EXCERPTED FROM

STEPHEN WOLFRAM

A NEW KIND OF SCIENCE

SECTION 9.9

Time and Causal Networks

But I very much doubt that any such obvious symmetry between space and time exists in the fundamental rules for our universe. And instead what I expect is much like we have seen many times before in this book: that even though at the lowest level there is no direct correspondence between space and time, such a correspondence nevertheless emerges when one looks in the appropriate way at larger scales of the kind probed by practical experiments.

As I will discuss in the next several sections, I suspect that for many purposes the history of the universe can in fact be represented by a certain kind of spacetime network. But the way this network is formed in effect treats space and time rather differently. And in particular—just as in a system like a cellular automaton—the network can be built up incrementally by starting with certain initial conditions and then applying appropriate underlying rules over and over again.

Any such rules can in principle be thought of as providing a set of constraints for the spacetime network. But the important point is that there is no need to do a separate search to find networks that satisfy such constraints—for the rules themselves instead immediately define a procedure for building up the necessary network.

Time and Causal Networks

I argued in the last section that the progress of time should be viewed at a fundamental level much like the evolution of a system like a cellular automaton. But one of the features of a cellular automaton is that it is set up to update all of its cells together, as if at each tick of some global clock. Yet just as it seems unreasonable to imagine that the universe consists of a rigid grid of cells in space, so also it seems unreasonable to imagine that there is a global clock which defines the updating of every element in the universe synchronized in time.

But what is the alternative? At first it may seem bizarre, but one possibility that I believe is ultimately not too far from correct is that the universe might work not like a cellular automaton in which all cells get updated at once, but instead like a mobile automaton or Turing machine, in which just a single cell gets updated at each step.

As discussed in Chapter 3—and illustrated in the picture on the right—a mobile automaton has just a single active cell which moves around from one step to the next. And because this active cell is the only one that ever gets updated, there is never any issue about synchronizing behavior of different elements at a given step.

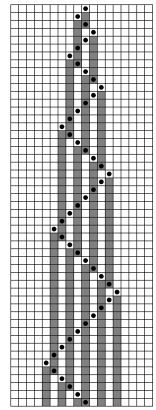

A mobile automaton in which only the single active cell indicated by a dot is updated at each step, thereby avoiding the issue of global synchronization.

Yet at first it might seem absurd to think that our universe could work like a mobile automaton. For certainly we do not notice any kind of active cell visiting different places in the universe in sequence. And indeed, to the contrary, our perception is that different parts of the universe seem to evolve in parallel and progress through time together.

But it turns out that what one perceives as happening in a system like a mobile automaton can depend greatly on whether one is looking at the system from outside, or whether one is oneself somehow part of the system. For from the outside, one can readily see each individual step in the evolution of a mobile automaton, and one can tell that there is just a single active cell that visits different parts of the system in sequence. But to an observer who is actually part of the mobile automaton, the perception can be quite different.

For in order to recognize that time has passed, or indeed that anything has happened, the state of the observer must somehow change. But if the observer itself just consists of a collection of cells inside a mobile automaton, then no such change can occur except on steps when the active cell in the mobile automaton visits this collection of cells.

And what this means is that between any two successive moments of time as perceived by an observer inside the mobile automaton, there can be a great many steps of underlying mobile automaton evolution.

If an observer could tell what was happening on every step, then it would be easy to recognize the sequential way in which cells are updated. But because an observer who is part of a mobile automaton can in effect only occasionally tell what has happened, then as far as such an observer is concerned, many cells can appear to have been updated in parallel between successive moments of time.

To see in more detail how this works it could be that it would be necessary to make a specific model for the observer. But in fact, it turns out that it is sufficient just to look at the evolution of the mobile automaton not in terms of individual steps, but rather in terms of updating events and the causal relationships between them.

The pictures on the facing page show an example of how this works. Picture (a) is a version of the standard representation that I have used for mobile automaton evolution elsewhere in the book—in which successive lines give the colors of cells on successive steps, and the position of the active cell is indicated at each step by a gray dot. The subsequent pictures on the facing page all ultimately give essentially the same information, but gradually present it to emphasize more a representation in terms of updating events and causal relationships.

