

Penrose Life: ash and oscillators

Margaret Hill¹, Susan Stepney¹, and Francis Wan²

¹ Department of Computer Science, University of York,
Heslington, York, YO10 5DD, UK.

susan@cs.york.ac.uk

² Ampleforth College, York, YO62 4ER, UK.

Abstract. We compare the long term behaviour of Conway’s Game of Life cellular automaton, from initial random configurations, on a bounded rectangular grid and a bounded Penrose tiling grid. We investigate the lifetime to stability, the final ‘ash’ density, and the number and period of final oscillators. Penrose grids have similar qualitative behaviour but different quantitative behaviour, with shorter lifetimes, lower ash densities, and higher occurrence of long-period oscillators.

Keywords: Conway’s Game of Life; Penrose tiles; ash; oscillators

1 Introduction

John Horton Conway’s *Game of Life* [1][3] is a simple two-dimensional, two state cellular automaton (CA), remarkable for its complex behaviour [1][8]. That behaviour is known to be very sensitive to a change in the CA rules. Here we investigate its sensitivity to changes in the grid, by the use of an aperiodic Penrose tiling grid [4][7].

2 Varieties of Life

In Conway’s *Game of Life* CA, the neighbourhood of each cell comprises the 8 nearest cells of the Moore neighbourhood. Each cell has two states, ‘dead’ and ‘alive’. If a cell is alive at time t , then it stays alive iff it has 2 or 3 live neighbours (otherwise it dies of ‘loneliness’ or ‘overcrowding’). If a cell is dead at time t , then it becomes alive (is ‘born’) iff it has exactly 3 live neighbours.

Life has grown its own extensive and idiosyncratic terminology over the years. Much of this is collected in the extensive on-line *Life Lexicon* [9]. In particular, an initial random starting state is called a *soup*, and, mixing metaphors somewhat, the final resulting configuration is called the *ash*.

We can run CAs such as Life on aperiodic grids, such as Penrose tilings, with a suitable definition of the ‘neighbourhood’.

There are two classic sets of Penrose tiles [4][7], *kites* and *darts* (so called because of their shapes) and fat and thin rhombuses. We use the kite and dart form. A plain kite and dart can be combined into a rhombus, and so tile the plane periodically. To force the tiling to be aperiodic, *matching rules*, marks on the tiles that must be matched together, are used.



Fig. 1. deflating a kite and a dart (a) minimal deflation; (b) deflation avoiding ‘holes’ in the result

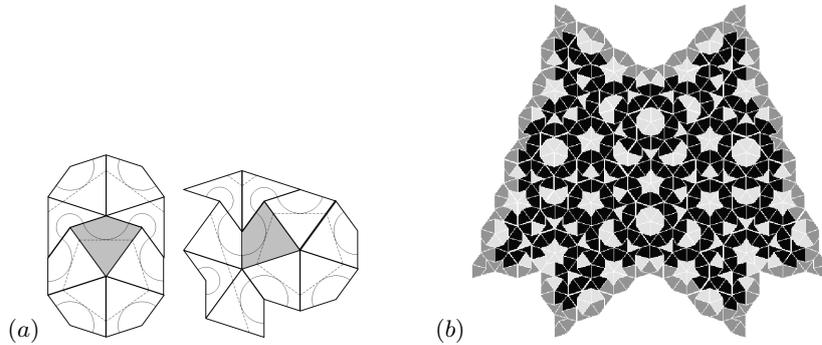


Fig. 2. variable Penrose neighbourhood: (a) example of kites with eight and nine neighbours; (b) numbers of neighbours in a small deflated grid: light coloured tiles have 8 neighbours; dark tiles have 9 neighbours; grey tiles are edge tiles with fewer than 8 neighbours.

