Tight Bounds on Periodic Cell Configurations in Life

David J. Buckingham and Paul B. Callahan

CONTENTS

This work was partially completed while Callahan was a postdoctoral assistant in the Institute for Theoretical Computer Science, ETH Zürich, Switzerland.

Periodic configurations, or oscillators, occur in many cellular automata. In an oscillator, repeated applications of the automaton rules eventually restore the configuration to its initial state. This paper considers oscillators in Conway's Life; analogous techniques should apply to other rules. Three explicit methods are presented to construct oscillators in Life while guaranteeing certain complexity bounds, leading to the existence of

- an infinite sequence K_n of oscillators of periods $n = 58, 59, 60, \ldots$ and uniformly bounded population, and
- an infinite sequence D_n of oscillators of periods $n = 58, 59, 60, \ldots$ and diameter bounded by $b\sqrt{\log n}$, where b is a uniform constant.

The proofs make use of the first explicit example of a stable glider reflector in Life, solving a longstanding open question about this cellular automaton.

1. INTRODUCTION

Periodic configurations, or oscillators, are a natural characteristic of many cellular automata. In an oscillator, repeated applications of the rules of the cellular automaton (CA) eventually restore the configuration to its initial state either in the same place, or in some cases shifted slightly. In the latter case, the configuration will seem to propagate across cellular space and is often referred to as a *spaceship* [Berlekamp et al. 1982, Chapter 25; Poundstone 1985]. In this paper, "oscillator" will refer to a stationary periodic configuration unless otherwise noted.

We define the *period* of an oscillator to be the smallest number of rule applications needed to restore it to its initial state. It is quite common to see the same small, low-period oscillators arising repeatedly from random initial states. There is

1058-6458/1998 $0.50 per page

no obvious connection between the shapes of these configurations and the underlying rules of the CA; rather, these shapes must be viewed as an emergent property of the system. Random experimentation often turns up a wealth of such configurations. There are usually other oscillators that are unlikely to arise from random initial states. In principle, the set of all oscillators can be enumerated using such techniques as the de Bruijn diagram [McIntosh 1991], or constraint satisfaction [Bell and Hickerson 1994]. When applicable, these techniques often yield remarkable results. However, they rapidly lead to combinatorial explosion, and are thus feasible only for small, low-period oscillators.

In this paper, we focus on the CA known as Conway's Life [Berlekamp et al. 1982, Chapter 25]. Our results could not have been accomplished without detailed knowledge of the behavior of this CA in particular. Nevertheless, it is almost certain that a variation of these techniques could be used to arrive at analogous results for other CA rules. We present a proof that for any $n \geq 58$, there is an oscillator of period n that operates under Life rules, and give three explicit methods to construct such oscillators. The first method, due to Buckingham [1996], results in large oscillators for large values of n, but is the only one that works for periods up to several hundred. The second and third are used to prove optimal asymptotic size bounds under two distinct definitions of size.

Why Life?

Though called a game and often considered a topic in recreational mathematics, Conway's Life is an important member of a very natural class of cellular automata. In such a CA, a cell has only two states, 0 and 1, and its transition rule is based on just its current state and the sum of the values of the eight cells surrounding it on the cartesian grid. (Specifically, in Life a cell with a state of 1 keeps this state if the sum of its neighbors is two or three. A cell with a state of 0 changes to 1 if the sum of its neighbors is exactly three. In all other cases, its next state is 0.) There are 2^{18} such CA rules possible, and of these, Life is arguably the simplest one that achieves its characteristic balance between rapid stabilization and uncontrollable chaotic activity (this corresponds to Class 4 of [Wolfram 1984]).

We believe there are no results of the type presented here for any other CA not designed with these results in mind. In addition, we are unaware of a simpler CA for which such claims are even plausible. The components we present, though simple to verify, often operate in a very counterintuitive, "chaotic" manner. They must, of course, be combined in a more predictable and structured manner.

Computers have been a useful tool in this research, both for enumerating combinations of subpatterns and verifying the results. It should be emphasized, however, that significant advances have all come from an informed judgment of what is possible, based on computer-aided manual experimentation. So far, computers have primarily been used to speed up the design process and fill gaps left by a manual search. Much potential remains for increasing the level of automation, suggesting that Life may merit more attention from the computer search community.

Conventions

We usually assume that Life operates on an infinite grid of cells whose state is 0 by default. A configuration, typically called a *pattern* in this context, is considered finite if it contains a finite number of cells with the value 1. The *population* is the total number of cells having state 1. The *bounding box area* of the pattern is the area of the smallest axis-parallel rectangle containing all cells having state 1. The *diameter* of the pattern is the maximum value of the Manhattan distance between any two cells having state 1. Life is often visualized as a grid covered with a pattern of round markers. A cell with a marker is considered to have a state of 1, an empty cell a state of 0. We will adhere to this convention in our illustrations.

2. RESULTS

We provide several methods to construct oscillators of all periods, with the consequence that oscillators of all periods 58 and above can be constructed. These methods can be used to produce instances that are small enough to observe in operation on any reasonable desktop workstation or PC. We have actually built explicit examples using every method, but for brevity we will omit many details from our discussion of the method needed to prove the second complexity bound.

We summarize our results as theorems. The proof is deferred until later.

Theorem 1. *There exists a constant a and a family of oscillators $K_{58}, K_{59}, K_{60}, \dots$ functioning under Life rules such that K_n has period n and the population of K_n never exceeds a.*

Theorem 2. *There exists a constant b and a family of oscillators $D_{58}, D_{59}, D_{60}, \dots$ functioning under Life rules such that D_n has period n and the diameter of D_n never exceeds $b\sqrt{\log n}$.*

Our bounds are within a constant of the best possible. This is obvious in Theorem 1. To see that this is true for Theorem 2, note that a period-n oscillator must have n distinct states, which is not possible unless it has a bounding box area of at least $\lceil \log_2 n \rceil$. Thus, its diameter must be $\Omega\left(\sqrt{\log_2 n}\,\right)$. With this in mind, it may briefly seem puzzling that a constant population bound is attainable. There is no contradiction here: in this case, the diameter of the oscillator grows linearly with n and the period is determined by the length of the path taken by a common type of 5-cell spaceship, called a glider (Figure 1), as it travels between two patterns that act as reflectors.

Gliders are an important component of many of our constructions, and were one of the first phenomena ever observed in Life. After two iterations of Life rules, a glider becomes identical to its reflection across a diagonal, shifted slightly. After two more iterations, the glider becomes identical to its original pattern but shifted by one diagonal cell. Because this pattern consists of just five cells, it occurs as a by-product of many other reactions. In fact, other small patterns with special properties often appear spontaneously, and all of our results rely heavily on this ability to exploit phenomena that might best be described as combinatorial coincidence.

FIGURE 1. The two shapes of a glider.

Ad hoc constructions are known for many lower periods (see [Hensel 1995]) and, as of this writing, the only gaps that remain under the definition used here are periods 19, 23, 27, 31, 37, 38, 41, 43, 49, 53, and 57. Periods 33, 34, 39, and 51 have only been realized as uninteresting combinations of lower-period oscillators; these are ruled out according to a stronger definition. The diagrams on the next three pages show examples of oscillators of all periods not covered by the constructions presented here. These ingenious designs are the work of many different people, and have been compiled by Dean Hickerson. (Credit for a particular oscillator is available on request.)

For completeness, we note that there is a class of oscillator-like patterns called *glider guns* that become identical to their initial state after n turns, with the exception that they leave an extra glider behind. This glider moves away in time for the next to appear, giving the pattern the effect of a machine gun firing off a stream of gliders at a regular frequency. This pattern can be interpreted as a simple oscillator if we consider it to be attached to an infinite stream of gliders, though this interpretation is not very useful for present purposes. A glider gun is an important device for many other constructions in Life, and we note here that both of the above theorems can be rewritten as results on glider guns. The only important change is that, for technical reasons, some oscillators of period 61 and below have not been modified to produce a stream of gliders. Thus, we must begin with $n = 62$.