Picture (b) is very similar to (a), but shows successive steps of mobile automaton evolution separated, with gray blobs in between indicating "updating events" corresponding to each application of the underlying mobile automaton rule. Picture (b) still has a definite row of cells for each individual step of mobile automaton evolution. But in picture (c) cells not updated on a given step are merged together, yielding vertical stripes of color that extend from one updating event to another.

So what is the significance of these stripes? In essence they serve to carry the information needed to determine what the next updating event will be. And as picture (d) begins to emphasize, one can think of these stripes as indicating what causal relationships or connections exist between updating events.

And this notion then suggests a quite different representation for the whole evolution of the mobile automaton. For rather than having a picture based on successive individual steps of evolution, one can instead form a network of the various causal relationships between updating events, with each updating event being a node in this network, and each stripe being a connection from one node to another.

A sequence of views of the evolution of a mobile automaton, showing how a network of causal relationships between updating events can be created. This network provides a very simple model for spacetime in the universe. Picture (a) is essentially the standard representation of mobile automaton evolution that I have used in this book. Picture (b) includes gray blobs to indicate updating events. Picture (c) merges cells that are not being updated. Picture (d) emphasizes the role of vertical stripes as connections between updating events. Pictures (e) through (g) show how a network can be formed with nodes corresponding to updating events. Pictures (h) and (i) demonstrate that with the particular underlying rule used here, a highly regular network is produced. ▸

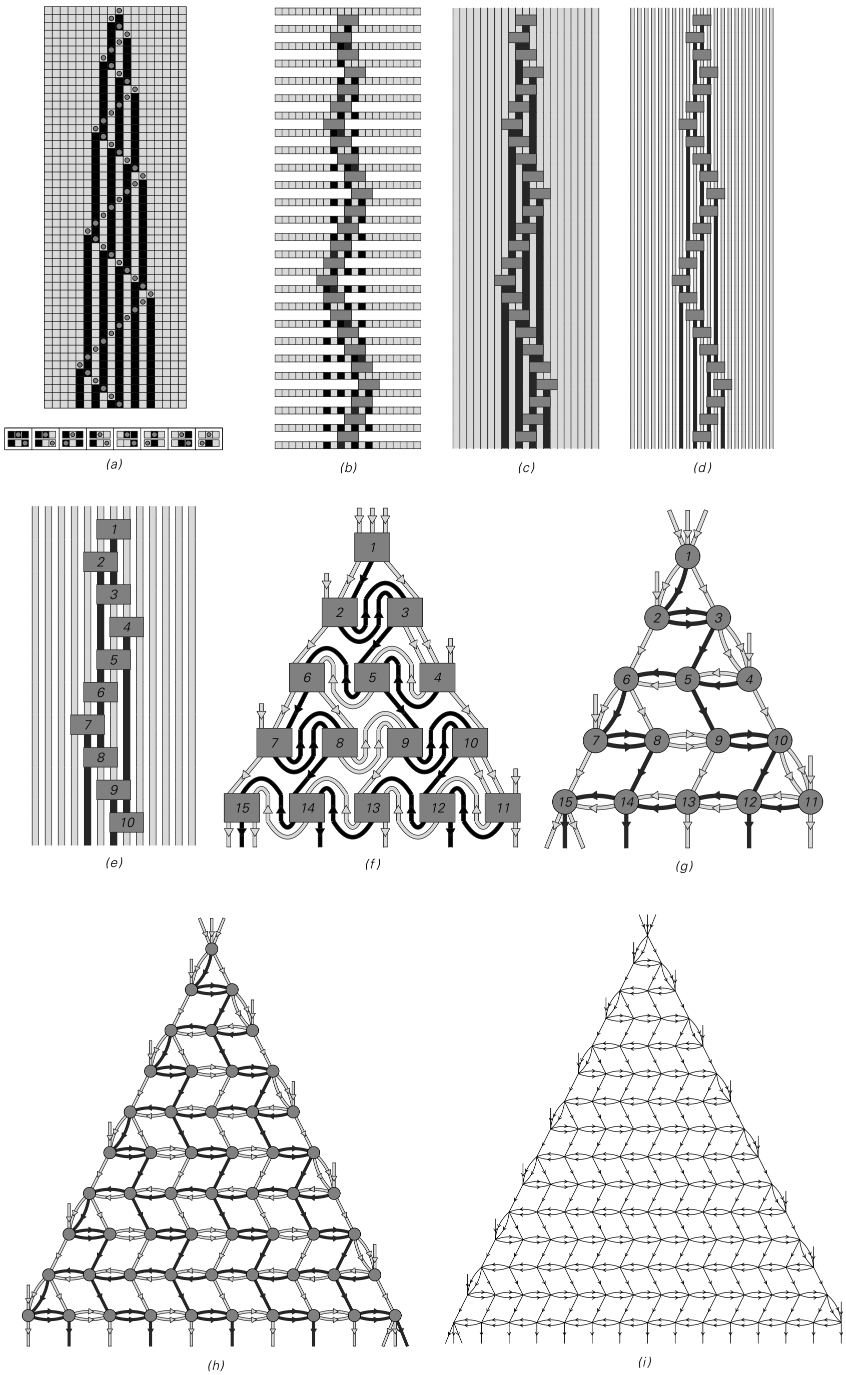
(a)
(b)
(c)
(d)
(e)
(f)
(g)
(h)
(i)

