The Effects of Coarse-Graining on One-Dimensional Cellular Automata

Alec Boyd
UC Davis Physics Deparment
alecboy@gmail.com

Abstract: Measurement devices that we use to examine systems often do not communicate all of the information hidden in the system, by coarse-graining it to give a simpler output. We examine how coarse-graining on 1-D cellular automata affects our understanding of the system by looking at the how the entropy rate of the spatial configuration changes with respect do different coarse-grainings. We find that the entropy rate of most of the cellular automata we simulated does not change significantly when we consider different elements of the simplest coarse-graining rules. However, for rule 18, there were fluctuations in the entropy rate that hinted at structure that is derivable using the strategies outlined by Hanson and Crutchfield [3].

Introduction:

Whenever we observe a system, we are viewing it through a measuring device. Often, systems are far too large for us to be aware of everything, so our measuring devices will reduce the amount of information by providing a "coarse-grained" version of the system to the observer. When we look at a television screen, our eyes average over the individual red, green, and blue pixels, returning an impression of the color, but reducing the amount of information available to the viewer. This is particularly true in statistical physics. Physicists will consider systems that contain huge numbers of particles, each of which have many degrees of freedom, but in their analysis of the system, they will only consider a few measurable bulk quantities, such as temperature, or chemical potential.

How does limiting your information about a system like this change your understanding of the system? What does it take to completely obscure the details of a system, such that it would be impossible to reconstruct a system that reproduces the same behavior? How hard is it to have a measuring device that will allow you to reconstruct the behavior exactly? Is it possible the system that is implied by the data that comes from our measurement device functions differently than the parent system? While I don't answer these questions, they provide the motivation for examining the nature of coarse-graining.

To tackle this problem, we consider one-dimensional cellular automata, because they are simple systems that have been studied and are well understood. To coarse-grain these systems, we take the spatial configuration of the system at a given time in the systems evolution and put it through a transducer, which functions as our measuring device. Just looking at the new grid after putting the system through a transducer doesn't reveal much about the differences between the systems. Thus, for each spatial configuration we calculate the entropy rate. This gives a metric for comparing the time evolution of a coarse-grained system with its parent system.

We consider five different rules for cellular automata. For each rule, we generate a list of one thousand random 1's and 0's as our start state. Then we iterate the rule on that state five hundred times to carry out five hundred time steps. This generates a one thousand by five hundred grid of ones and zeros, which has one spatial axis, and one time axis. Then, for each time step, we take the spatial pattern, and we coarse-grain it by putting every other pair of bits through a function that maps to a zero or a one.

Of the rules 18, 22, 54, 90, and 110, we find that rules 22, 54, and 90 don't yield significant changes in the entropy rate for the coarse-grainings that we use. For rule 110, we find that the entropy rate decreases with time more slowly for the course-grained system than for the parents system. For rule 18, we find that there is a variety of behavior for the different coarse-grainings that we consider. One of the course-grained systems of rule 18 hints at the propagating particles that exists between different domains in the cellular automata, as illustrated by Hanson and Crutchfield [3].

Background:

A one-dimensional cellular automata is a system in which the current state is expressed by a string of ones and zeros. Then, you deterministically operate on your current state to get the next state. We consider the subset of cellular automata in which the rules are based on nearest neighbors. This means that if you are trying to find the next state of a location in your string, you only use your current state, and that of your nearest neighbors. There are only 256 rules for how to do this. Figure 1 and 2 illustrates how this is done with rule 18.

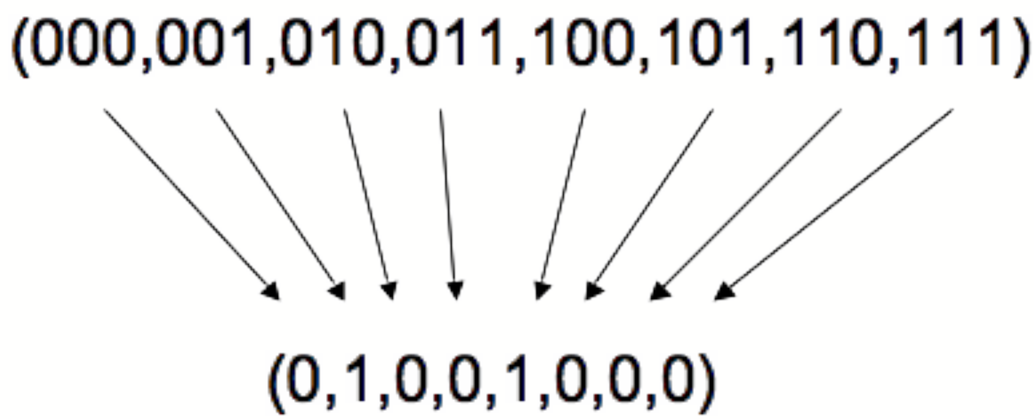

Figure 1: Each triplet maps to a number which is defined by the rule.