A valid Penrose tiling has no gaps or overlapping tiles. The deflation algorithm [6] guarantees a valid tiling. At each round of deflation, each kite and dart tile is replaced with smaller kites and darts (figure 1). This leads to overlapping tiles, but the overlap is exact, and so the extra tiles can be safely removed. Since we use this deflation algorithm, we are restricted to the sizes (number of cells) of Penrose grids produced by the successive deflation generations.

In a rectangular grid, four cells meet at every vertex, and every cell has eight neighbours. In a Penrose grid, three, four or five cells can meet at a vertex, and Penrose grid cells can have either eight or nine neighbours (figure 2). We have found no algorithmic way of reducing the neighbourhood of all tiles to eight whilst maintaining the undirected nature of the neighbourhood graph. So we leave the neighbourhood as it is, and apply the Life rules to it unchanged.

We need to cope with the edge of the deflated Penrose grid. There are two conventional ways in CAs of removing the effect of the edge of the grid.

1. periodic boundary conditions: the grid has the topology of a torus, finite but unbounded (has no edges). This is the approach usually taken for investigating statistical properties of soups, with the results more or less tentatively extrapolated to infinite grids. However, this approach is impossible for aperiodic grids such as a Penrose grid.

2. lazy infinite grid: implemented by lazily expanding a finite grid as activity nears its edges [5]. This is the approach usually taken for investigating the properties of particular structures, such as glider guns. It is not practical for

| | S, small | M, medium | L, large | X, extra-large |
|-------------|--------------|---------------|---------------|-----------------|
| Penrose | 688 | 1907 | 5170 | 13900 |
| rectangular | $676 = 26^2$ | $1936 = 44^2$ | $5184 = 72^2$ | $13924 = 118^2$ |

Table 1. The four grid sizes investigated: four Penrose deflations and the corresponding nearest regular square grid size

implementation on Penrose grids produced by deflation, since the generation n grid does not clearly appear as a subpart of the larger generation $n + 1$ grid.

Since neither of these standard approaches is suitable for investigating Penrose soups, we choose to investigate the effects of having a *bounded* grid, explicitly noting the effect of the edges. We have to decide how to handle the boundary. We can choose the border cells to stay ‘dead’, no matter what their neighbours’ states, or choose them to have a reduced neighbourhood of five (or three at the corners). These choices are equivalent for CA rules like those of Life, where the state transition depends only on the *total* number of live neighbours (so permanently dead neighbours are equivalent to no neighbours).

3 Experimental set-up

For both the regular and Penrose grids, to compare like with like, we investigate the behaviour of a finite grid, initially empty except for a smaller patch of soup. Given that we are restricted to certain Penrose grid sizes by the deflation algorithm, we restrict the rectangular grid to the nearest similar sizes (table 1).

So our investigations are parameterised by the initial soup patch size S , initial soup density D , and fixed grid size G . The questions we pose are:

- what are the lifetimes of the initial random configurations?
- what are the final ash densities?
- what are the periods of the oscillators at the end of the lifetime?

A pattern is called a period n (pn) oscillator if it repeats after n generations. We define the *lifetime* of a configuration as the number of generations from the initial random starting configuration until it stabilises to ash, where a *stable* state comprises only oscillators, including ‘Still Life’ $p1$ oscillators. So we wait until any gliders have been absorbed by the boundary.

For example, on an infinite grid, the well-known ‘r-pentomino’ initial configuration reaches a stable state after 1103 generations, when it comprises several Still Lifes and $p2$ *blinkers*, and six escaped gliders. On a finite but large enough grid it stabilises to ash once the six gliders reach the boundary.

4 Lifetime results

We use 1000 runs with each parameter set, but different random starting configurations.

