in the system of one-dimensional iterative automata," *Information Sciences* **5** 199-216 (1973).
- [66] Zvi Kohavi, *Switching and Finite Automata Theory* (second edition), McGraw-Hill Book Company, New York, 1978 (ISBN 0-07-035310-7).
- [67] A. N. Kolmogorov, *Foundations of the Theory of Probability*, Chelsea Publishing Company, New York, 1950.
- [68] K. B. Krohn and J. L. Rhodes, "Algebraic Theory of Machines," *Mathematical Theory of Automata*, Polytechnic Press, Brooklyn, New York, 1963. pp. 341-384.
- [69] Christopher G. Langton, "Self-reproduction in cellular automata" *Physica D* **10** 135-144 (1984).
- [70] Christopher G. Langton, "Studying artificial life with cellular automata," *Physica D* **22** 120-149 (1986).
- [71] Christopher G. Langton, "Computation at the edge of chaos: phase transitions and emergent computation," *Physica D* **42** 12-47 (1990).
- [72] Stephen Levy, *Hackers: Heroes of the Computer Revolution*, Anchor Press/Doubleday, Garden City, New York, 1984 (ISBN 0-385-19195-2), chapter 7.
- [73] Wentian Li, "Power spectra of regular languages and cellular automata," *Complex Systems* **1** 107-130 (1987).
- [74] Wentian Li, Norman Packard, and Christopher G. Langton, "Transition phenomena in cellular automaton rule space," *Physics D* **45** 77-94 (1990).
- [75] G. G. Lorentz, *Bernstein Polynomials*, University of Toronto Press, Toronto, 1953.
- [76] Barry F. Madore and Wendy L. Freedman, "Computer simulations of the Belousov-Zhabotinsky reaction," *Science* **222** 615-616 (1983).
- [77] O. Martin, A. Odlyzko, and S. Wolfram, "Algebraic aspects of cellular automata," *Communications in Mathematical Physics* **93** 219-258 (1984).
- [78] Marvin L. Minsky, *Computation: Finite and Infinite Machines*, Prentice-Hall, Englewood Cliffs, New Jersey, 1967.

- [79] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics* **5** 115-133 (1943).
- [80] Henryk Minc, *Nonnegative Matrices*, John Wiley and Sons, New York, 1988 (ISBN 0-471-83966-3).
- [81] Edward F. Moore, "Gedanken Experiments on Sequential Machines," in C. E. Shannon and John McCarthy (eds), *Automata Studies*, Princeton University Press, Princeton, New Jersey, 1956.
- [82] Edward F. Moore, "Machine models of self reproduction," *American Mathematical Society Proceedings of Symposia in Applied Mathematics* **14** 17-33 (1962).
- [83] Edward F. Moore, "The firing squad synchronization problem," in Edward F. Moore (ed), *Sequential Machines: Selected Papers*, Addison-Wesley, Reading, Massachusetts, 1964. pp. 213-214.
- [84] Edward F. Moore, "Machine models of self-reproduction," in A. Burks (ed), *Essays on Cellular Automata*, University of Illinois Press, Urbana, 1970.
- [85] Scott Morris, "Games - The game of Life," *Omni*, vol. 7, no. 1 (October, 1984), pp. 188-189.
- [86] Masakazu Nasu, "Local Maps Inducing Surjective Global Maps of One Dimensional Tessellation Automata," *Mathematical Systems Theory* **11** 327-351 (1978).
- [87] Masakazu Nasu, "Indecomposable Local Maps of Tessellation Automata," *Mathematical Systems Theory* **13** 81-93 (1979).
- [88] Masakazu Nasu, "An Interconnection of Local Maps Inducing Onto Global Maps," *Discrete Applied Mathematics* **2** 125-150 (1980).
- [89] Masakazu Nasu, "Uniformly finite-to-one and onto extensions of homomorphisms between strongly connected graphs," *Discrete Mathematics* **39** 171-197 (1982).
- [90] John von Neumann, *Theory of Self-reproducing Automata* (edited and completed by A. W. Burks), University of Illinois Press, 1966.
- [91] Mark D. Niemiec, "Life Algorithms," *Byte* vol. 4, no 1 (January, 1979) pp. 90-97.
- [92] Norman H. Packard and Stephen Wolfram, "Two-Dimensional Cellular Automata," *Journal of Statistical Physics* **38** 901-946 (1985).
- [93] Asher Peres, "Reversible logic and quantum computers," *Physical Review A* **32** 3266-3276 (1985).
- [94] M. Perles, M. O. Rabin, and E. Shamir, "The Theory of Definite Automata," *IEEE Transactions on Electronic Computers* **EC-12** 233-243 (1963).