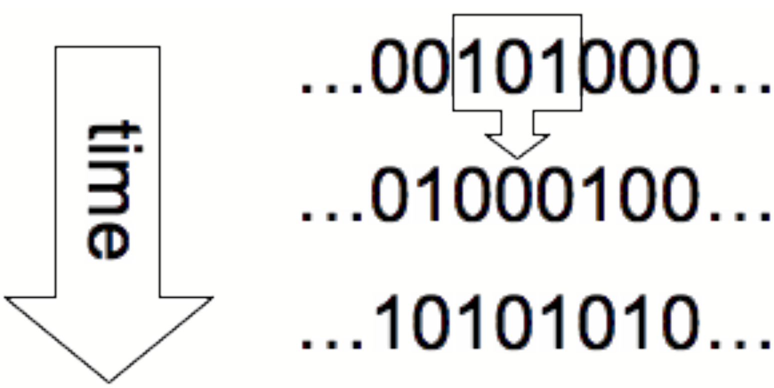

Figure 2.

In this case, we start with a list of random 1's and 0's, and for rule 54 the result looks like Figure 3., where the oranges cells represent 1's and the white cells represent 0's.

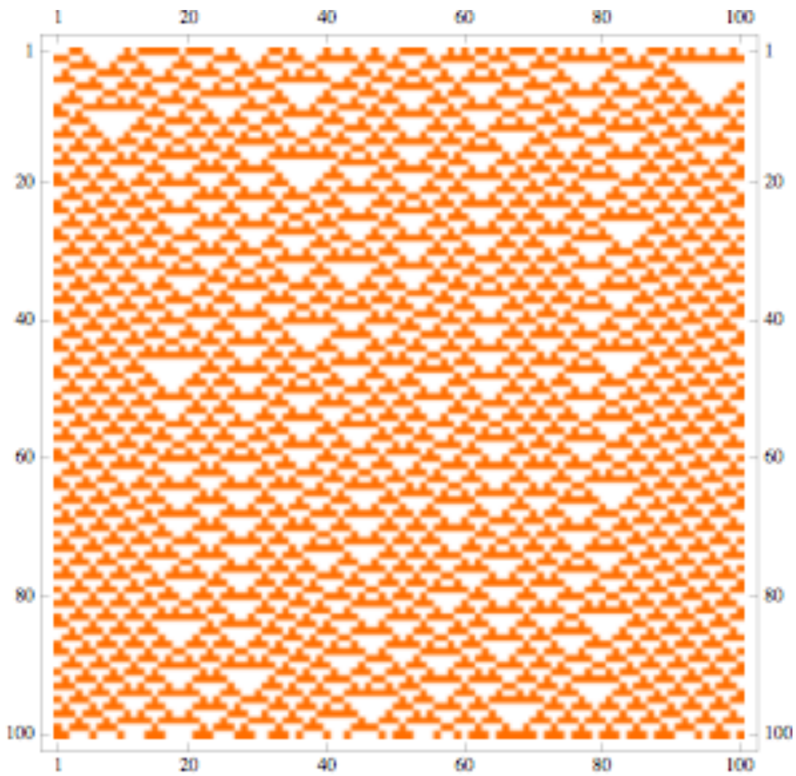
Figure 3.

Methods:

For each of the rules 18, 22, 54, 90, and 110, we start with a different randomized initial state of one thousand 0's and 1's. Then, we iterate five hundred times to get a 1000x500 matrix. We set up the boundaries of our list states to be circular, which means that the 0^{th} cell is the right nearest neighbor of the 999^{th} cell.

For each of the constructed space-time diagrams, we coarse-grain it in five different ways. We coarse-grain a diagram in a similar way to determining the next element in a cellular automata. In every spatial configuration, for every other pair of cells we map those two elements to a one or a zero, which effectively halves the spatial extent of our system. There are 16 possible rules to do this type of coarse-graining. Figure 4 shows how rule 4 works, and Figure 5 shows how it is implemented.

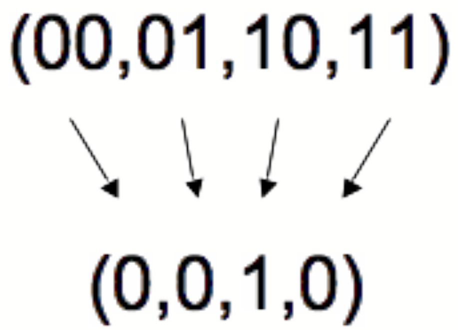

Figure 4: Each possible pair maps to a number that's defined by rule 4.

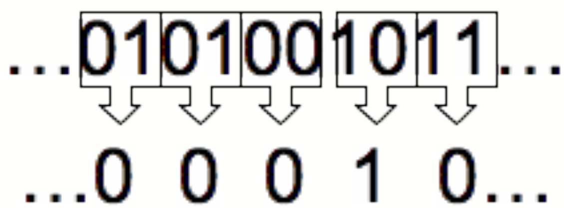

Figure 5: (Top) The parent spatial configuration. (Bottom) The coarse-grained spatial configuration for rule 4.

When we apply this process to Figure 3., yielding Figure 6.