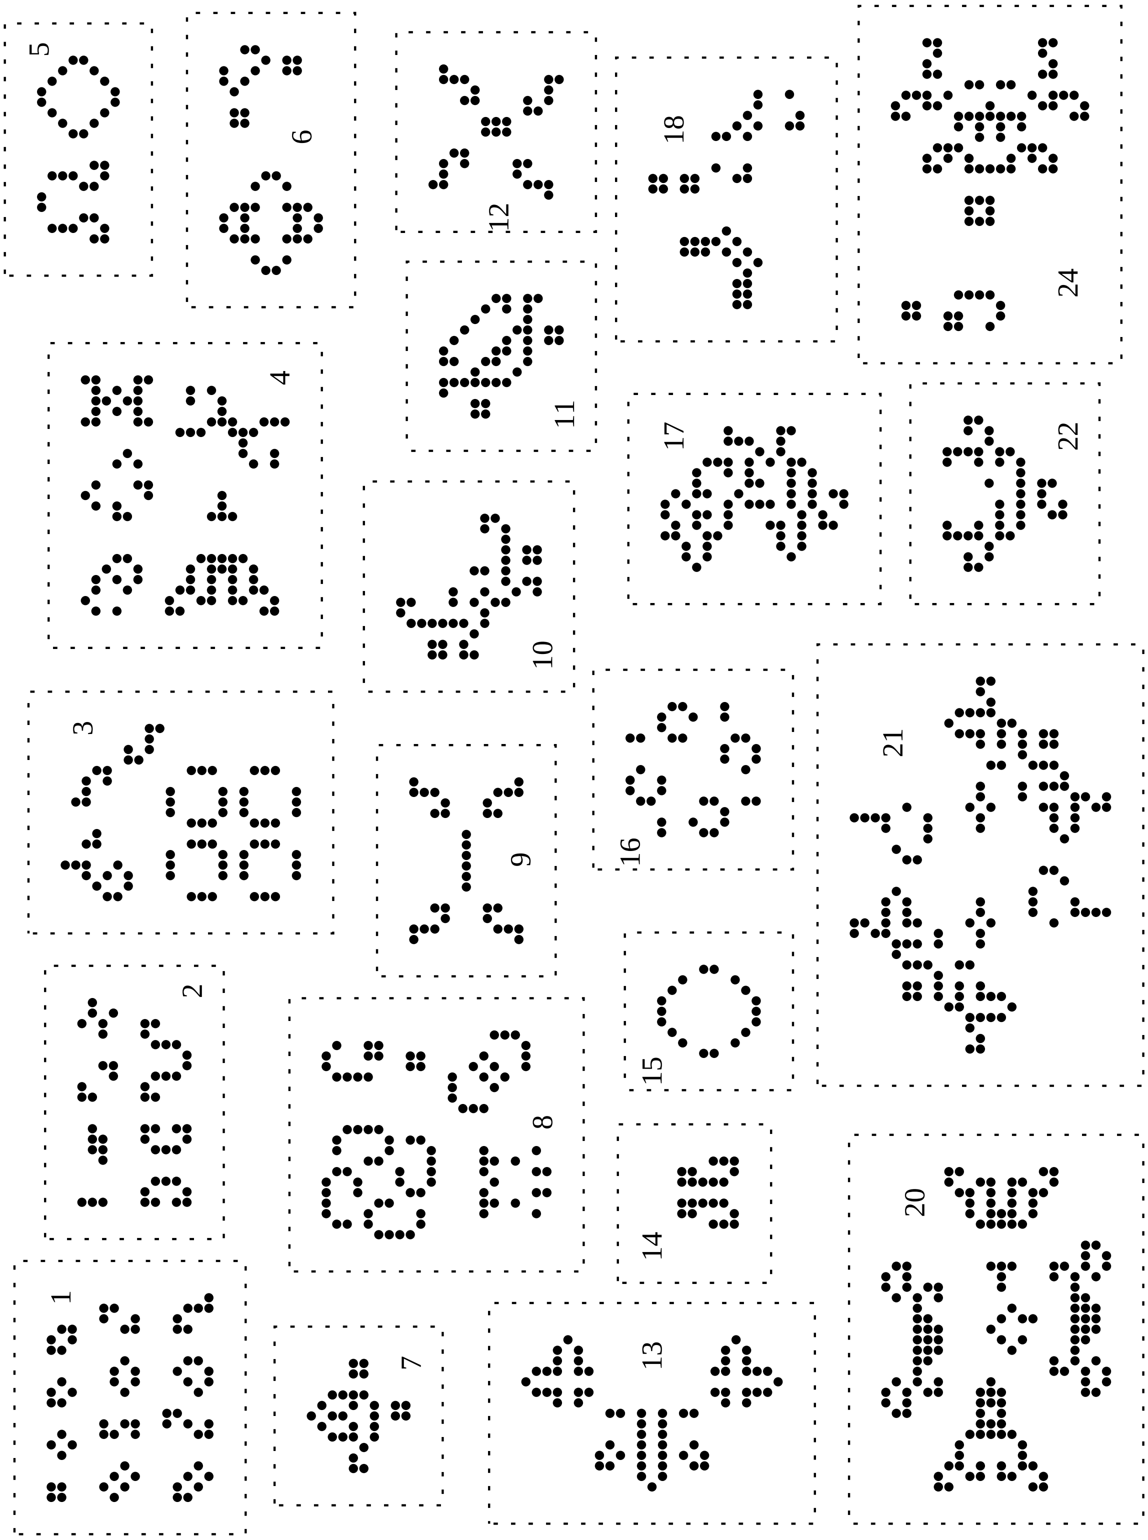
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
20
21
22
24

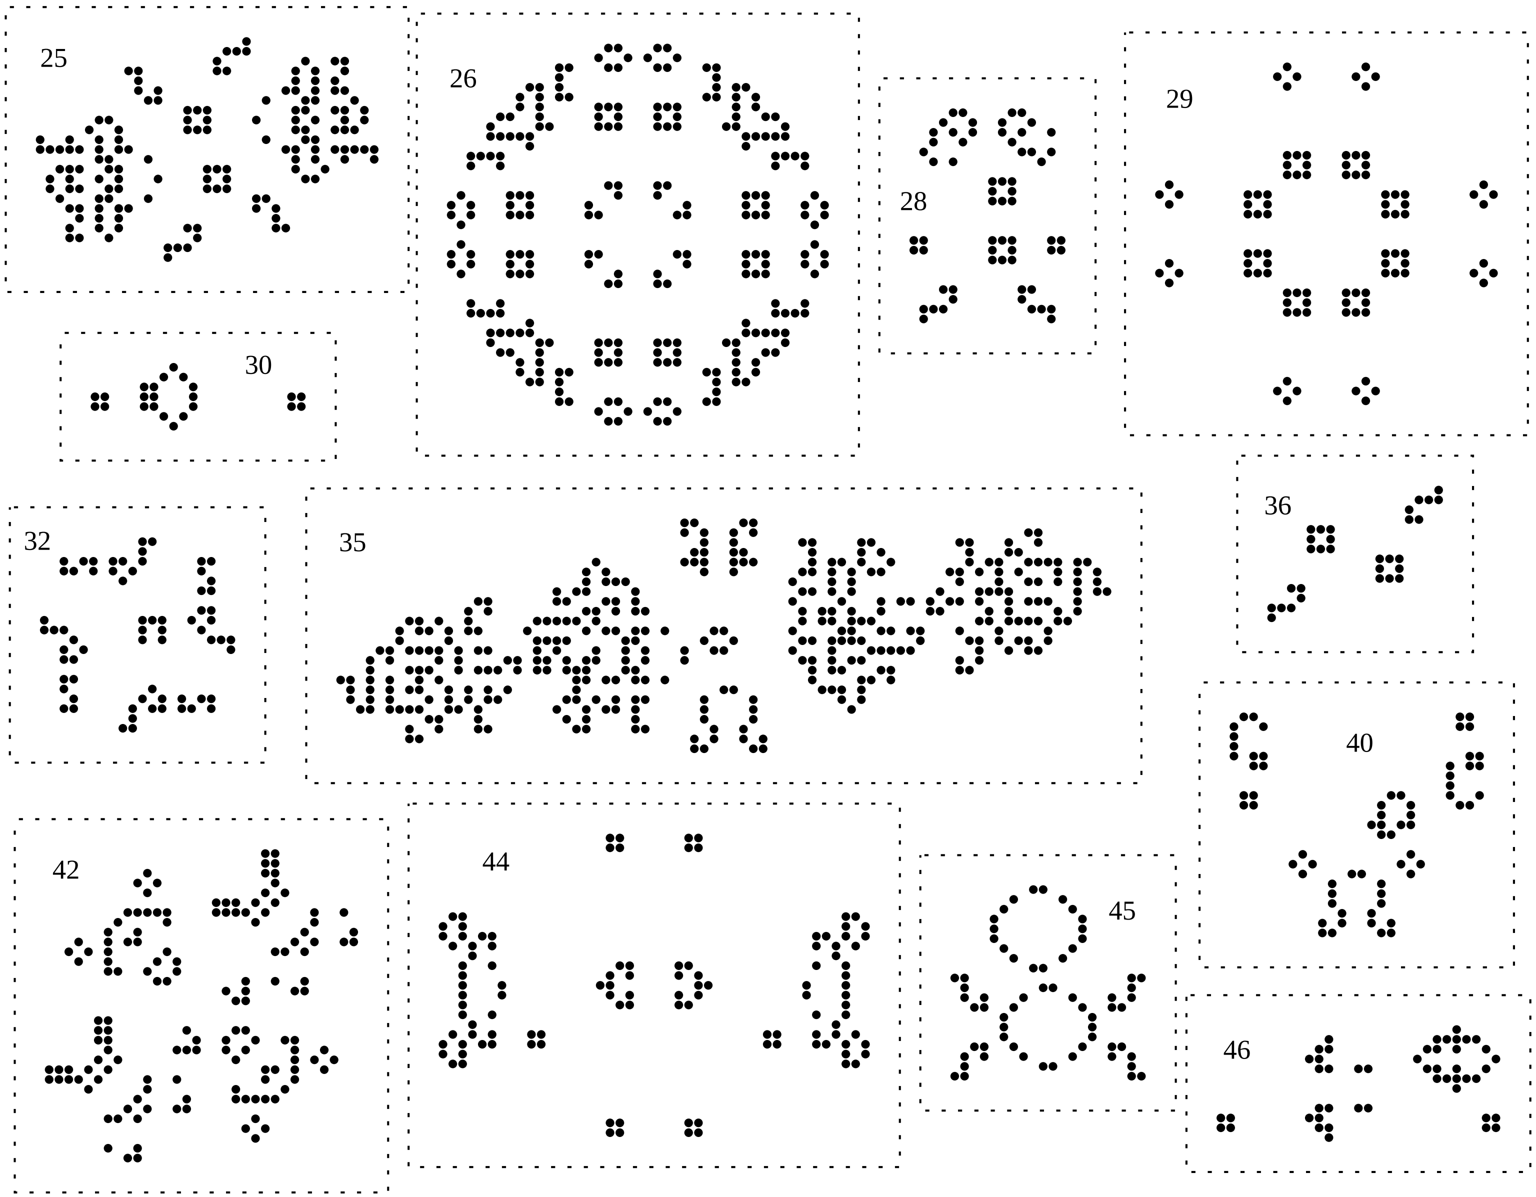

FIGURE 2. Oscillators realizing periods in the range not covered by our algorithm. (Continued on next page.) Box 1 shows twelve still lifes in three rows; box 2 has four oscillators on the top row and two on the bottom row (the lower left one is symmetric and in two pieces); box 3 shows the familiar pulsar at the bottom, and there are two oscillators side by side above it; box 4 has three on the top row and and two on the bottom row; boxes 5 and 6 show two oscillators each, side by side; box 8 has two rows of two oscillators each.