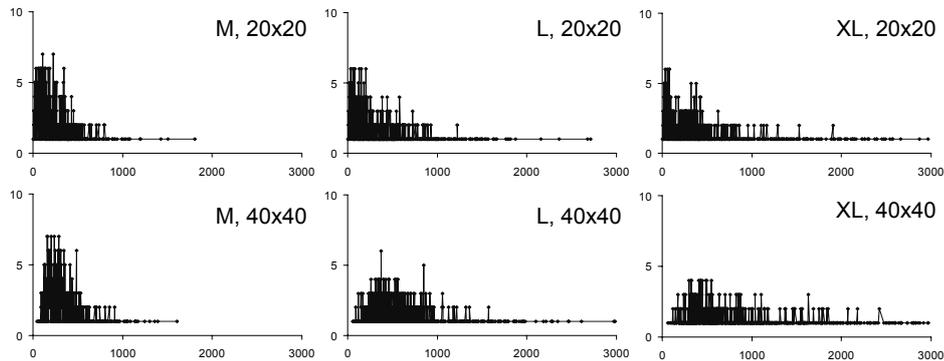


Fig. 3. Frequency against lifetime to stability on a regular grid, for initial density $D = 20\%$. Columns are increasing grid sizes G , showing lengthening lifetime tails; rows are increasing soup patch sizes S , peaks move to longer lifetimes.

4.1 Classic Life

Previous results Achim Flammenkamp has an extensive list of ash objects grown from soup [2]. He starts with initial densities $D = 0.371 - 0.375$, on toroidal (periodic boundary condition) grids of sizes $2^{12} \times 2^{12}$ (4096×4096) and $2^{14} \times 2^{14}$ (16384×16384). These experiments give an asymptotic ash density of 0.0287115 bits per cell. There are no reports of the time taken to stabilise.

The grids we test here are much smaller (the largest grid we test is X = 118×188), but we test a much broader range of grid sizes and initial densities.

Distribution The distributions of lifetimes are highly skewed (figure 3). There is a peak in the distributions at low lifetimes, with a long tail of high lifetimes. This tail is longer on larger grids, that is, larger grids can support longer lifetime structures, implying that there is a correlation between a structure's lifetime and its size. The peak lifetime is higher with larger initial patches: small patches tend to die more rapidly than larger patches.

Effect of initial density We expect a short lifetime at low density, because there are too few live cells to cause more to be born. We might also expect short lifetimes at high density, as everything dies of overcrowding.

We can see these effects qualitatively from the following approximate argument. (It is only an approximation to the real behaviour, because we treat the distribution as smooth, yet clumpiness has an important effect; however, the intuition it provides is sound.) Let the average density at time t be ρ_t . Then the average density at time $t + 1$ will be (approximately):

$$\rho_{t+1} = \rho_t P(\text{staying alive}) + (1 - \rho_t) P(\text{being born}) \quad (1)$$

$$= \rho_t (P(2 \text{ nbrs}) + P(3 \text{ nbrs})) + (1 - \rho_t) P(3 \text{ nbrs}) \quad (2)$$

$$= \rho_t ({}_nC_2 \rho_t^2 (1 - \rho_t)^{n-2}) + {}_nC_3 \rho_t^3 (1 - \rho_t)^{n-3} \quad (3)$$

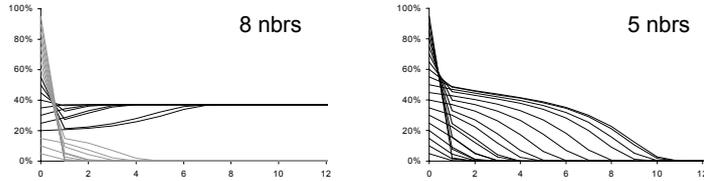


Fig. 4. Density against timestep for smoothed evolution of initial densities

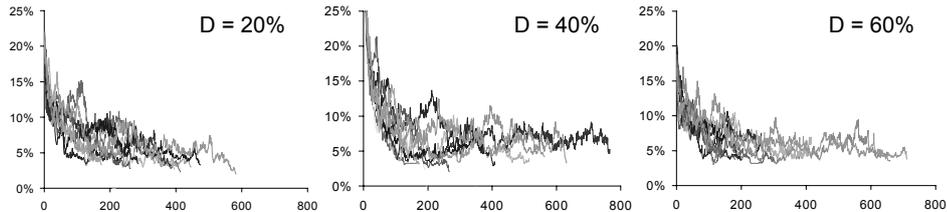


Fig. 5. Density against time on an M grid, of 10 initial random densities D

where $P(m \text{ nbrs})$ is the probability that m neighbours are alive. For a neighbourhood of $n = 8$, we get