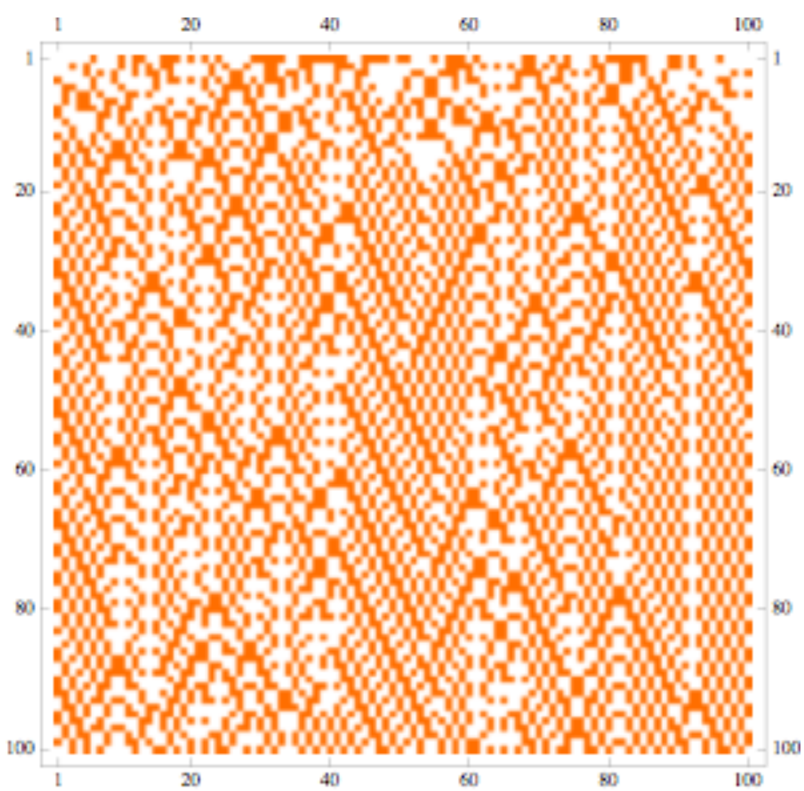

Figure 6.

The five coarse-graining rules that we use are 0001, 0010, 0011, 0110, and 0111. Rule 1111 and 0000 are uninteresting, because they map everything to a single value. The rules are 1-0 symmetry, because it doesn't matter if we exchange all 1's with all 0's for evaluating the Shannon entropy. Lastly, in the computation, I assumed left-right symmetry, which is not appropriate for rule 110. For future work on this, I will extend the number of rules so that I am not assuming left-right symmetry.

For each of the grids that we construct, we evaluate the approximate entropy rate for the spatial configuration. The entropy rate is defined as

$$h_u = \lim_{L \to \infty} H(L) - H(L-1) \quad \text{(eq. 1)}$$

where H(L) is the Shannon entropy of words of length L. This is calculated by evaluating the probability of all words of length L, and plugging into the following equation.

$$H(L) = -\sum_{w \ni words} \Pr(w) \log_2(\Pr(w)) \quad \text{(eq. 2)}$$

H(L) represents the uncertainty in the which word you expect to read at length L. If you have been reading the spatial configuration from right to left for a long time (or left to right, it doesn't matter), the entropy rate tells you the uncertainty you have about the next number you read from the spatial configuration. This is an interesting metric, because it is conceptually similar to thermodynamic entropy, which grows as the log of the available state space of the system. If a system can only be measured as a single element of the state space, the uncertainty in the system grows the same as the log of the state space. So, the entropy rate can be thought of as the thermodynamic entropy density (entropy-per-cell) of the system.

Because of limited data and time, we can only approximate the entropy rate as H(L)-H(L-1) for some L. Because of the limited size of our lattice, the statistics of for generating our probability distribution won't be very good past a certain L. The highest value of L for which the graph of H(L) was still linear for all cases was L=5. Thus, in this study we approximated entropy rate to be

$$h_u \approx H(5) - H(4) \quad \text{(eq. 3).}$$

We then plotted the entropy rate.

Results:

When we coarse-grained our cellular automata, there was one change that was consistent among all the different cases. When we plotted the entropy rate as a function of time, the variance of the entropy rate would be higher for the coarse-grained system than for the parent system, as illustrated by figure 7.

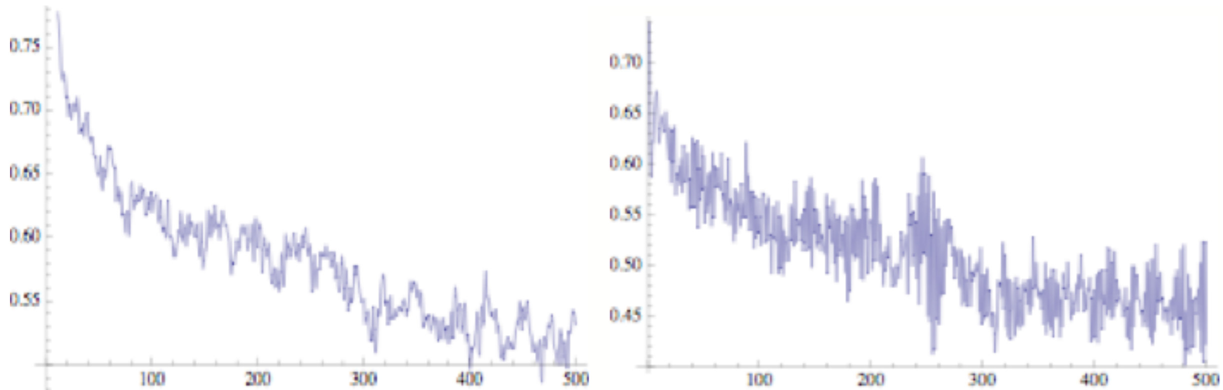

Figure 7: The vertical axis is time entropy rate in both cases, and the horizontal axis is time. (left) rule 54 entropy rate curve. (right) rule 54 coarse-grained with rule (0010).

