BINARY MORPHOLOGY AND CELLULAR AUTOMATA

I can't leave this subject without mentioning *cellular automata* (CAs). Conway's "Game of Life" is an example of a *cellular automaton* (CA). In each generation (or iteration), a set of rules is applied to a binary image to generate another image. Conway used a very simple set of rules, where the value of a pixel on the next iteration depends only on the number of ON pixels adjacent to it. Too few neighboring ON pixels (starvation) or too many (overpopulation) caused the pixel to be OFF (die) in the next iteration. The result was a complicated and unexpected evolution of patterns.

Now I find it interesting that people only seem to use 2-D CAs. Not much happens in 1-D. Using a 3-neighborhood consisting of the binary pixel and its two neighbors in 1-D, you can have $2^3 = 8$ different patterns, each of which can have a rule that sets the new pixel value to either ON or OFF. Thus, you can have $2^8 = 256$ different sets of rules for how each pixel changes in the next iteration. A large number, but certainly manageable. But in 2-D, if your rule depends on the 3-neighbors, you can have $2^9 = 512$ different patterns, leading to 2^{512} different sets of rules. That's a very large set, much too large to be usefully explored at random.

With such large numbers, it is easy to see that people might be impressed with the computational power. And the numbers continue to explode in higher dimensions. For example, in 3-D, with a 3-neighborhood you have 2^{27} patterns, and approximately $2^{130000000}$ different rule sets! I haven't seen anything done in 3-D, though I imagine it would be much more interesting than 2-D, though much harder to show visually what is happening.

For amusement, you can see some interesting 2-D CA at this web site. Each CA is described by its rule, which enumerates the output pixel result for each possible combination of neighbor pixels.

There is a simple relation between 2-D CAs and binary morphology. The *hit-miss* operation is a rule that gives a *single pattern* of ON and OFF pixels to be matched at each location, with the result that if the pattern is matched the output value is ON, and otherwise it is OFF. *Any cellular automaton can thus be built as a generalization of the hit-miss operation, where in general more than one pattern is checked for a match*. So in computational power, the hit-miss operation (thus generalized) is equivalent to the cellular automaton!

Things get even more interesting. Alan Turing's fundamental discovery was that all programmed computers are equivalent to a very simple *Turing Machine* that could read and write binary data on an infinitely long tape.

Any machine that is equivalent to a Turing Machine is called *Turing complete*. And it is easily shown that cellular automata are Turing complete. Consequently, the hit-miss operator can be used to implement a Turing Machine! In principle, you can compute *anything* with this generalized hit-miss operator!

People who have spent a lot of time with CAs tend to become captivated by their power. I believe it is because the 2-D CA have both general computational power and, at the same time, are able to show us the computational evolution directly through our visual system, rather than analytically in some abstract mathematical representation.

The laws of physics can be expressed in many ways; among them, a local description of fields and their derivatives being the most common. The fields, which can be related to physical measurements, are the solutions to *partial differential equations*. These are solved on digital computers, typically by discretizing space onto a lattice and time into discrete increments, as a set of *difference equations*. A very simple example is the Laplace equation for the electrostatic potential in some region surrounded by a closed boundary on which the potential values are known everywhere. The solution on the lattice points inside is found by applying, over and over, a very simple rule: *replace the value of the potential at each point by the average of the four closest neighboring values*.

This is a *relaxation* method; eventually the potential at each point arrives at its final value. You can implement this by a CA; for 2-D geometries, you can even implement it on a spreadsheet! (You may be worried because the world -- and spreadsheet cells -- appear to be described by real numbers instead of binary numbers. We'll sweep this objection under the rug by noting that real numbers can be approximated to arbitrary precision by binary numbers, patterns of 0s and 1s.) Because the laws of nature seem to be expressible locally by such very simple rules, people naturally wonder if the entire universe is, at the root, one big CA with some very simple rules.

There are two areas of physics in particular that have proven to be very difficult to explain, and that have recently forced physicists to develop new fields of mathematics. One is elementary particles, where things are not understood at the very small scales and mathematical descriptions tend to blow up. The other is in the macroscopic regime where we have particle interactions (e.g., in fluids) leading to complex nonlinear behavior such as turbulence and associated chaotic dynamics. In both cases, people naturally search for a set of simple underlying rules to explain the complex behavior. In particle physics, Wheeler speaks of the "quantum foam" at the scale of the *Planck length*, about 10^{-33} cm, where space-time itself is strongly perturbed by quantum gravity. Can rules at that tiny scale, perhaps given by a CA on a lattice, lead to the observed phenomenology of particle physics at much larger scales such as 10^{-17} cm? Likewise, people have observed chaotic behavior arising out of 2-D CAs, along with universal scaling parameters for the phenomena, such as period-doubling in chaos, that arise from very simple dynamical rules for the system evolution. So it is natural to hope that CAs can model and perhaps even explain the most difficult fields of physics. And why stop there? What about conscious intelligence, a mystery so slippery and deep that the mind is completely boggled at its contemplation? Could *we* possibly be CAs?

Edward Fredkin, a physicist at Boston University, has spent much of his career on the search for the rules of the universe based on the assumption that space and time are discrete. A good introduction to his thinking is his 1992 paper, *A new cosmogony*. In 2002, Stephen Wolfram published his Opus Magnus, *A new kind of science*, a 1200 page treatise on cellular automata based on unpublished work he did over the past 10 years. Wolfram, who got his PhD in physics from CalTech at age 20, has a history of brilliant work, including founding the company (of his name) that makes Mathematica. Like Julian Schwinger, his motto could well be "If you can't join 'em, beat 'em," but unlike Schwinger, Wolfram is still in the game. I expect Wolfram to make some important observations in this field.