$$\rho_{t+1} = 28\rho_t^3(1 - \rho_t)^5(3 - \rho_t) \quad (4)$$

For an initial ρ that is not too large or too small, the density rapidly converges to $\rho_\infty \approx 37\%$ (figure 4a). However, for low initial densities ($\rho_0 \lesssim 20\%$), there is not enough activity to sustain Life, and the density rapidly falls to zero. For high initial densities ($60\% \lesssim \rho_0$), there is massive death in the first generation, and the resulting density ρ_1 is less than the critical value, and so again rapidly converges to zero.

Similarly, for a neighbourhood of $n = 9$ (relevant to some cells in the Penrose grid), we get

$$\rho_{t+1} = 12\rho_t^3(1 - \rho_t)^6(10 - 3\rho_t) \quad (5)$$

For an initial ρ not too large or too small, the density rapidly converges to $\rho_\infty \approx 35\%$.

We compare these calculations with actual runs, on an M grid (44×44 cells), with initial densities of 20%, 40%, and 60% initially covering the entire grid, for 10 runs each. The evolutions of the densities are shown in figure 5. We see that the evolution is qualitatively similar, but the actual densities are lower than the calculation, at closer to 10% whilst still evolving, and 2–5% once stabilised to ash. (This latter figure is consistent with Flammenkamp’s asymptotic value of $\approx 2.9\%$.) The lower densities demonstrate the importance of clumpiness. There is no correlation between lifetime to stability, and ash density (figure 8a).

Figure 6 shows graphs of mean lifetime against density, for various grid sizes and initial patch sizes. At low densities, they have low lifetimes, and as the

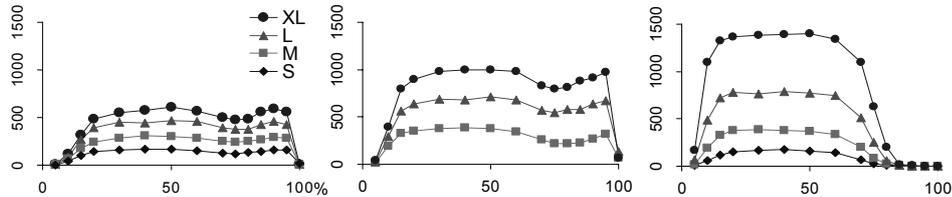


Fig. 6. Mean lifetimes to stability against initial density, on a regular grid, for different soup patch sizes $S = 20 \times 20$, 40×40 , full grid

density rises, so does the lifetime. Once the density gets too high, lifetimes start to drop again. But then, for very high densities, the mean lifetime starts to rise once more (except when the initial patch completely fills the grid). Why?

These experiments are run on patches of initial random soup that are smaller than the total grid size. For very high densities, this initial patch is essentially a solid square. Note that the average density at the edge of a patch is half the interior density. So for solid patches, their edges are at a density suitable for sustaining Life. The centre rapidly dies, but the edges survive and propagate for a long time. So the entire graph can be thought of as having two components: one due to the central region, peaking near $D_{max} \approx 50\%$, and one due to the edges, peaking at $\approx 2D_{max}$.

Detailed effect of patch and grid size The earlier figure 3 shows that both average and maximum lifetimes increase with initial soup patch size S , and grid size G .

The bigger the grid, the longer the lifetime: the boundary does seem to be ‘killing’ the life. Clearly, if some central region shoots out gliders, then bigger grids will give longer lifetimes, because the lifetime is taken once all the gliders have hit the boundary, which will take longer for larger grids. But ours are all relatively small grids, and that is not the dominant effect: the soup is ‘boiling’ over the whole grid.