The increased variance is most likely because of the fact that when we coarse-grain our data, the length of the spatial configuration that we analyze shortens. Thus, the word distribution we use becomes less exact, yielding more noise in the calculation of the evaluation of the entropy.

However, when I compare the relative entropy rates of the systems, I average over ten time steps to eliminate the wrinkles, and make the graphs more readable. Figures 8, 9, and 10 show the results for rules 22, 54, and 90.

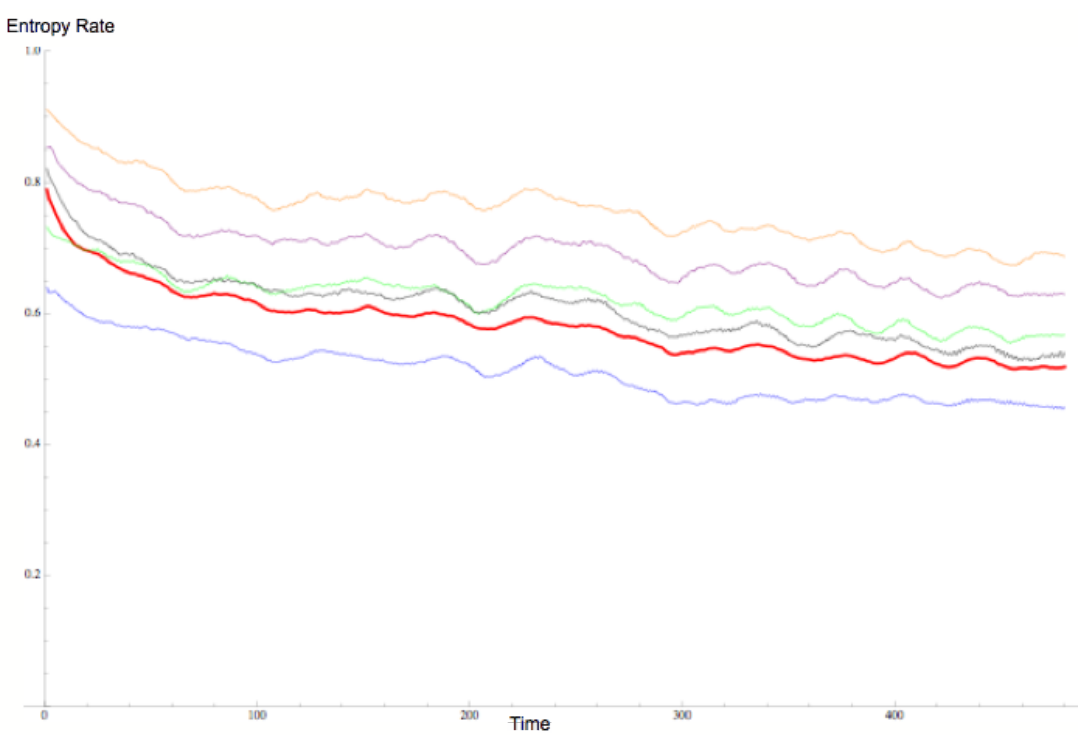

Figure 8: Rule 54. The bold red line is the parent system.

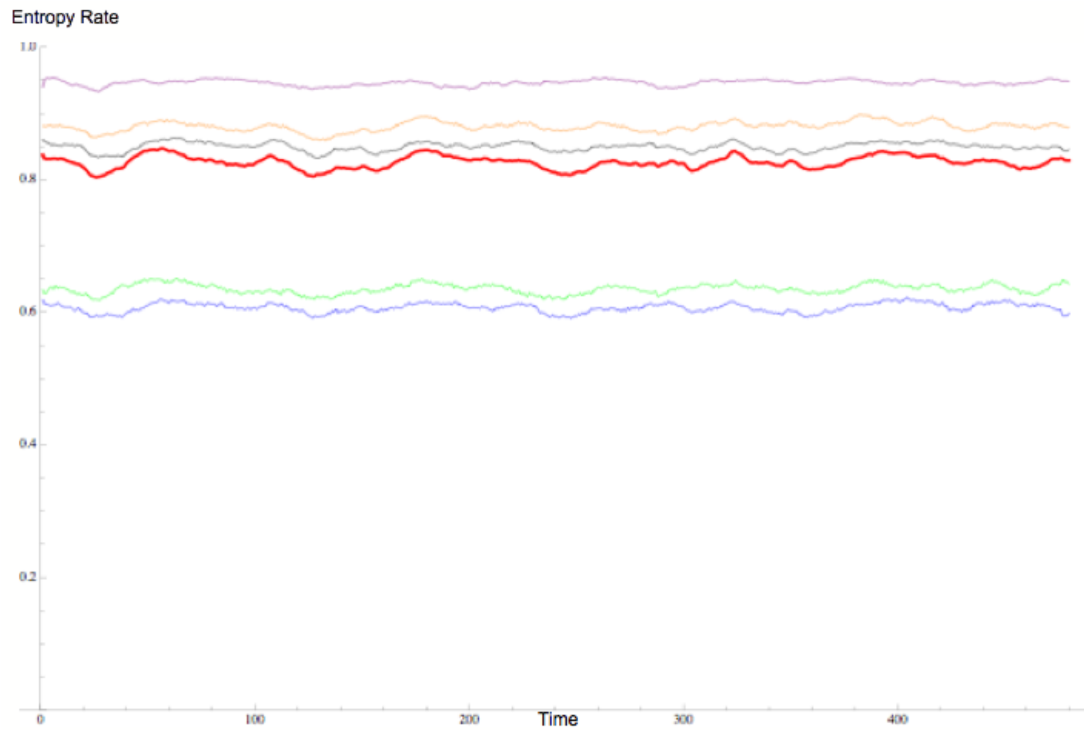

Figure 9: Rule 22.