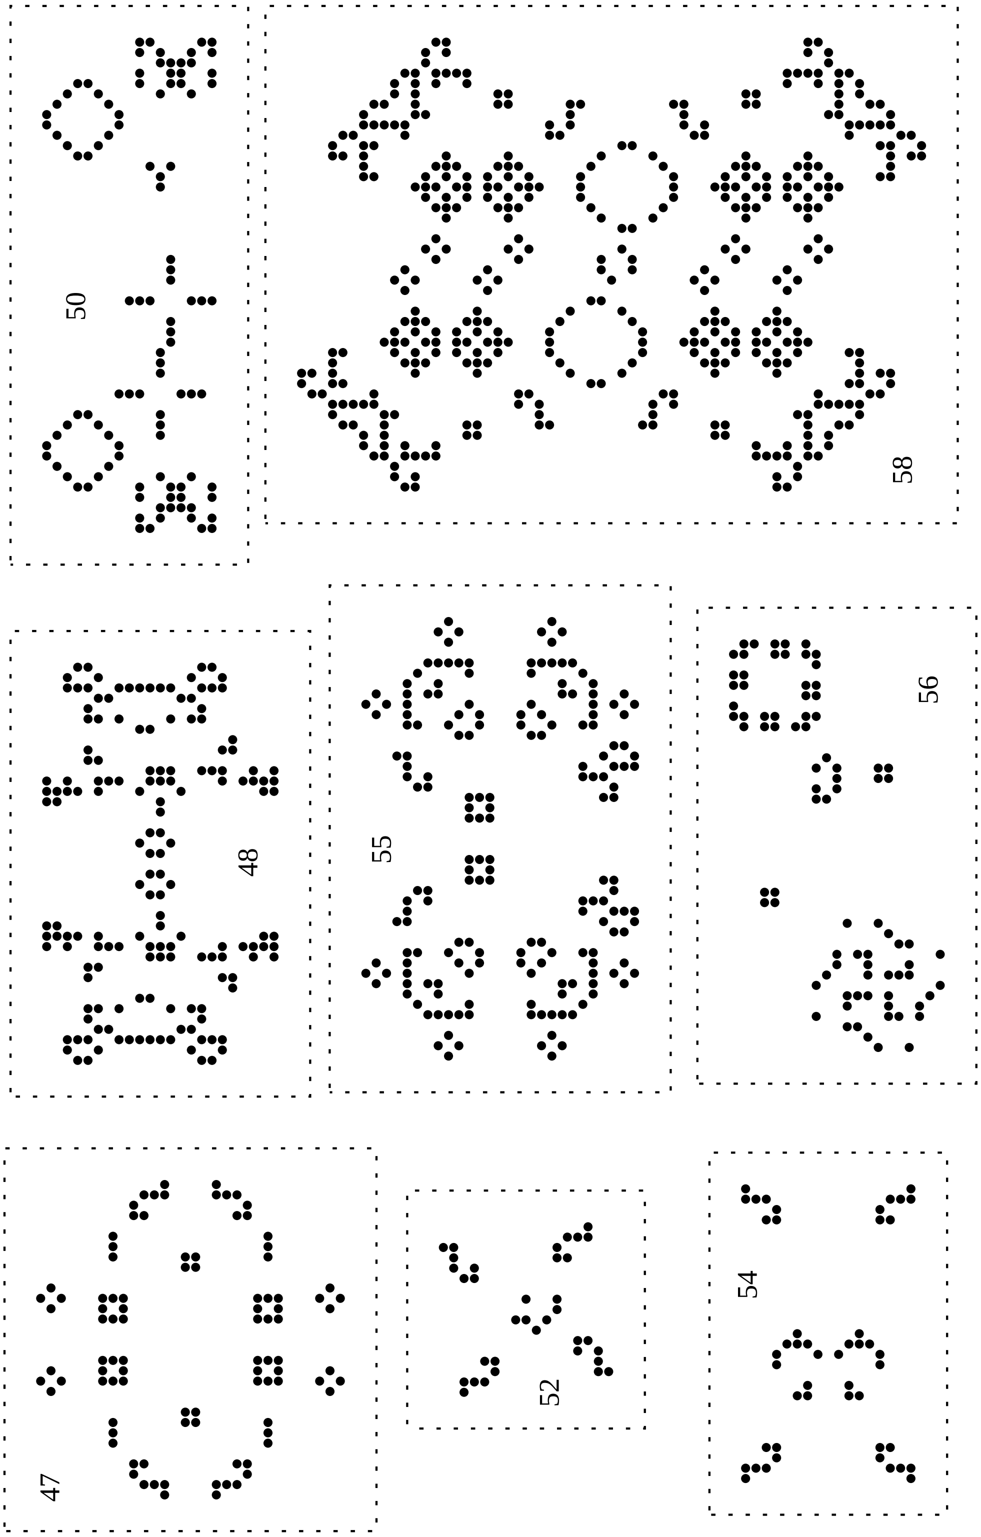
50
58
48
55
56
47
52
54

3. PREVIOUS RESULTS IN LIFE

There is a long history of constructive proofs to show that Life exhibits certain properties. Initially, it was not even clear that unbounded population growth was possible in Life. While many similar CA rules result in fast, explosive growth, small random patterns in Life tend to stabilize eventually, though they can remain active for a long time. The question of unbounded growth was later answered in the positive with the discovery of the glider gun, discussed above, and several other patterns discussed elsewhere [Berlekamp et al. 1982, Chapter 25; Poundstone 1985].

Another early problem in Life concerned so-called Garden of Eden patterns. These are patterns that cannot be the result of applying the Life rule to any previous pattern. A general proof by Moore [1962] suffices to prove the existence of such a pattern in Life, and indeed in many cellular automata of interest. Nevertheless, the problem of finding a reasonably small instance remained open for several years until an explicit example was verified; see [Berlekamp et al. 1982; Poundstone 1985], for example.

Finally, with the accumulation of many versatile interactions between gliders and other patterns, it became clear that a universal computer could be embedded in Life. This is hardly a surprising property to find in a complex system, though it is not necessarily an easy one to demonstrate. More significantly, it was shown how to embed a *universal constructor* in Life [Berlekamp et al. 1982], a pattern that allows the self-replication of arbitrarily complex computers. The idea of a universal constructor was introduced by von Neumann [1966] who designed a very complex CA rule in order to support such a configuration. One of Conway's original aims in developing the rules to Life was to find a more elegant cellular model for which von Neumann's results would hold. With this result, it became clear that Conway's Life was one of the simplest mathematical models of universal computation and self-replication.

Since that time, a large body of unpublished work has refined and extended these results. Many of the new results are more specialized and concern, for example, the design of patterns with unusual growth rates such as $\Theta(n \log n)$, or linear but whose growth coefficient has an irrational limit (compiled in [Callahan 1997; Hensel 1995]). Certain special-purpose computers have been embedded: for example, a pattern by Hickerson outputs spaceships at prime-numbered intervals [Hensel 1995]. This uses an implementation of the sieve method based on the geometry and dynamics of the cellular space. Besides these large constructions, a collection of small, "natural" oscillators and spaceships has been accumulated by applying intuition and a variety of automated search techniques [Hensel 1995].

It is not clear that much is to be gained by having a collection of computational devices that work under Life rules. However, the general design techniques should carry over to other systems. With increasing miniaturization, we must eventually reach the point at which the device we are able to manufacture is determined by the laws of physics rather than the wishes of the designer. While this point may be far in the future, it suggests at least one instance in which it may be of practical importance to find small embeddings of computational devices in complex systems whose rules we cannot modify.

4. THE OSCILLATOR PROBLEM

The results of this paper differ from past results in the "time grain" of the embedding we exhibit. That is, in all past instances, a single time step in the system to be embedded could only be simulated by a sequence of several time steps in the embedding, typically at least 30, which is the period of the simplest known glider gun. Consequently, most known constructions that are useful for general computation cannot be used to embed arbitrary-period oscillators. While it has long been known how to build any oscillator of period $30n$, it has been clear that the techniques used do not extend to arbitrary periods.

Alternative Possibilities

To see that the present result is noteworthy, we consider alternatives. At one extreme, we might have a CA that usually settles down into a disjoint collection of small oscillators of "natural" periods. In this case, a plausible conjecture is that there are a finite number of such natural periods and that the highest attainable oscillator period is their least common multiple, the result of considering the disjoint collection to be a single oscillator. This may be the case for some cellular automata. Early observations in Life, however, suggested a more interesting possibility.

The *pentadecathlon* is a well known period-15 oscillator that was found very early by hand. It was observed that a glider hitting this along a particular diagonal at a particular phase will be reflected 180°, as shown in Figure 3; the reflected glider is indicated by hollow markers. This reflection can be repeated in the other direction. Thus a single glider can be sent on a loop between these patterns, forming an oscillator, the period of which depends on the total trip time. Because the reflecting device has period 15, the trip time for a glider is restricted to be a multiple of 15. Other restrictions are related to the speed of the glider, which takes four steps to move one diagonal. As a result, the oscillators that we can construct all have periods of the form $60 + 120k$. This is sufficient to refute the pessimistic conjecture considered above, but it is a long way from obtaining all periods. Thus, a second plausible conjecture is that given enough reflector mechanisms, all multiples of the minimum reflector periods can be obtained.