We can see that 5 neighbours (the case on the boundary) is not enough to maintain Life, using equation 3 for $n = 5$ neighbours. We get

$$\rho_{t+1} = 10\rho_t^3(1 - \rho_t)^2(2 - \rho_t) \quad (6)$$

In this case, whatever the initial value of ρ , it quickly converges to zero (figure 4b): all Life dies.

4.2 Penrose grid lifetimes

Now that we understand the effect of grid size, patch size, and initial density on the regular grid, we can investigate the effect of using an aperiodic Penrose grid.

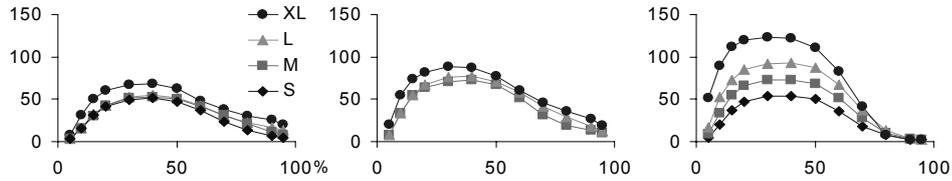


Fig. 7. Mean lifetimes to stability against initial density, on a Penrose grid, for different soup patch sizes $A = 400, 16040$, full grid

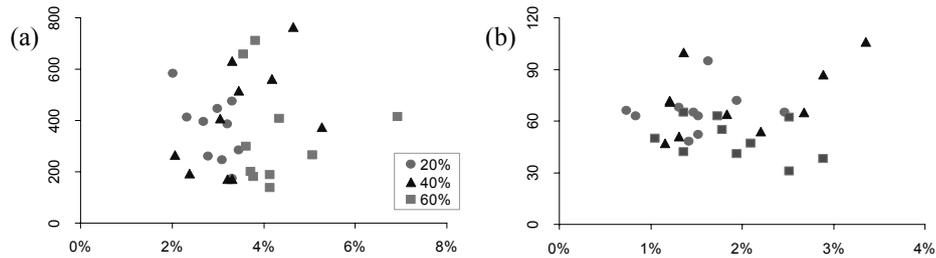


Fig. 8. lifetime to stability against ash density, for initial densities $D = 20\%, 40\%, 60\%$, on M size grids: (a) rectangular grid (b) Penrose grid

Distribution The lifetime distributions show the same qualitative behaviour as for a regular grid: the skewed distribution, and the increasing lifetimes with grid and patch size. However, the mean lifetimes to reach stability are approximately an order of magnitude *smaller* (figure 7).

The densities drop much faster than they do on the rectangular grid, last for a much shorter time at the low density before stabilising, and produce lower density ash, 1–3% (figure 8b).

5 Oscillator distribution

It appears that the aperiodic grid stops structures propagating any distance and affecting distant objects. So everything becomes ‘Still P-Life’, or an oscillator, much sooner. We investigate this further, by looking at the distribution of oscillator periods in the ash.

5.1 Life oscillators

The Flammenkamp web site [2] has an extensive list of oscillators grown from soup. The vast majority are $p2$. Why do we find so very few $p3$ and higher oscillators in the ash?

There are many kinds of $p2$ oscillators. The *blinker* has only three active cells; there are also three 6 cell $p2$ oscillators (*beacon*, *clock*, and *toad*). These

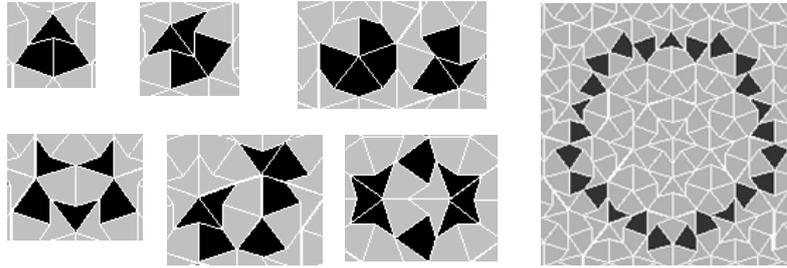


Fig. 9. Examples of Penrose still lifes. All the small examples were found in the ash; the large ring was hand constructed.

readily form by chance in the ash. Similarly, the commonly-occurring *glider* has only five active cells.