Picture (e) shows the updating events and stripes from the top of picture (d), with the updating events now explicitly numbered. Pictures (f) and (g) then show how one can take the pattern of connectivity from picture (e) and lay out the updating events as nodes so as to produce an orderly network. And for the particular mobile automaton rule used here, the network one gets ends up being highly regular, as illustrated in pictures (h) and (i).

So what is the significance of this network? It turns out that it can be thought of as defining a structure for spacetime as perceived by an observer inside the mobile automaton—in much the same way as the networks we discussed two sections ago could be thought of as defining a structure for space. Each updating event, corresponding to each node in the network, can be imagined to take place at some point in spacetime. And the connections between nodes in the network can then be thought of as defining the pattern of neighbors for points in spacetime.

But unlike in the space networks that we discussed two sections ago, the connections in the causal networks we consider here always go only one way: each connection corresponds to a causal relationship in which one event leads to another, but not the other way around.

This kind of directionality, however, is exactly what is needed if a meaningful notion of time is to emerge. For the progress of time can be defined by saying that only those events that occur later in time than a particular event can be affected by that event.

And indeed the networks in pictures (g) through (i) on the previous page were specifically laid out so that successive rows of nodes going down the page would correspond, at least roughly, to events occurring at successively later times.

As the numbering in pictures (e) through (g) illustrates, there is no direct correspondence between this notion of time and the sequence of updating events that occur in the underlying evolution of the mobile automaton. For the point is that an observer who is part of the mobile automaton will never see all the individual steps in this evolution. The most they will be able to tell is that a certain network of causal relationships exists—and their perception of time must therefore derive purely from the properties of this network.

So does the notion of time that emerges actually have the familiar features of time as we know it? One might think for example that in a network there could be loops that would lead to a deviation from the linear progression of time that we appear to experience. But in fact, with a causal network constructed from an underlying evolution process in the way we have done it here no such loops can ever occur.

So what about traces of the sequential character of evolution in the original mobile automaton? One might imagine that with only a single active cell being updated at each step different parts of the system would inevitably be perceived to progress through time one after another. But what the pictures on page 489 demonstrate is that this need not be the case. Indeed, in the networks shown there all the nodes on each row are in effect connected in parallel to the nodes on the row below. So even though the underlying rules for the mobile automaton involve no global synchronization, it is nevertheless possible for an observer inside the mobile automaton to perceive time as progressing in a synchronized way.