Stable Reflectors

The above observation leads naturally to the *stable reflector problem*, namely, to find a stable pattern that, when hit by a glider along some path, reflects one or more gliders, after which its original state is restored allowing it to be reused. (A *stable* pattern is a period-one oscillator, also known by Conway's punning term "still life." We use "stable pattern" here because this notion can be generalized to other CA rules.)

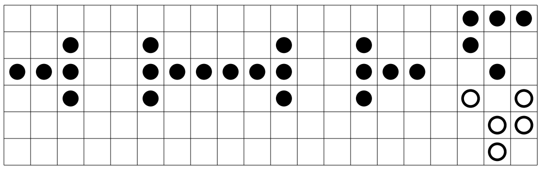

FIGURE 3. A pentadecathlon reflects a glider.

The existence of such a pattern was known to follow from the existence of a universal constructor, but for many years, its explicit construction appeared beyond the scope of known techniques. In this paper, we provide an explicit example of a stable reflector, and use it for our second method of constructing oscillators. It is not hard to see that stable reflectors satisfying some additional properties can be used to realize any desired oscillator period by adjusting glider trip times. We defer the discussion of these properties until Section 6.

Though a stable reflector is a very desirable construction, there is a reason why the first reflectors found tended to be of higher periods. A reflector has to interact with a glider without being itself affected by the interaction — or if it is, it must somehow repair any damage. Many higher-period oscillators produce *sparks*, or collections of cells that appear at every cycle but can be removed without changing the future of the remaining cells in the pattern. Left alone, these sparks vanish harmlessly. Sometimes, however, there is a way to collide a glider with sparks such that the product is a glider in another direction. The result is a reflector.

A stable pattern obviously produces no sparks. Thus, *a stable reflector must be damaged by the glider that hits it.* It must then produce a new glider, as well as repair the damage caused by the old one. In fact, there are small stable patterns that exhibit limited capacity for self-repair. This occurs, as usual, by coincidence and not deliberate design. Best known among these is the so-called *eater*, a 7-cell pattern that can destroy a glider

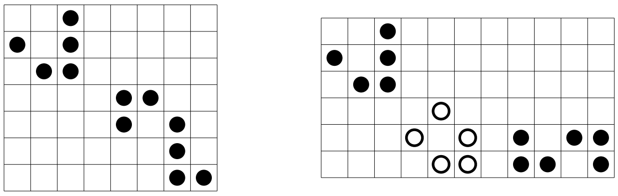

FIGURE 4. Left: a glider on the path to being eaten. Right: the boat bit reaction (boat shown with hollow markers).

without itself being destroyed. Less well known, but very useful, is the *boat-bit* reaction, in which a glider is converted by a 6-cell stable object called a *snake* into a 5-cell stable object called a *boat*. Another glider along the same path deletes the boat cleanly, and this combination can be used to build a toggle memory. These reactions are shown in Figure 4. No similar reaction has been found that produces a new glider as a result, and a fairly complicated mechanism for self-repair was ultimately required to produce a stable reflector, as we will later discuss.

An existence proof has already been mentioned for a stable reflector, based on the existence of a universal constructor. (This reasoning was probably arrived at independently by several people, so appropriate attribution of the proof is difficult.) The idea is that a certain field of stable objects can be triggered by a glider in a chain reaction, to convert it into any desired collision of gliders. This collision, in turn, can be one that synthesizes a universal constructor. In particular, it must synthesize one that is programmed to (a) rebuild the field of stable objects, (b) output the reflected glider, and (c) annihilate itself cleanly with a hail of gliders on carefully chosen paths. While this scenario carries an undeniable Alice in Wonderland appeal, the fact remains that a universal constructor has never been built explicitly. Furthermore, any pattern based on it would certainly be too large to view in operation on a reasonable machine at a reasonable pixel resolution. Thus, for those seeking concrete verification of such theoretical claims, a more parsimonious mechanism is required.

5. HERSCHEL-LOOP OSCILLATORS

The first explicit method of building oscillators of any period was discovered very recently by Buckingham [1996]. This method does not use the stable reflector. However, some of the same components are required to build a functioning stable reflector. In addition, the current fastest stable reflector requires 623 steps to restore its state and accept another glider, limiting its use to oscillators of period 623 and above. Buckingham's method is the only uniform construction of an infinite family of all oscillators of period 58 and above.

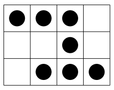

FIGURE 5. The Herschel.

Like the stable reflector, Buckingham's construction uses the idea of self-repairing stable components. Instead of redirecting the path of a glider, his components reposition a 7-cell active pattern called a Herschel (Figure 5) that by itself does not even function as a spaceship. Rather, it proceeds chaotically, emitting two gliders and eventually leaving behind a small amount of stable debris. Careful observation suggested that the Herschel was one of the most likely candidates to tame with an arrangement of self-repairing stable patterns. Several such arrangements cause it to reproduce itself cleanly at a new position and orientation, producing gliders as a by-product.

Buckingham conceived of the above notion in 1973 (see [Buckingham 1996]), but was unable to fit all the pieces together until over twenty years later, after collecting many interactions that came close to satisfying this goal. Callahan was able to extend the set of conduits using a more automated approach, though in principle all of the known components could have eventually been found using a computer-aided manual search within a feasible time frame.

The Herschel is the product of a better known pattern called the *B-heptomino*, which often arises

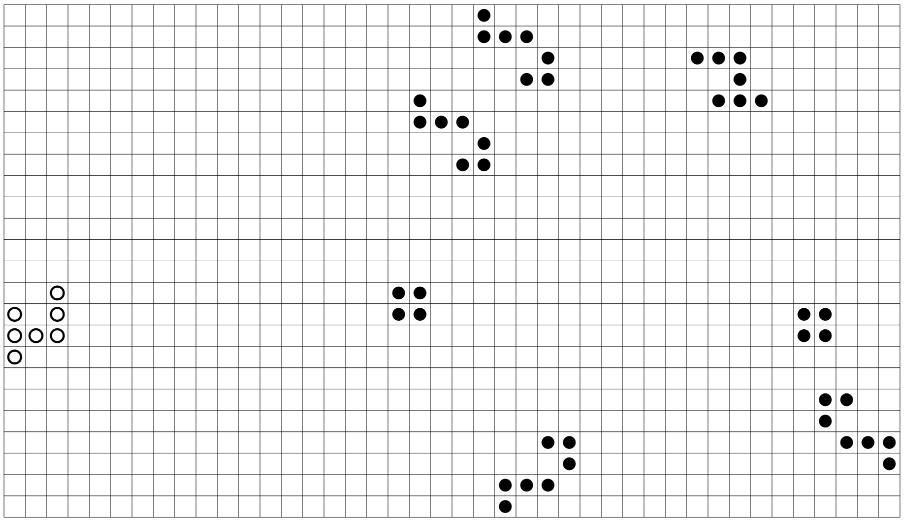

FIGURE 6. 112-step right turn Herschel conduit.

spontaneously in random configurations. The latter pattern evolves into a Herschel and a *block*, which is a 2×2 stable arrangement of live cells appearing very commonly in debris. Most Herschel synthesis techniques work by first producing a B-heptomino and then eliminating the extra block one way or another. The B-heptomino can itself be synthesized from the well known R-pentomino. All of these pathways are significant in the development of new devices. For brevity, however, we will restrict our discussion to the Herschel, and view the stable Herschel-moving patterns (*Herschel conduits*) as black boxes.

One such Herschel conduit (fourteen are now known) is shown in Figure 6. The Herschel is shown here in the standard orientation that we use for reporting the transformations realized by each conduit. The position of a Herschel is taken to be the position of its center cell (i.e., the cell whose removal leaves two disconnected sets of three cells each). In the figure, the original pattern is superimposed with the position of the Herschel after 112 time steps. The objects other than the Herschel (four eaters and two blocks) look exactly as they do initially at this point, while the original Herschel is eliminated and replaced by another Herschel, shown with hollow markers. In addition, a glider is emitted, traveling north-east from the original position of the Herschel. All the Herschel conduits function in a similar fashion. The emitted gliders can be suppressed with additional eaters, but often they are useful in other constructions.

The list of known Herschel conduits is summarized in Table 1. Using this table, we can compute the effect of composing Herschel conduits in sequence. Not all sequences are possible, of course. Some lead to self-intersection, others to more subtle forms of interference. However, with so many degrees of freedom, we can generally find a way to move a Herschel along a path — if sometimes a complicated one — to a position and orientation that we need for a particular construction.

For Buckingham's oscillator-building method, we are mainly interested in moving the Herschel back to its original position and orientation. The number of steps it takes to accomplish this will be the sum, n, of the delay values for all the components in the path. We refer to such a path as a *Herschel loop*. After n steps, the Herschel will be in a position to repeat the same trip, and will continue this circuit forever. By suppressing all the emitted gliders, we obtain a period-n oscillator. The simplest instances of Herschel loops are obtained by composing four 64-step left turns for a period-256

Transformation	Shift	Delay	Found by
left turn	$(9, 11)$	64	DJB
x flip	$(-8, 25)$	77	DJB
right turn	$(-33, 12)$	112	DJB
identity	$(1, 32)$	116	PBC
identity	$(-6, 40)$	117	DJB
x flip	$(14, 20)$	119	DJB
x flip	$(-4, 48)$	153	PBC
right turn	$(-41, 17)$	156	DJB
identity	$(3, 49)$	166	PBC
x flip	$(-5, 27)$	158	DJB
x flip	$(0, 45)$	176	PBC
left turn	$(16, 24)$	190	DJB
NE diagonal flip	$(40, 17)$	200	PBC
NW diagonal flip	$(32, 7)$	202	DJB

TABLE 1. Known Herschel conduits.

oscillator, or four 112-step right turns for a period-448 oscillator.