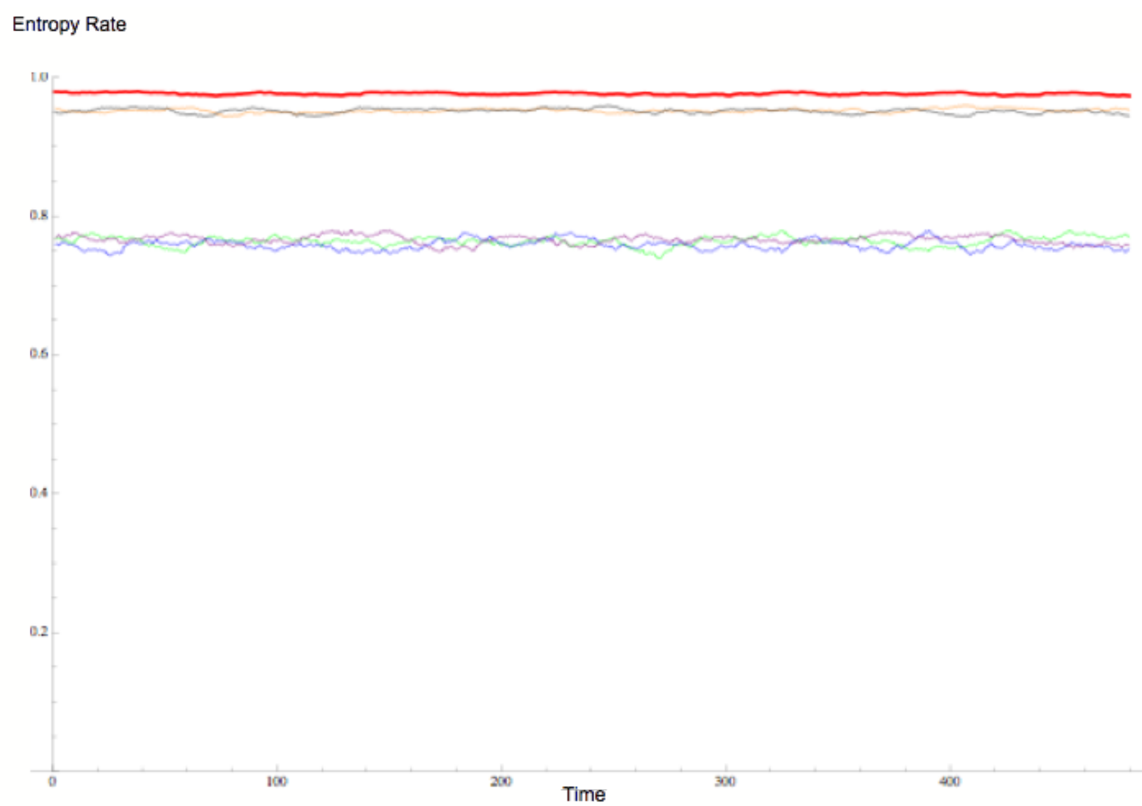

Figure 10: Rule 90

In each case, the bold red line represents the entropy rate of the parent system, and each of the other colored lines represents the entropy rate of one of the five coarse-grainings. For rule 54, the entropy rate of the system appears to be decreasing. For rule 90 and rule 22, the entropy rate doesn't appear to change appreciably over the time scale that we observe. However, in all three cases, the coarse-grainings of each system appear to have relatively the same features as the parents system, with an overall shift in the entropy rate. It seems that the coarse-grainings aren't changing anything that our metrics can easily detect. The behavior of rule 110 is shown in Figure 11.