Later in this chapter I will discuss how space works in the context of causal networks—and how ideas of relativity theory emerge. But for now one can just think of networks like those on page 489 as being laid out so that time goes down the page and space goes across. And one can then see that if one follows connections in the network, one is always forced to go progressively down the page, even though one is able to move both backwards and forwards across the page—thus agreeing with our everyday experience of being able to move in more or less any direction in space, but always being forced to move onward in time.

So what happens with other mobile automata?

The pictures on the next two pages show a few examples.

Rules (a) and (b) yield very simple repetitive networks in which there is in effect a notion of time but not of space. The underlying way any mobile automaton works forces time to continue forever. But with rules (a) and (b) only a limited number of points in space can ever be reached.

The other rules shown do not, however, suffer from this problem: in all of them progressively more points are reached in space as time goes on. Rules (c) and (d) yield networks that can be laid out in a quite

(a) (b)

(c)

(d)

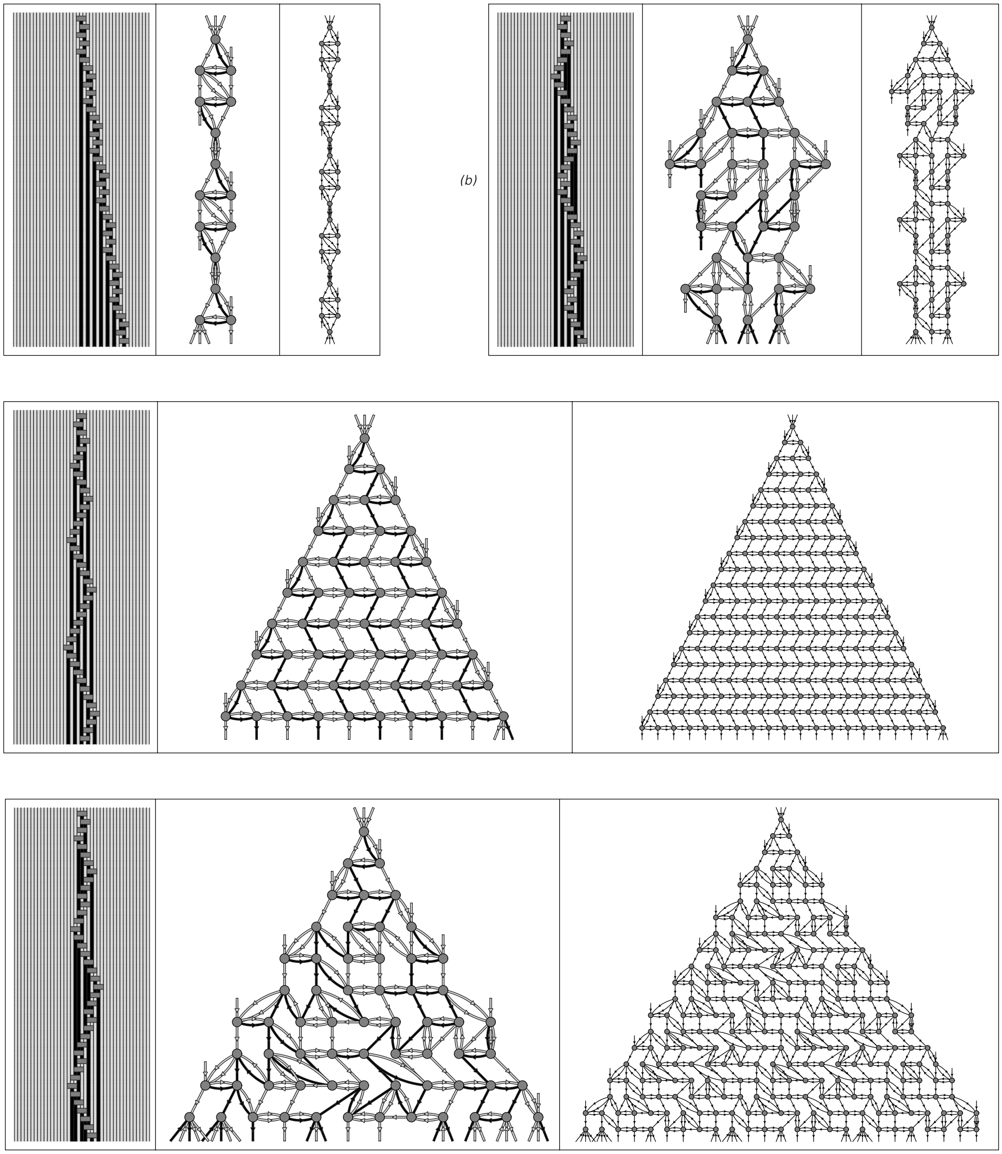

Examples of mobile automata from Chapter 3 and the causal networks they generate. In each case the picture on the left is essentially the standard representation of mobile automaton evolution used in Chapter 3. The pictures on the right are then causal network representations of the same evolution. The networks are laid out in analogy to the space networks on page 479, with nodes being placed on successive rows if they take progressively more connections to reach from the top node.

(a)
(b)
(c)
(d)

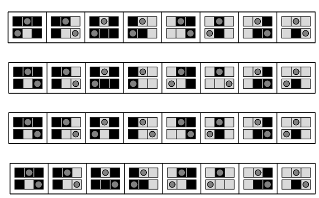

(e)

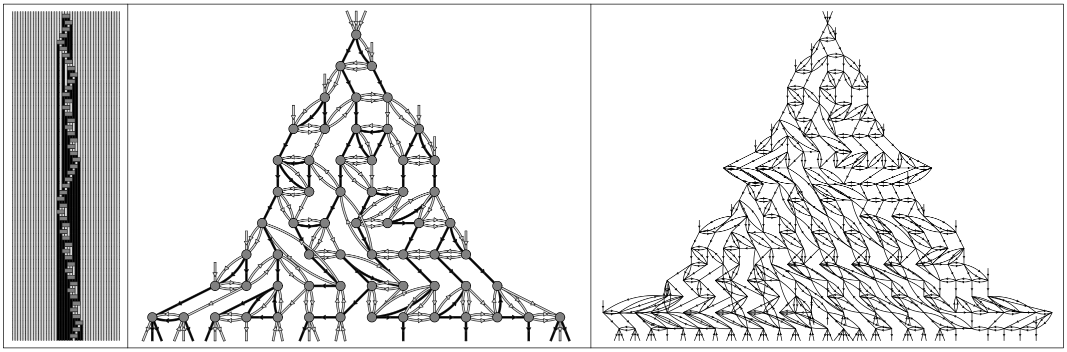

(f)

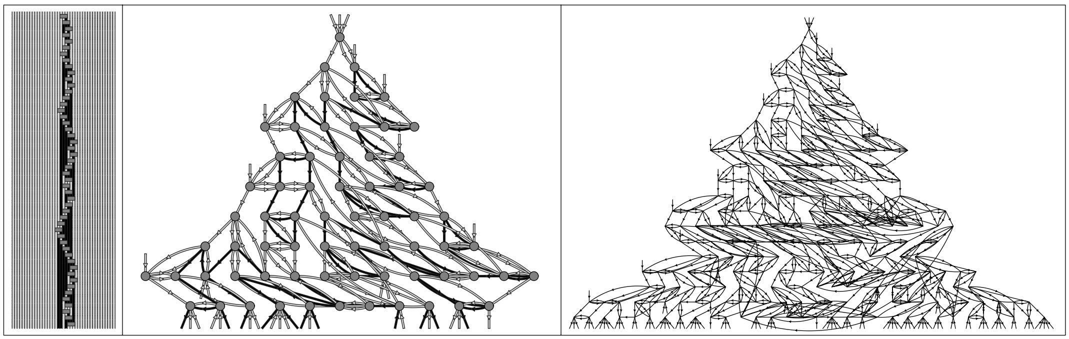

(g)