Infinitely many paths can be constructed that translate the Herschel by some amount and rotate it 90° left or right of its original orientation. These compound right and left turns can be replicated four-fold in the same manner as atomic turns, giving us a Herschel loop whose period is four times the number of steps through one of these paths. Two issues remain to be resolved if we wish to attain an oscillator period as low as 58. First, it is clear that the time around a Herschel loop is always much greater than 58 steps. Second, if a loop is composed of four identical turns, then the time around it once must be a multiple of four.

The key to resolving these issues is the notion of *period dividing*, an idea we will use again when we discuss the next type of oscillator. Period dividing works as follows. If the time around a Herschel loop is an integer kn, then, provided some spacing conditions are met, we can divide its period by k simply by replacing the single Herschel with k Herschels spaced n steps apart. (Spacing conditions must not be overlooked. For example, it is not possible to divide the period-256 oscillator by 4, because its turn elements require incoming Herschels to be spaced at least 153 steps apart. Nevertheless, a period-64 oscillator can be constructed using other conduits that allow closer packing.) Then, after n steps, any given Herschel is replaced by its predecessor in the sequence of duplicates, and overall the pattern is identical to the one we started with and is therefore a period-n oscillator.

The problem of building a period-n oscillator can be reduced in this way to the problem of finding a Herschel loop such that the total time around the loop is some multiple of n. This problem in turn can be solved by applying the following lemma.

Lemma 1. *For any integer $n \geq 1$, there exists a k such that a functioning Herschel path can be built that translates the Herschel by some amount and rotates it by 90° in kn steps. Moreover, this path can be composed with three copies of itself to form a $4kn$-step Herschel loop.*

Idea of the proof. A fully detailed proof would require an analysis of the Herschel conduits themselves. The basic idea, however, is analogous to the fact that given any two relatively prime integers a and b, we can find, for any n, a multiple of n that has the form $ai + bj$ where i and j are nonnegative integers. It is relatively easy to build a non-self-intersecting Herschel path that translates without turning. Two such compound translations can be built whose delays are relatively prime. These can be used in combination, along with one of the turns, to build a compound turn such that the time through this turn is a multiple of n. □

This gives us a mechanical method of obtaining a Herschel-loop oscillator of any sufficiently large period n that we desire. We simply find a turn that is a multiple of n, self-compose it into a Herschel loop, and populate the loop with enough Herschels to reduce the total period to n. We must still be careful about close packing. Buckingham observed that Herschels could be packed as close as 58 steps from each other in the 77-step and 112-step conduits. All other conduits require greater spacing. Such close packing is only possible if we suppress

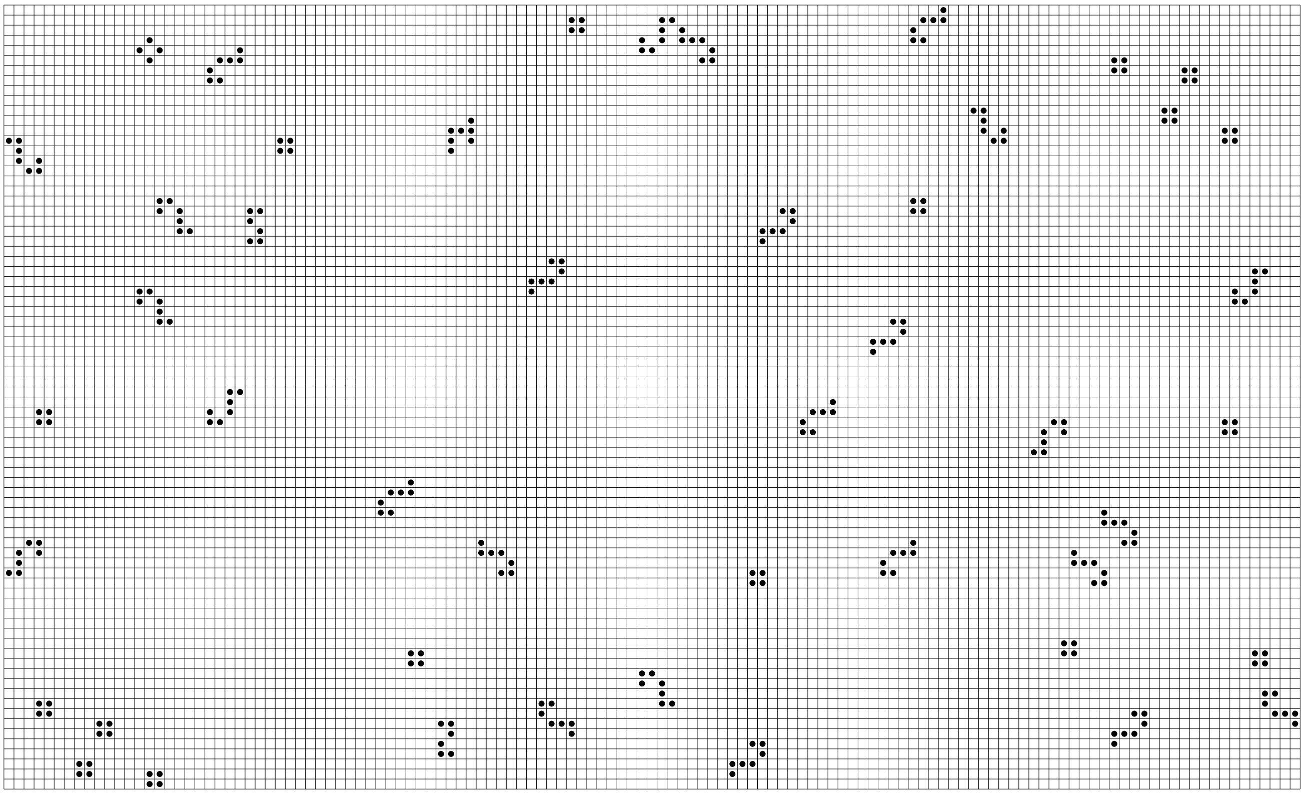

FIGURE 7. A 937-step Herschel loop.

gliders in these stages as soon as possible. Originally, Buckingham [1978] used a pattern called *eater*-2 to allow packing as close as 61. He later reduced this using a new eater found very recently by Hickerson [1997].

No glider gun of period less than 62 has been constructed using this method. However, the 77-step and 112-step conduits are sufficient to build oscillators of all periods 58 and up. The size of the period-61 oscillator can be reduced somewhat if we use the 153-step conduit as well, and much smaller period-58 and period-60 oscillators have been constructed using other techniques. Period 60 is particularly easy to realize, for example, using two pentadecathlons to reflect a glider, as noted earlier, or using one of a number of interactions between period-30 glider guns. Buckingham gave explicit examples of glider guns of period 62 through 69 that can easily be turned into oscillators. We summarize these results as follows.

Lemma 2. *For any $n \geq 58$, a Herschel loop can be constructed and populated with Herschels to realize a period-n oscillator.* □

Herschel loops allow the construction of the lowest period oscillators of all our methods. However, it is clear that the size of such a loop is $\Omega(n)$ measured by population and at least $\Omega(\sqrt{n})$ measured by diameter. Thus, while this is a useful method, and can often be used to build the smallest known oscillators of certain periods, it is not the best method asymptotically.

It is worth observing that the above construction always gives a Herschel loop whose time around is a multiple of four. There is no reason why odd-period — or, for that matter, prime-period loops — should not exist; however, they must necessarily be asymmetric and more difficult to find for that reason. One can enumerate all possible loops up to a few thousand steps using exhaustive search.

In this way, it was possible to find a Herschel loop with a total trip time of 937 steps (see Figure 7). The result is a fairly small prime-period oscillator (or alternatively, a glider gun) that contains just one Herschel.

6. STABLE-REFLECTOR OSCILLATORS

Compared to the design of an appropriate Herschel path, it is a relatively simple matter to build an oscillator out of a stable glider reflector. The problem, of course, is building such a reflector in the first place. We consider the simpler problem first and then return to the design of a stable reflector.

Suppose we have a stable reflector such that its reflected glider is turned 180° from the incoming glider, and there is enough clearance between the paths of these gliders that another incoming glider along the same path as the first will not interact with its reflected glider. This is the simplest kind of reflector that can be used to build arbitrary-period oscillators. Without the condition of clearance, we could only send one glider at a time between reflectors, making it impossible to apply the method of period dividing and obtain all periods above a certain value. A stable reflector that realized a 90° turn would also work, since it is possible to build a 180° reflector by composing two of them. We consider the 180° reflector here because it leads to the simplest construction.

The above conditions place no constraint on the phase of the reflected glider, which turns out to be quite important. We define the *phase shift* of a 180° reflector as follows. Let g be the original glider as it is approaching the reflector and let g' be its reflected glider. Apply the Life rules $4m$ times to the pattern containing both glider and reflector, where m is large enough that g' has appeared. Now apply the Life rules again either 0, 1, 2, or 3 times until g' is identical to the original shape of g but rotated 180°. That number, between 0 and 3, is the phase shift of the reflector. To simplify our analysis, we assume the stable reflector in question has a phase shift of 0.

Lemma 3. *Suppose we have a stable* 180° *reflector with a phase shift of* 0*. For all sufficiently large* n*, we can build a period-*$8n$ *oscillator by reflecting a glider between two copies of this reflector.*

Proof. Consider the closest spacing of the reflectors in such a pattern. After $4m$ turns, for some m, the glider g will be replaced by another glider g', identical but rotated 180°. So, after an additional $4m$ turns, by symmetry, the glider g will be restored to its original position and orientation. Thus, this pattern is a period-$8m$ oscillator. Now if we separate these reflectors by just one additional cell diagonally, the time of the total trip taken by the glider back to its starting position will be increased by 8 steps (4 steps for the increased distance each way), giving us a period-$8(m+1)$ oscillator. More generally, if we increase the separation by k diagonal cells, we obtain a period-$8(m+k)$ oscillator. □

The preceding lemma shows that we can build an oscillator of any sufficiently large period that is a multiple of 8. Now we need only apply the method of period dividing, as in the previous section. Instead of using multiple Herschels, in this case we use multiple gliders. Thus, we generalize the above lemma as follows.