The smallest $p3$ oscillator, the *caterer* (discovered by Dean Hickerson in 1989), however, has 12 cells active in its smallest configuration. This is much less likely to occur by chance than smaller period oscillators, and so is unlikely to be found in the ash. Our tests never discovered a caterer oscillator.

The $p3$ oscillator that does occur in our tests is the *pulsar*, which has 24 cells active in its smallest configuration. Surely this is even *less* likely? However, the pulsar has a 10 state predecessor, which is more likely to occur than the caterer.

The smallest $p4$ oscillators (*mazing* and *mold*) have 12 cells active in their smallest configuration, and the smallest $p6$ oscillator (*unix*) has 16 cells, so these are also unlikely to occur by chance, unless they have small state progenitors, too. They never occurred in our tests.

5.2 Penrose oscillators

Although the Penrose grid is aperiodic, any Penrose oscillator is of general interest, because a Penrose tiling has the *recurrence property*: any given finite patch of Penrose tiling recurs in infinitely many other patches, in any Penrose tiling. (This does not conflict with what we said earlier about being unable to find the deflated generation n grid within the generation $n + 1$ grid: these grids are finite, whereas the recurrence property applies to full, infinite, tilings.) Hence a given oscillator confined to a particular patch of a given tiling can also occur in infinitely many other patches, in any Penrose tiling.

This property also means that we are justified in confining our experiments to just those Penrose grids obtained by deflating a single dart: any finite patch that occurs in any Penrose tiling will occur in our grids (provided that they are big enough, of course).

We have discovered a rich zoo of small-period P-Life oscillators in the ash. There are many $p1$ still lifes (a few examples are shown in figure 9), including one with three cells, some forming closed loops, and some disconnected.

Because of the two different cell shapes, and the different neighbourhood sizes, many oscillators come in several variants. For example, the ash contained

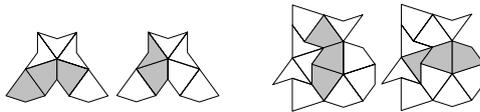


Fig. 10. Two of the $p2$ Penrose oscillators, named *plinkers* by analogy to the Life three cell *blinker*

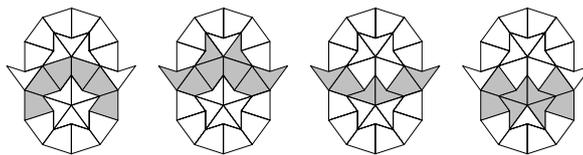


Fig. 11. A 6 cell $p4$ Penrose oscillator, the *bat*

six different three-cell $p2$ oscillators, analogous to the single *blinker* (figure 10). Some of these visually distinct variants look more similar when considering just the neighbourhood topology, rather than the different tile shapes.

There is a 6 cell $p4$ oscillator that we dub the *bat* (figure 11); the ash exhibited four variant bats. The ash also threw up a symmetrical 8 cell $p8$ oscillator (figure 12).

Perhaps most amazingly, we discovered an 8 cell $p15$ oscillator (figure 13). When animated it appears that the live cells are ‘dancing’ around the central star formation. This movement is reminiscent of the behaviour of a glider on the regular grid: the $p15$ *dancer* is moving, but is confined to perpetually move in a circle because of the nature of the Penrose grid. We found no long range propagating structures, because of the aperiodic nature of the grid. However, the possibility of arbitrarily large rings (figure 9) hints at the possibility of arbitrarily large period dancers around such rings: such will almost certainly need to be hand constructed.