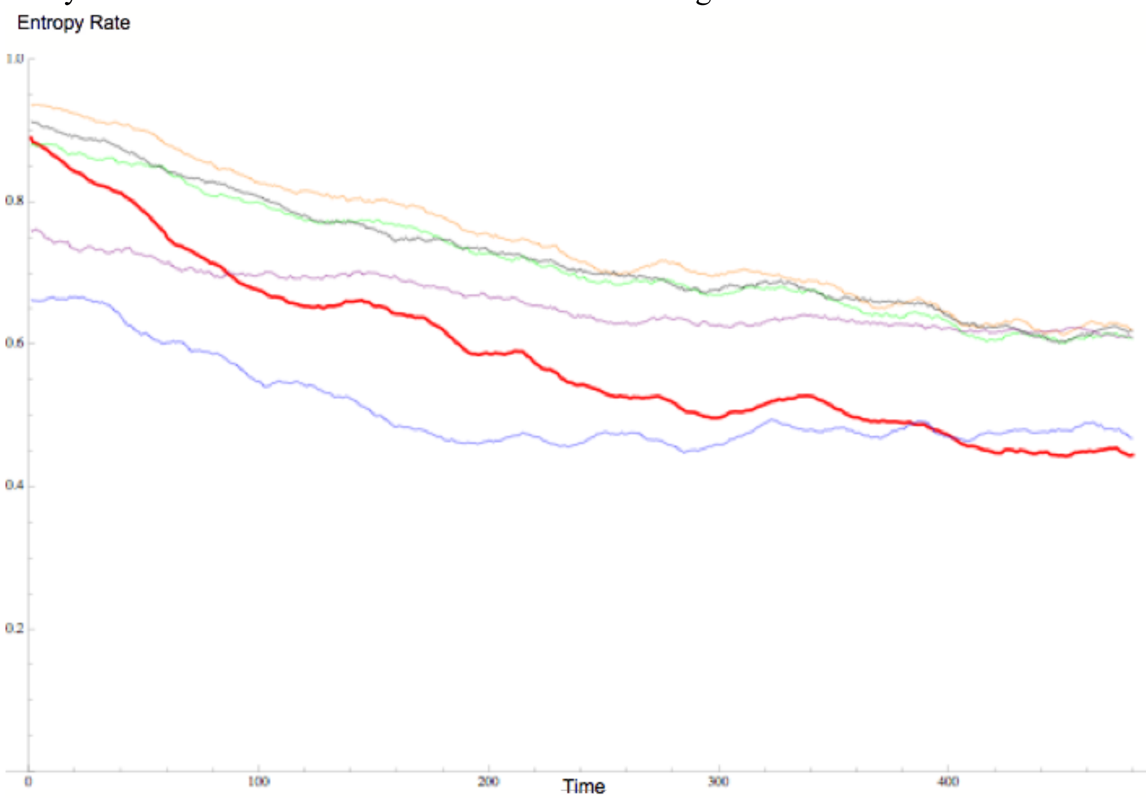

Figure 11: Rule 110.

In this case, the coarse-grainings display behavior which is different from the parent system. The parent system's entropy rate is decreasing faster than it is for the coarse-grainings. This might imply that the systems being described by the coarse-grainings obey different rules than the parent system.

Rule 18 is the most interesting case of all five. The entropy rate for rule 18 is shown in Figure 12.

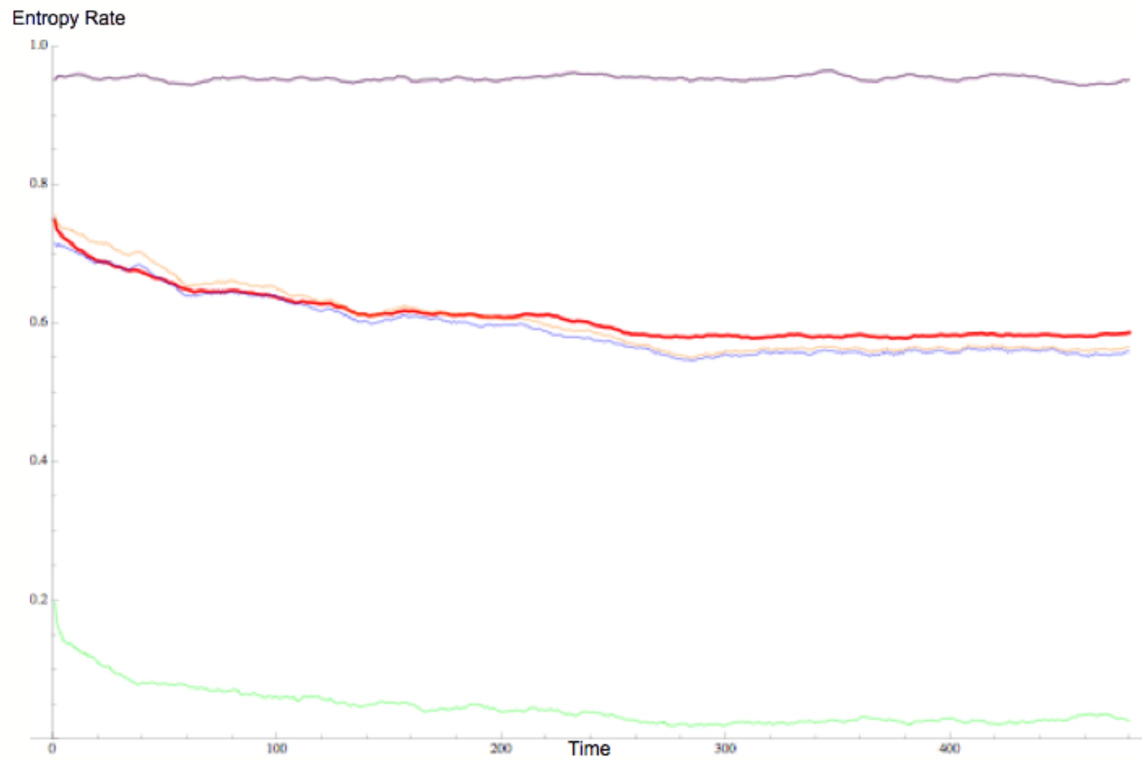

Figure 12: Rule 18.

As you can see from the graph, the coarse-grainings express three different behaviors. Two of the course-grainings appears to be almost completely random, with the entropy rate being close to one. The another coarse-graining appears to have almost no uncertainty, because the entropy rate appears to drop to zero. Then, there are two coarse-grainings (0010 and 0011) which look almost exactly like the parent system. These two are actually very interesting, because, it turns out that they have a very large variance.