Lemma 4. *Suppose we have a stable* 180° *reflector with a phase shift of* 0 *and sufficient clearance to allow passing of incoming and reflected gliders. For all sufficiently large* n*, we can build a period-*n *oscillator by reflecting* 1, 2, 4, *or* 8 *gliders between two copies of this reflector. Moreover, there is a constant* k *such that the population of each such oscillator is* k *or less.*

Proof. We use the preceding lemma and a straightforward application of period dividing. The size bound follows from the fact that the oscillator consists of a constant number of gliders and two reflectors. □

Explicit Construction of a Stable Reflector

Recall that a stable reflector is inevitably going to be damaged by the glider that hits it. It must repair this damage and also output a glider. Our strategy for constructing such a pattern is to limit the amount of damage to something we can realistically expect to repair, but at the same time produce an initial reaction of sufficient duration to produce a new object (at least one new glider).

This balance is very difficult to obtain. Many stable objects will produce a long-lived reaction when hit by a glider, but they are destroyed in the process. Some other patterns can stop a glider in its path, either by destroying it as in the case of the eater or even leaving a stable object in its place, as in the case of the boat bit. These patterns are not destroyed in the process, but the reaction they produce is very rapid and an unlikely candidate for producing a glider.

The most likely scenario is that the glider would have to destroy at least one simple stable object in order to "borrow" enough activity to produce a new glider. The explosive reaction caused by this collision could then be amplified and controlled by surrounding stable objects, eventually producing the means to "pay back" the loan by reconstructing the stable object and cleaning up any debris. Assuming this is possible at all, there should be little difficulty in emitting one or more new gliders in the process. In fact, this idea is similar in flavor to the proof that assumes a universal constructor. The difference is that we would like to minimize the initial damage so it is reparable by a small, special-purpose device.

In practice, "minimize the initial damage" means "first collide the glider with a block." This was observed by others working on the problem [Schroeppel 1994] and independently by Callahan. The block is the most common stable object found in random residue, and seems the most likely candidate to restore at a particular position and orientation, the latter guaranteed by the block's eightfold symmetry.

There are several ways to collide a glider with a block. The most promising way to begin with seemed to be one that produces the *pi heptomino*, a symmetric pattern that expands rapidly for over 150 steps before stabilizing into a symmetric arrangement of stable and period-2 objects. We began with this collision and sought ways of interacting it nondestructively with surrounding stable objects, primarily eaters and blocks, both known for their limited self-repair mechanisms. This is known as *perturbing* a reaction. The goal was to find such an interaction that at least restored the block in its original location. This event is so unlikely that any search is bound to produce a small enough list of candidates to examine by hand.

The Block-Repair Mechanism

Figure 8 shows the most promising block-repair mechanism obtained so far. Callahan found this early in the process of developing an effective computerized method of enumerating perturbations by multiple objects. Only the upper block and lower eater are needed to restore the collision block, however, placing it well within the range of computer-aided manual search. In fact, this arrangement had been considered previously by Buckingham, but at the time there was no way to eliminate a major difficulty, discussed below.

The glider is shown at the last step before the collision. In the figure, the pattern is superimposed

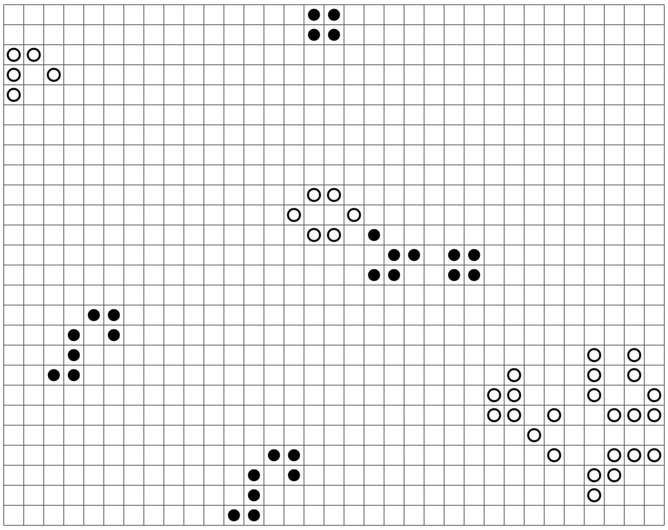

FIGURE 8. Initial stage of the stable reflector.

with its product 104 steps later, shown with hollow markers. At this step, the two eaters and block surrounding the collision have perturbed the explosion, resulting in

- the restoration of the block involved in the collision,
- a glider traveling in the opposite direction,
- a small, chaotic pattern to the lower right of the block, and
- a 6-cell stable object known as a *beehive*, adjacent to the block and obstructing the path of a glider sent on the same path as the original one.

Only the last of these products is undesirable. Most of our later effort will be spent eliminating it. The chaotic pattern is not useful as such, but can be perturbed into something useful; we have restored the block, so anything we can obtain from this pattern is effectively obtained for free. The reflected glider is one of the things we were hoping to obtain in the first place, though it is of limited use since it has a phase shift of 1.

The beehive appears at first a fatal error. Its position is deep in the interior of the active region, though this may not be obvious from the picture. Thus, there seems to be no way to place a stable object to suppress this beehive without interfering with earlier steps. We could eliminate the beehive cleanly with certain glider collisions originating from outside the active region, but this leaves the problem of producing such a glider.

The solution appeared in the form of Buckingham's announcement of Herschel conduits. Using a computer enumeration, Callahan found a pair of blocks that convert the chaotic reaction product into the R-pentomino. The latter is also chaotic but can be perturbed into a Herschel using several patterns found by Buckingham, exactly one of which works in this context. Using Herschel conduits, a Herschel can be used to generate many new gliders, and the paths of these gliders can be adjusted as desired using an appropriate sequence of conduits. The solution, then, was simply to send the Herschel along to an appropriate diagonal at which it could emit a glider to eliminate the beehive. While this scheme may appear somewhat cumbersome, one must admit it compares favorably to the solution that requires a universal constructor.

Construction of an appropriate Herschel path turned out to be relatively straightforward. In this way we constructed the first explicit stable reflector to operate under Life rules. In the original solution, it took several thousand steps to eliminate the beehive, limiting the packing of incoming gliders. Callahan later refined this to a reflector that reflects gliders as close as 894 steps apart. In this construction, the final Herschel can be cleanly eliminated or else used to emit as many additional gliders as desired. Hickerson [1997] designed a reflector using the block repair mechanism (see Figure 9) which works in 747 steps. It terminates with a special-purpose perturbation that destroys the Herschel in the process of converting it into a glider along a path that would take longer to reach using standard conduits. In either case, the use of the Herschel leaves us with a reflector that produces more than one reflected glider, resulting in an interesting form of glider gun in which we recirculate one glider on a closed path, producing an unbounded number of gliders as a product.

Buckingham designed an even faster (672-step) stable reflector using Herschel conduits without the block-repair mechanism.[1] This relies on the fact that a Herschel can be moved to nearly any desired position with conduits, and converted directly into a variety of simple, stable objects, several of which are quite readily converted into a Herschel after a collison with a glider. The new reflector is based on a collision with the 5-cell object called a boat.

The block-repair mechanism remains useful in constructions in which we need to convert a glider into a Herschel that is not destroyed when repairing the collision damage. Such constructions are used in the next section.

[1]The record holder as of this writing, at 623 steps, was designed by Stephen Silver [1997] using the block-repair mechanism and a new R-pentomino-to-Herschel conduit.

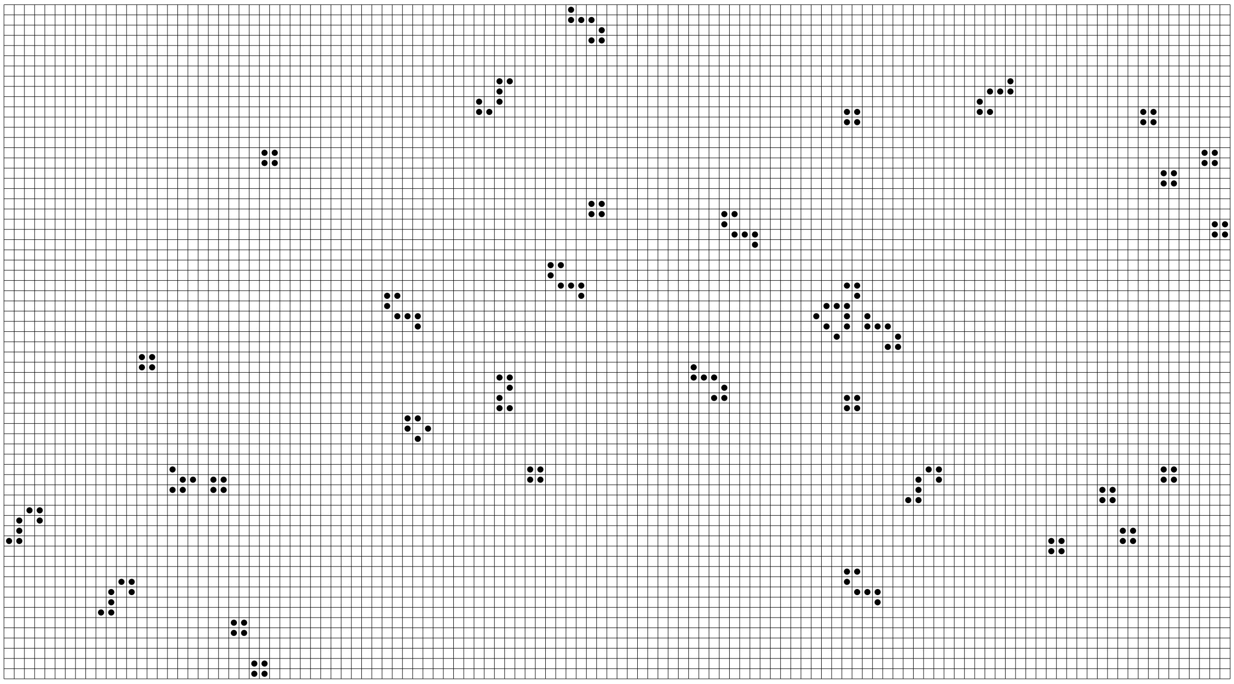

FIGURE 9. 747-step stable reflector.