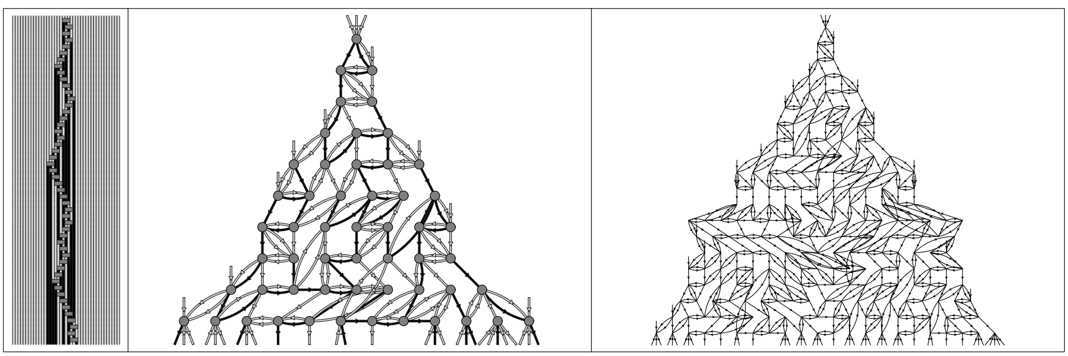

Note that a single connection can join events that occur at very different steps in the evolution of the underlying mobile automaton. And indeed to construct even a small part of the causal network can require an arbitrarily long computation in the underlying mobile automaton. Thus for example to make the causal networks in pictures (e), (f) and (g) requires looking respectively at 2447, 731 and 322 steps of mobile automaton evolution. And indeed in some cases there can be connections that are in effect never resolved. And thus for example in picture (a) there are downward connections that never reach any other node—reflecting the presence of positions on the left in the mobile automata evolution to which the active cell never returns.

(e) (f) (g)

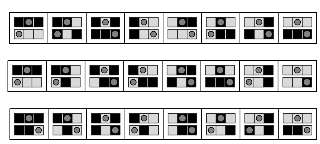

regular manner. But with rules (e), (f) and (g) the networks are more complicated, and begin to seem somewhat random.

The procedure that is used to lay out the networks on the previous two pages is a direct analog of the procedure used for space networks on page 479: the row in which a particular node will be placed is determined by the minimum number of connections that have to be followed in order to reach that node starting from the node at the top.

In cases (a) and (c) the networks obtained in this way have the property that all connections between nodes go either across or down the page. But in every other case shown, at least some connections also go up the page. So what does this mean for our notion of time? As mentioned earlier, there can never be a loop in any causal network that comes from an evolution process. But if one identifies time with position down the page, the presence of connections that go up as well as down the page implies that in some sense time does not always progress in the same direction. Yet at least in the cases shown here there is still a strong average flow down the page—agreeing with our everyday perception that time progresses only in one direction.

Like in so many other systems that we have studied in this book, the randomness that we find in causal networks will inevitably tend to wash out details of how the networks are constructed. And thus, for example, even though the underlying rules for a mobile automaton always treat space and time very differently, the causal networks that emerge nevertheless often exhibit a kind of uniform randomness in which space and time somehow work in many respects the same.

But despite this uniformity at the level of causal networks, the transformation from mobile automaton evolution to causal network is often far from uniform. And for example the pictures at the top of the facing page show the causal networks for rules (e) and (f) from the previous page—but now with each node numbered to specify the step of mobile automaton evolution from which it was derived.

And what we see is that even nodes that are close to the top of the causal network can correspond to events which occur after a large number of steps of mobile automaton evolution. Indeed, to fill in just twenty rows

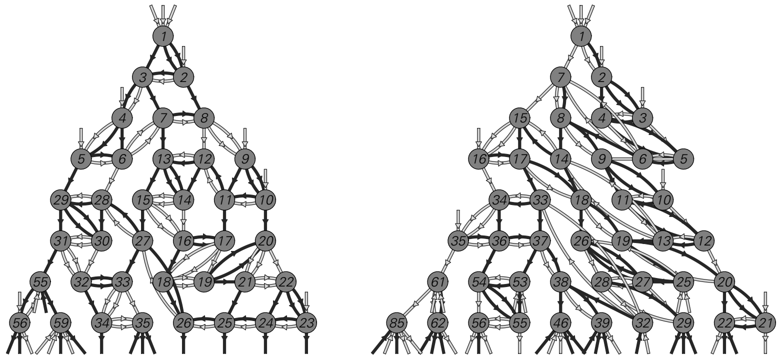

Causal networks corresponding to rules (e) and (f) from page 493, with each node explicitly labelled to specify from which step of mobile automaton evolution it is derived. Even to fill in the first few rows of such causal networks, many steps of underlying mobile automaton evolution must be traced.