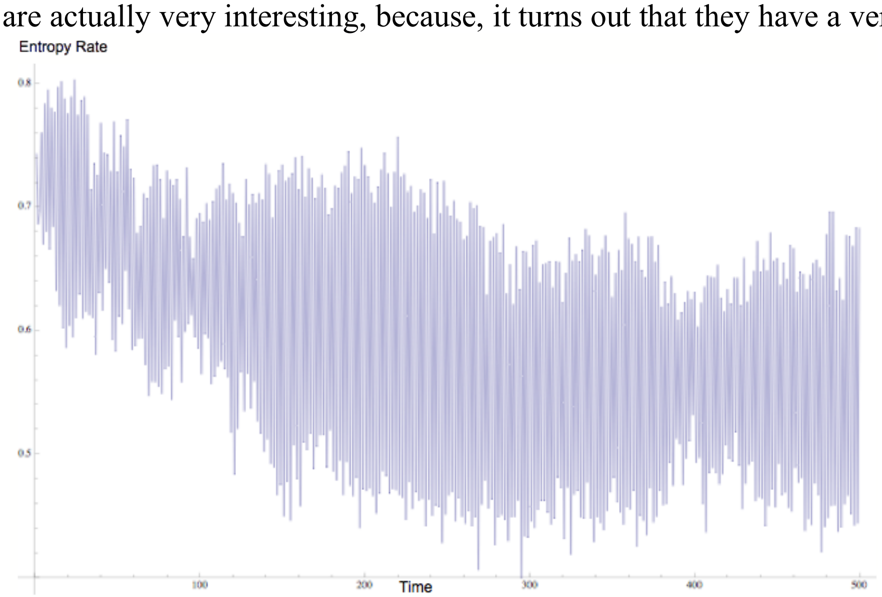

Figure 13: Rule 18 coarse-graining 011.

It appears that this variance has period 2, so I compare the system that exists at even time-steps to the system at odd time-steps. Figure 14 shows the space-time diagram of a rule 18 cellular automata after being coarse-grained with rule 0011. The cells that correspond to even time-steps are labeled by blue, and the cells that correspond to odd time-steps are labeled by orange.