Proof of Theorem 1. Hickerson's reflector produces several gliders reflected 90° with an even phase shift. Thus two such reflectors can be combined into a 180° reflector with a phase shift of 0 and sufficient path clearance to satisfy the conditions of Lemma 4. Combining this with Lemma 2 we have, finally, a proof of Theorem 1. □

7. LOW-DIAMETER OSCILLATORS

While the stable-reflector oscillator is optimal in terms of cell population, its diameter grows linearly with period. There is no reason why this must be the case in general. In fact, there are constructions that achieve very high periods and very low diameters. These have been known for many years, though they can only be used to obtain a limited set of periods. A useful component for such a device is a toggle memory that is triggered by a glider and that reflects a new glider on 1–0 transitions. Such devices can be chained together into a ripple counter, a well known element of digital design [Wakerly 1994].

Thus, the first idea for obtaining a high period is to send a stream of gliders into a chained sequence of such toggle memories. If the input period is n, the first toggle will have output period $2n$, the next $4n$, etc. By eliminating the final gliders, we obtain an oscillator of period $2^k n$ using a chain of k toggles. Assuming a compact layout technique, such as adding reflectors at intervals to snake the path into a more compact form, this can all be placed in a pattern with diameter $O(\sqrt{k})$.

Something like this can be built entirely of devices that have been known for years. Until recently, these periods had to be multiples of some known set of base periods n. More significantly, this construction seems to limit us to multiples of large powers of two. In this section, we show how to get around these problems. This requires a more complex design than either of our previous oscillator construction methods, so we present our results at a higher level of abstraction. We have built a working example of an oscillator using this technique that is compact enough (fitting in a 467×429 box) to view on any PC.

Glider-Triggered Glider Gun

The first new component we need is a glider gun that can lie dormant indefinitely as a collection of stable objects, but when triggered by a single glider becomes active, producing a sequence of gliders at a regular period until stopped (for example by colliding another glider with it). When stopped it returns to its original state, ready to be activated by the next glider.

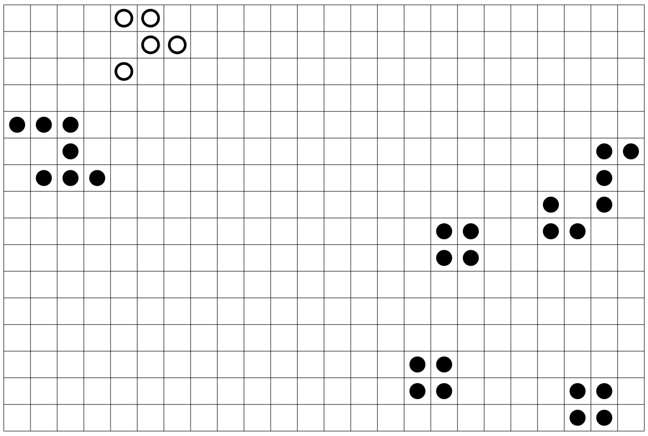

FIGURE 10. 119-step conduit leaves clearance behind its north-east glider.

We will skip the low-level details, but such a device can easily be built along the same lines as a stable reflector. The only additional property it must have is that it reflects a glider back into itself along the same path as the glider that triggered it, thus repeating the triggering process indefinitely. This requires some care, because most Herschel conduits do not leave any clearance behind the emitted gliders. An exception is the 119-step x-flip, shown in Figure 10. Using this device one can obtain the effect of capturing an entering glider and recirculating it around the same loop indefinitely. The loop is composed partially of Herschel conduits. Thus, it is a simple matter to output an additional glider each cycle, resulting in a glider gun. Moreover, one can stop the gun simply by blocking the path of the recirculating glider.

Herschel-Based Ripple Counters

Hickerson [1997] was the first to consider and implement a bit toggle based on a Herschel conduit. The idea is that the first Herschel travels through the conduit, leaving behind a small piece of stable debris. The next Herschel interacts destructively with the debris, cleanly annihilating both in the process. This gives us a conduit through which every second Herschel emerges. Chaining k of these together, we obtain a multi-stage conduit through which every 2^k-th Herschel emerges. The state as each Herschel enters is determined by the subset of stages containing destructive debris.

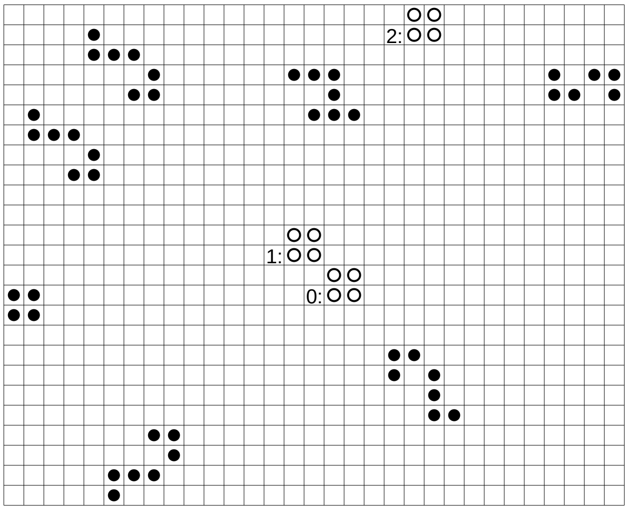

FIGURE 11. A quaternary counter stage.

Hickerson looked for a single conduit-like construction that accomplished this task, but was unable to find one. Instead, he composed a sequence of stages in such a way that the last stage emitted a glider that used the boat-bit reaction to leave debris in the first stage. This is an effective solution, but results in a rather large pattern. Unfortunately, a single-stage Herschel toggle seems to be a rarer object than intuition would suggest: after some trial and error using a computer to enumerate perturbations, we were unable to find one. (Dieter Leithner [1997] constructed such a device shortly after we completed the initial draft of this paper.) Callahan, however, found a single-stage quaternary counter based on the 112-step turn stage. This is shown in Figure 11.

The quaternary counter is like the bit toggle, but only allows every fourth Herschel to pass through. This is actually preferable for current purposes, since we would like to store a counter state as compactly as possible. The debris consists of a block in each case, and this block can be in one of three positions. When the block is absent, a Herschel passes through, leaving a block in the position marked 0 in the figure. When the next Herschel reaches the stage, it reacts with the block destructively, leaving only a new block, now in the position marked 1. Similarly, the next Herschel moves this to the position marked 2. At this last position, an ancestor of the Herschel entering from certain stages cleanly annihilates the block and vanishes in the process. This clears the counter stage so that the next Herschel can pass.

Lemma 5. *For any k we can build a stable pattern with diameter $O(\sqrt{k})$ that emits one Herschel for every 4^k Herschels that it absorbs.*

Proof. We need to build a Herschel path containing k quaternary counters. When these allow a Herschel to pass, they operate equivalently to 112-step turns, so the geometry is identical to building a conventional Herschel path. It is straightforward to build such a path in a snaking pattern such that its length grows asymptotically with the square of its diameter. □

Burst Reflectors

We can combine such a 4^k counter with a triggerable glider gun to obtain a 4^k-*burst reflector*. This absorbs one glider and emits a stream consisting of exactly 4^k gliders in succession. Such a device works as follows. We assume we have a triggerable glider gun that leaves its Herschel intact after recycling a glider into its input. We send this Herschel through a 4^k counter. The Herschel emitted by this counter eventually is further sent on such a path as to interfere with the recirculating glider and de-activate the gun. Thus the gun will output exactly 4^k gliders before emitting the Herschel to de-activate itself. Clearly, a lot of details are left to be resolved in practice. It was straightforward to construct a 16-burst reflector explicitly using two quaternary stages, and there does not seem to be any difficulty in extending this method.

It should be kept in mind that the delay through the counter stages grows at least linearly with k. A conservative analysis would require the period of gliders in a burst to increase linearly with k. Hickerson [1997] has observed that this is unnecessary. To see this, assume that the counter counts from 0 to $4^k - 1$, emitting a Herschel to de-activate the gun when its state changes to 0. By the time the gun is de-activated, it has inserted some number i of additional Herschels into the counter. Hence, the counter stabilizes at state i rather than 0. Regardless of the value of i, the burst reflector based on it (initialized to state i) must output exactly one glider for every state transition from i back to i again. This is exactly 4^k gliders.