of the causal networks for rules (e) and (f) requires following the underlying mobile automaton evolution for 2447 and 731 steps respectively.

One feature of causal networks is that they tell one not only what the consequences of a particular event will be, but also in a sense what its causes were. Thus, for example, if one starts, say, with event 17 in the first causal network above, then to find out that its causes were events 11 and 16 one simply has to trace backwards along the connections which lead to it.

With the specific type of underlying mobile automaton used here, every node has exactly three incoming and three outgoing connections. And at least when there is overall apparent randomness, the networks that one gets by going forwards and backwards from a particular node will look very similar. In most cases there will still be small differences; but the causal network on the right above is specifically constructed to be exactly reversible—much like the cellular automata we discussed near the beginning of this chapter.

Looking at the causal networks we have seen so far, one may wonder to what extent their form depends on the particular properties of the underlying mobile automata that were used to produce them.

For example, one might think that the fact that all the networks we have seen so far grow at most linearly with time must be an inevitable consequence of the one-dimensional character of the mobile automaton rules we have used. But the picture below demonstrates that even with such one-dimensional rules, it is actually possible to get causal networks that grow more rapidly. And in fact in the case shown below there are roughly a factor 1.22 more nodes on each successive row—corresponding to overall approximate exponential growth.

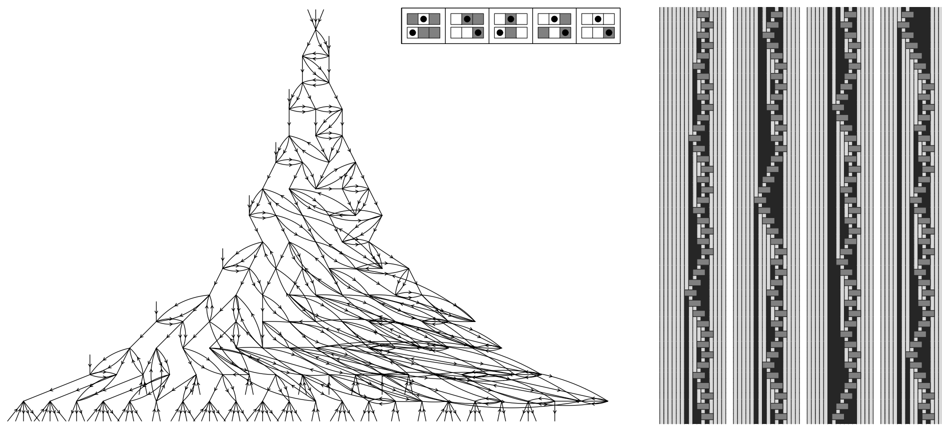

A one-dimensional mobile automaton which yields a causal network that in effect grows exponentially with time. The underlying mobile automaton acts like a binary counter, yielding a pattern whose width grows logarithmically with the number of steps. The three cases not shown in the rule are never used with the initial conditions given here.

The causal network for a system is always in some sense dual to the underlying evolution of the system. And in the case shown here the slow growth of the region visited by the active cell in the underlying evolution is reflected in rapid growth of the corresponding causal network.

As we will see later in this chapter there are in the end some limitations on the kinds of causal networks that one-dimensional mobile automata and systems like them can produce. But with different mobile automaton rules one can still already get tremendous diversity.

And even though when viewed from outside, systems like mobile automata might seem to have almost none of the familiar features of our universe, what we see is that if we as observers are in a sense part of such systems then immediately some major features quite similar to those of our universe can emerge.