These relatively long period Penrose oscillators are more common than their regular Life equivalents, because they are so small. This still raises the question: why are smaller oscillators possible? The occasionally larger neighbourhood is a possible contributor. It certainly allows constructs such as arbitrarily large rings. Future work will investigate oscillator distributions, and the precise effect of the extra neighbour.

6 Conclusions

Life on a Penrose grid has similar qualitative behaviour to regular Life, but different quantitative behaviour. The lifetime to stability is an order of magnitude shorter, the ash density is about half, and there are more spatially small long-period oscillators.

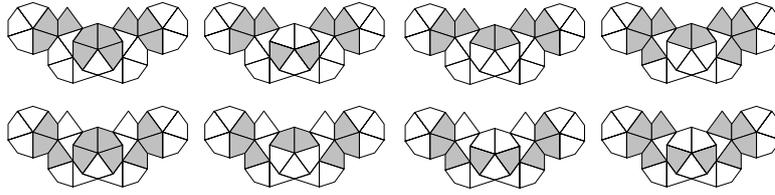


Fig. 12. An 8 cell $p8$ Penrose oscillator (read across the rows, then down the columns)

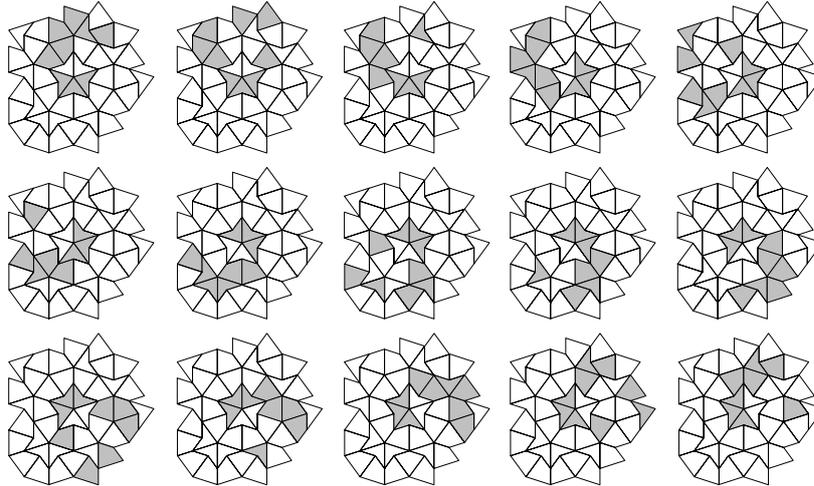


Fig. 13. An 8 cell $p15$ Penrose oscillator, the *dancer* (read across the rows, then down the columns)

References

1. Elwyn R. Berlekamp, John Horton Conway, and Richard K. Guy. *Winning Ways for Your Mathematical Plays Volume 2: games in particular*. Academic Press, 1982.
2. Achim Flammenkamp. Achim's game of life page. <http://wwwhomes.uni-bielefeld.de/achim/gol.html>, 2004.
3. Martin Gardner. Mathematical games: The fantastic combinations of John Conway's new solitaire game "life". *Scientific American*, 223(4):120–123, October 1970.
4. Martin Gardner. Mathematical games: extraordinary non-periodic tiling that enriches the theory of tiles. *Scientific American*, 236(1):110–121, January 1977.
5. R. Wm. Gosper. Exploiting regularities in large cellular spaces. *Physica D*, 10:75–80, 1984.
6. B. Grünbaum and G. C. Shephard. *Tilings and Patterns*. W. H. Freeman, 1987.
7. Roger Penrose. Pentaplexity. *Eureka*, 39:16–32, 1978.
8. Paul Rendell. Turing Universality of the Game of Life. In Andrew Adamatzky, editor, *Collision-Based Computing*. Springer, 2002.
9. Steven Silver. Life lexicon, release 24. <http://www.argentum.freemove.co.uk/lex.htm>, February 2005.