In any such construction, we must be careful about timing the de-activation of the gun. We can eliminate most timing constraints by using the final Herschel to create a block to suppress the recirculating glider. We have to insure that a glider is not present when the block is being formed, but this can be accomplished using a sequence of conduits to realize an appropriate delay. Clearly, there are also lower limits on the period of a burst reflector, based on Herschel packing and the lowest attainable period for a triggerable glider gun. Combining all of the above, we have:

Lemma 6. *For sufficiently large q and for all $k \geq 0$ there is a 4^k-burst reflector of diameter $O(\sqrt{k})$ that absorbs one glider and emits a period-q burst of 4^k gliders.* □

Programmable Delay Reflector

The burst reflector is useful as a sort of system clock, but suffers from its dependence on powers of 4. It is still not clear how to use it to build an arbitrary-period oscillator. We now consider the component that allows us to adjust the period more or less arbitrarily up to a coarse approximation that we will later refine. The component is actually considerably simpler than the burst reflector, and also uses a ripple counter:

Lemma 7. *There is a constant c such that for any $k \geq 0$, and any 4^k-burst reflector, there exists an integer $p \leq ck$ such that for any m between 1 and 4^k there is a stable pattern that absorbs a burst of 4^k gliders from the reflector, emits a single glider p steps after absorbing the m-th glider in the burst, and returns to its initial state after receiving the last glider in the burst.*

Proof. The pattern converts each incoming glider into a Herschel and sends it through k quaternary counter stages. The initial state of this counter is set to adjust the step at which the glider is emitted. Because the burst consists of 4^k gliders, the state is reset to its initial value obliviously after the last glider is absorbed. $\square$

Pairing this component with a burst reflector, we can now build a small-diameter device that reflects a glider 180° after an adjustable delay. In other words:

Lemma 8. *There is a constant c such that, for sufficiently large q and for all $k \geq 0$, there exists $p \leq ck$ such that for any m between 1 and 4^k, there is a 180° stable glider reflector of diameter $O(\sqrt{k})$ that emits its reflected glider at step $p + mq$.* $\square$

We combine two of these reflectors to form an oscillator in the usual manner. However, we need to remember that regardless of the step at which a reflector emits its glider, it must still spend over $4^k q$ steps resetting the state of the adjustable component, and it is not ready to receive another glider during this time. There is no problem here as long as we insist that one of the reflectors is set to output its final glider as late as possible. This construction will allow us to realize every q-th period within some range. We can treat q as a constant, in practice the lowest period of a triggerable glider gun. This construction is depicted schematically in Figure 12. We summarize the set of oscillators realizable this way as follows.

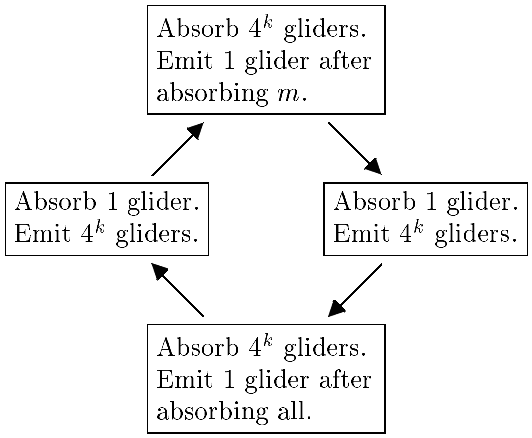

FIGURE 12. Schematic of an oscillator of period $(4^k{+}m)q + O(1)$.

Lemma 9. *For any n, there is an oscillator of period n' and diameter $O(\sqrt{\log n})$ such that $n' \leq n$ and $n - n' = O(1)$.* $\square$

This construction is a bit counterintuitive, so it is worth considering a cycle of this oscillator, setting m to various values in the first (adjustable) reflector. When m is set to 1, the first glider is reflected very soon after its collision. The reflected glider then enters the second reflector, set to maximum delay. These two reflectors continue to operate concurrently for most of the cycle. The first finishes shortly before the second. The second emits its glider and then finishes, repeating the cycle. As the initial state of the first reflector is adjusted, the number of steps of concurrent activity is reduced accordingly. When the first reflector is set to maximum delay $m = 4^k$, there is relatively little concurrency between the reflectors; the first finishes shortly after emitting the reflected glider.

Proof of Theorem 2. To complete the proof of Theorem 2 we need to be able to adjust the period n to any desired value. This is possible using a more conventional delay technique, because the total error is now constant. It is relatively straightforward to prove the following using a snaking pattern of stable reflectors of appropriate phase shifts.

Lemma 10. *For any sufficiently large n, there is a pattern with $O(\sqrt{n})$ diameter that absorbs a glider and later emits another glider equivalent to the first glider delayed by n steps.* □

So, to complete the construction, we first place a pair of reflectors to attain a period within some constant of that desired. Next, we add a delay component, as above, to the path of the recirculating glider, adjusting the period so it is exactly equal to n.

The construction of this section works for sufficiently large periods n. For lower periods, we can build an oscillator some other way, without affecting the asymptotic bounds of the theorem. This completes the proof of Theorem 2. □

8. CONCLUSION

Conway's announcement of the rules of Life and its subsequent popularization by Martin Gardner (collected in [Gardner 1983]) led to much interest in this elegant cellular automaton. Many fundamental questions were answered during an early burst of activity, much of which was carried out entirely by hand. Life remains a source of interesting open questions to this day. In the meantime, our understanding has been increased by over 25 years of careful study, and our tools include refined search techniques and greatly enhanced computing power.

The study of Life has moved beyond its early stages, in which a limited number of known interactions were combined in an ingenious fashion to prove the existence of patterns of interest. Often such patterns (notably, the as-yet-unrealized universal constructor) were far too unwieldy to build explicitly. Since that time, useful interactions and atomic patterns have been found in such abundance as to defy all attempts to catalog. The question of whether a certain kind of pattern can be built at all has been replaced in many instances by the puzzle of finding a particularly small or elegant example.

This shift in focus is demonstrated by the problem of realizing all oscillator periods and the related problem of finding a stable reflector. In this case, the early proof's reliance on the universal constructor would have made it hard to build even a very large instance of a stable reflector. By combining new interactions, however, we were able to exhibit a small, convincing example of such a pattern.

The search for Life patterns combines computer search, ad hoc puzzle solving, and a variety of analytical techniques that apply to well-behaved patterns such as oscillators, reflectors, and conduits. Life's rich structure makes it an ideal testbed both for combinatorial search and more interactive forms of computer-aided mathematics. We believe the techniques demonstrated here will lead to more methods for finding useful configurations in other CAs of interest. Future directions include the following questions:

Is there a smaller, faster stable reflector (or glider-to-Herschel conduit) — particularly one that does not require a multiple-stage Herschel track?

Can a feasible automated method be developed to enumerate all Herschel conduits within a given size range?

Can the gaps in known oscillator periods be filled, either by ad hoc constructions or a more general uniform construction?

Can similar constructions be realized in other CAs, particularly those not designed for this purpose?

Is there a CA for which nontrivial negative results hold for the oscillator or stable reflector problems?

Can the techniques used here be applied to CAs of more general interest, such as those motivated by physical systems?

ACKNOWLEDGEMENTS

We thank David Bell, Noam Elkies, Dean Hickerson, Dieter Leithner, Harold McIntosh, and others on the Life mailing list for their encouragement and useful comments on this paper.

ELECTRONIC AVAILABILITY

A machine-readable archive of all patterns discussed in this paper can be obtained by contacting Paul Callahan.

REFERENCES

[Bell and Hickerson 1994] D. Bell and D. Hickerson, "Life search", 1994. See ftp://ftp.germany.eu.net/pub/research/ci/Alife/packages/lifesearch/. ANSI C program to enumerate oscillators in Life by David Bell, based on an algorithm by Dean Hickerson.

[Berlekamp et al. 1982] E. R. Berlekamp, J. H. Conway, and R. K. Guy, *Winning ways for your mathematical plays*, Academic Press, 1982.

[Buckingham 1978] D. J. Buckingham, "Some facts of life", *Byte* **3**:12 (1978), 54–67.

[Buckingham 1996] D. J. Buckingham, "My experience with B-heptominos in oscillators", 1996. See http://www.cs.jhu.edu/~callahan/patterns/bhept/bhept.html.

[Callahan 1997] P. B. Callahan, "Conway's life miscellany web page", 1997. See http://www.cs.jhu.edu/~callahan/lifepage.html. Patterns and mathematical results on Life.

[Gardner 1983] M. Gardner, *Wheels, life, and other mathematical amusements*, W. H. Freeman, New York, 1983.

[Hensel 1995] A. Hensel, "Life pattern archive", 1995. See http://www.mindspring.com/~alanh/lifep.zip. A collection of Life patterns and descriptions.

[Hickerson 1997] D. Hickerson, 1997. Personal communication.

[Leithner 1997] D. Leithner, 1997. Personal communication.

[McIntosh 1991] H. V. McIntosh, "Linear cellular automata via de Bruijn diagrams", 1991. See http://www.cs.cinvestav.mx/mcintosh/cf/debruijn.html.

[Moore 1962] E. F. Moore, "Machine models of self reproduction", pp. 17–33 in *Mathematical problems in the biological sciences*, Proceedings of Symposia in Applied Mathematics **14**, Amer. Math. Soc., Providence, 1962.

[von Neumann 1966] J. von Neumann, *Theory of self-reproducing automata*, University of Illinois Press, Urbana, Illinois, 1966.

[Poundstone 1985] W. Poundstone, *The recursive universe*, William Morrow, New York, 1985.

[Schroeppel 1994] R. Schroeppel, 1994. Personal communication.

[Silver 1997] S. Silver, 1997. Personal communication.

[Wakerly 1994] J. F. Wakerly, *Prentice-Hall*, Englewood Cliffs, NJ, 1994.

[Wolfram 1984] S. Wolfram, "Universality and complexity in cellular automata", *Phys. D* **10**:1-2 (1984), 1–35.

David J. Buckingham, David Buckingham Consulting, 174 Louisa Street, Kitchener, Ontario N2H 5M5, Canada (Dave.Buckingham@sigmastar.com)

Paul B. Callahan, Hyseq, Inc., 670 Almanor Avenue, Sunnyvale, CA 94086 (callahanp@acm.org)

Received June 20, 1997; accepted in revised form December 3, 1997