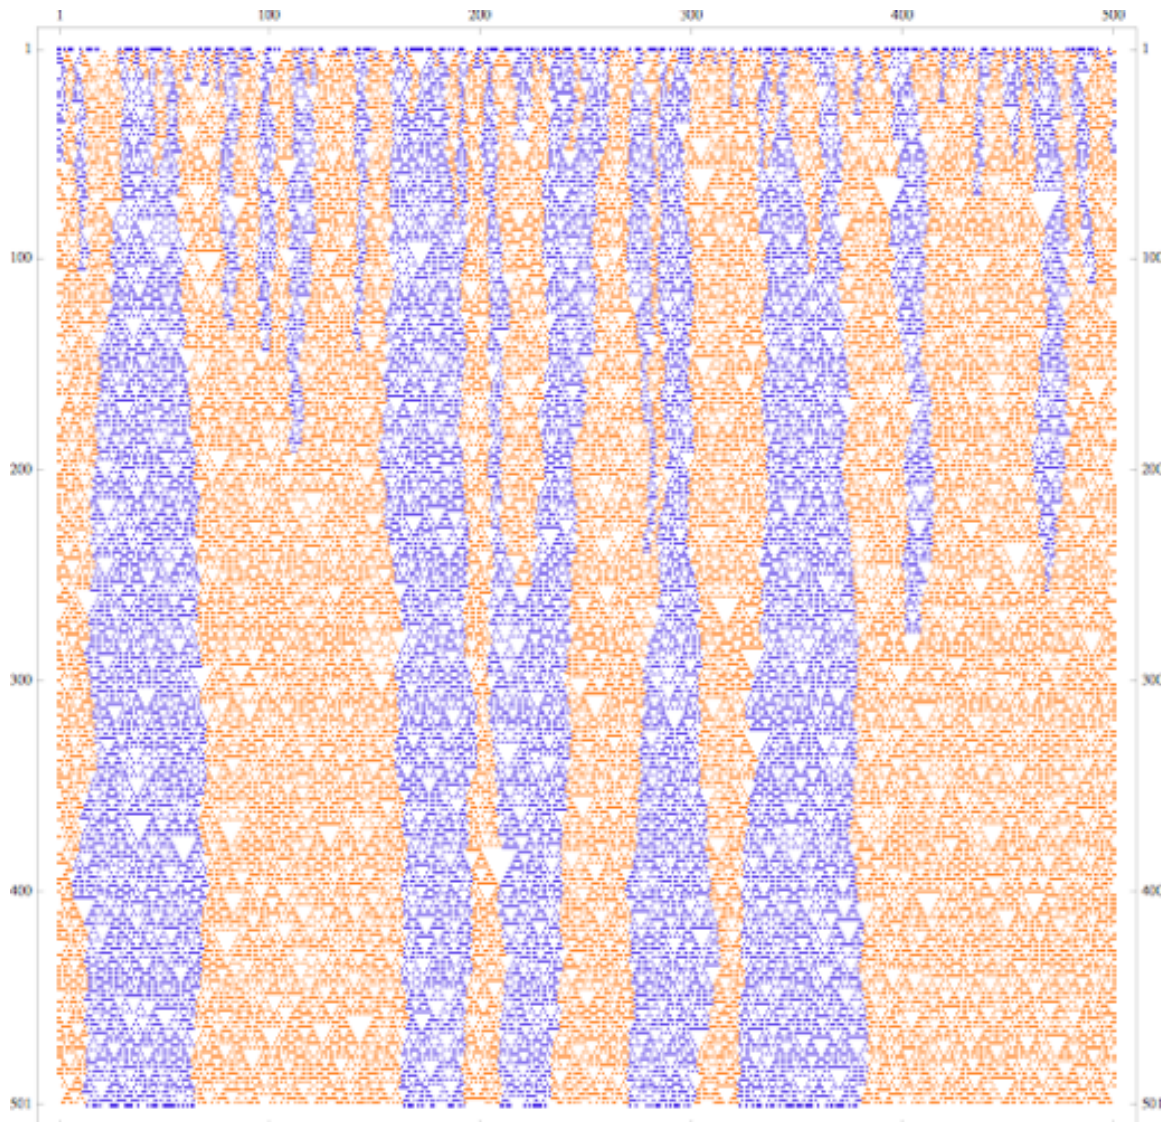

Figure 14.

It looks like we have two distinct regions that are separated by a boundary, which have the potential to merge with other boundaries as you increase time. This is very similar to the results which are described for rule 18 by Hanson and Crutchfield [3]. They found that spatial configurations of rule 18 can be divided into domains, which are invariant under the operation of rule 18. This domain is given by the string …0A0A0A… which is alternating 0's and A's, where A represents a choice between 0 and 1. This domain can be expressed with on the 0's on either odd cells or the even cells corresponding to two different phases. In places where domains of different phases meet, we have a particle to account for the phase difference. These particles propagate forward in time, and can merge, but cannot spontaneously generate. This description matches the picture above, where we can see the orange and blue regions as the odd or even domains of the cellular automata.

I should reiterate that in the even spatial configurations, we have two domains: one of blue and white cells, and one of entirely white cells. We have a similar situation for the odd spatial configurations, except that instead of blue we have orange. To understand why this is, recall that the coarse-graining rules that yield this behavior are 0011 and 0010. This means that the pairs 00 and 01 both get mapped to 0. Because this course-graining rule takes every other pair of cells, it has an inherent phase, and therefore treats the strings 0A0A0A… and A0A0A0… differently. Both 0011 and 0010 will map 0A to 0 and A0 to A. Thus when given 0A0A0A…, the coarse-graining will be 000…, and when given A0A0A0…, the coarse-graining will be AAA…. This implies that the observed regions of all white are really just regions of 0A0A… and the regions of blue (or orange) and white are really just regions of A0A0…. So, the blue and orange regions we observe in Figure 14 are the same as the domains that are apparent in the parent system. The fact that blue domains are continuous with time and occur on the even time-steps, while the orange domains occur on the odd time-steps implies that these domains oscillate between A0A0… and 0A0A…. periodically with time. They just do so out of temporal phase, so they are always out of phase for any given spatial configuration. Thus, by viewing these course-grainings, are better able to visualize the functionality of rule 18, which is propagation particles separating two different types of domains.

Conclusion and Future Work:

For most of the cellular automata that we consider, it is unclear whether coarse-graining on this level yields different behavior. However, coarse-graining rule 18 yielded very different behavior, which was revealing about the system. It is unclear if this is something special about rule 18, or whether this is something special about the rules we are using to coarse-grain the cellular automata. In either case I believe that this subject merits further work to better understand how systems will change when coarse-grained.

It’s possible that the coarse-grainings we used had an anomalous affect on rule 18. One of the greatest weaknesses of our coarse-graining is that by mapping from every other pair of elements to a single element we have introduced measurement device which is sensitive to the length two phase of the diagram. It could be that these rules only interacted in an interesting way with rule 18, because the system had regions that had two-cell periodicity. In order to resolve this problem, we might consider coarse-grainings that map from three or more cells to one. However, if the system has domains with period three or more, we could have the same problems. It would be ideal to find some sort of coarse-graining that doesn’t have any sort of phase choice inherent in its application.

Also, more work should be done to consider this coarse-graining from and information theory perspective. It seems that part of what is interesting about coarse-graining a system like this, is that we have the potential to either raise or lower the entropy rate, making a system appear more or less ordered. In the case of a two-cell coarse-graining, this is based on the fact that if we have two variables X and Y, which have some uncertainty, and we have Z, which is a deterministic function of X and Y, then the uncertainty in Z can be both greater or less than the average of the uncertainty in X and Y. For example, given H(X)=1 and H(Y)=0, if we choose the coarse-graining function Z=Y, then H(Z)=0, which is less than the average of H(X) and H(Y), which is ½. If Z=X, then H(Z)=1, which is greater than the average. In these cellular automata, the average entropy of X and Y is the average of the entropy of two adjacent cells in the parent system. This average is related to the entropy rate of the parent system. H(Z) is conceptually similar to what the entropy rate should be for the coarse-grained system if Z is the value of a cell in that system. It's possible that we could learn more about how the entropy rate changes when you coarse-grain by considering this type of analysis.

References:

[1] J. E. Hanson and J. P. Crutchfield, "Computational Mechanics of Cellular Automata: An Example", Physica D **103** (1997) 169-189

[2] J. P. Crutchfield and J. E. Hanson, "Turbulent Pattern Bases for Cellular Automata", Physica D **69** (1993) 279-301.

[3] J. E. Hanson and J. P. Crutchfield, "*The Attractor-Basin Portrait of a Cellular Automaton", J. Statistical Physics* ***66*** *(1992) 1415 - 1462.*