- [95] Kenneth E. Perry, "Abstract mathematical art," *Byte*, vol. 11, no. 13 (December 1986), pp. 181-192.
- [96] William Poundstone, *The Recursive Universe*, William Morrow and Company, New York, 1985 (ISBN 0-688-03975-8).
- [97] Kendall Preston, Jr., and Michael J. B. Duff, *Modern Cellular Automata*, Plenum Press, New York, 1984 (ISBN 0-306-41737-5).
- [98] Anthony Ralston, "De Bruijn Sequences—A Model Example of the Interaction of Discrete Mathematics and Computer Science," *Mathematics Magazine* **55** 131-143 (1982).
- [99] D. Richardson, "Tessellations with local transformations," *Journal of Computer and System Sciences* **6** 373-388 (1972).
- [100] D. R. Rutherford, *Substitutional Analysis*, Edinburgh, at the University Press (1948).
- [101] Tadakazu Sato and Namio Honda, "Certain Relations between Properties of Maps of Tesselation Automata," *Journal of Computer and System Sciences* **15** 121-145 (1977).
- [102] L. S. Schulman and P. E. Seiden, "Statistical mechanics of a dynamical system based on Conway's game of Life," *Journal of Statistical Physics* **19** 293-314 (1978).
- [103] Eugene Seneta, *Non-Negative Matrices*, John Wiley and Sons, New York, 1973 (ISBN 0-470-77605-6).
- [104] Claude E. Shannon and John McCarthy (editors), *Automata Studies*, Annals of Mathematics Studies #34, Princeton University Press, Princeton, 1956. (ISBN 0-691-07916-1).
- [105] J. A. Shohat and J. D. Tamarkin, *The Problem of Moments*, American Mathematical Society, Providence, Rhode Island, 1943.
- [106] Sven Skyum, "Confusion in the Garden of Eden," *Proceedings of the American Mathematical Society* **50** 332-336 (1975).
- [107] Alvy Ray Smith III, "Simple computation-universal cellular spaces," *Journal of the Association for Computing Machinery* **18** 339-353 (1971).
- [108] Shinji Takesue, "Reversible Cellular Automata and Statistical Mechanics," *Physical Review Letters* **59** 2499-2502 (1987).
- [109] Tommaso Toffoli, "Computation and Construction Universality of Reversible Cellular Automata," *Journal of Computer and System Sciences* **15** 213-231 (1977).
- [110] Tommaso Toffoli and Norman Margolus, *Cellular Automata Machines*, The MIT Press, Cambridge, Massachusetts, 1987. (ISBN 0-262-20060-0).

- [111] Stanislaw Ulam, "On some mathematical problems connected with patterns of growth of figures," in A. Burks (ed.) *Essays on Cellular Automata*, University of Illinois Press, 1970.
- [112] Richard S. Varga, *Matrix Iterative Analysis*, Prentice Hall, Inc, Englewood Cliffs, New Jersey, (1962).
- [113] Robert T. Wainwright (editor), *Lifeline*, a quarterly newsletter with 11 issues published between March 1971 and September 1973.
- [114] A. Waksman, "An optimum solution to the firing squad synchronization problem," *Information and Control* **9** 66-78 (1966).
- [115] Norbert Wiener, *Cybernetics*, John Wiley and Company, New York, (1948).
- [116] W. John Wilbur, David J. Lipman, and Shihab A. Shamma, "On the prediction of local patterns in cellular automata," *Physica* **19D** 397-410 (1986).
- [117] Stephen J. Willson, "Growth rates and fractional dimensions in cellular automata," *Physica* **10D** 69-74 (1984).
- [118] Stephen Wolfram, "Statistical mechanics of cellular automata," *Reviews of Modern Physics* **55** 601-644 (1983).
- [119] Stephen Wolfram, "Universality and complexity in cellular automata," *Physica* **10D** 1-35 (1984).
- [120] Stephen Wolfram, "Computation theory of cellular automata," *Communications in Mathematical Physics* **96** 15-57 (1984).
- [121] Stephen Wolfram (Ed.), *Theory and Applications of Cellular Automata*, World Scientific Press, Singapore, 1986 (ISBN 9971-50-124-4 pbk).
- [122] Takeo Yaku, "The constructibility of a configuration in a cellular automaton," *Journal of Computer and System Sciences* **7** 481-496 (1973).