Life Without Death is P-complete

David Griffeath[1] and Cristopher Moore[2]

[1] Mathematics Department, University of Wisconsin, Madison 53706
e-mail: griffeat@math.wisc.edu
[2] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501
e-mail: moore@santafe.edu

Abstract. We show that if a cellular automaton in two or more dimensions supports growing "ladders" which can turn or block each other, then it can express arbitrary Boolean circuits. Then the problem of predicting the CA for a finite amount of time becomes **P**-complete, the question of whether a finite configuration grows to infinity is **P**-hard, and the long-term behavior of initial conditions with a periodic background is undecidable.

This class includes the "Life without Death" rule, in which cells turn on if exactly three of their neighbors are on, and never turn off.

1 Introduction

Given the initial conditions of a d-dimensional cellular automaton, suppose we want to know the state at a site t time-steps in the future. We can do this in $\mathcal{O}(t^{d+1})$ steps on a serial computer, or $\mathcal{O}(t)$ steps on a parallel one, simply by simulating the CA explicitly and filling in the light-cone above the site in question. Thus CA Prediction is in the class **P** of problems solvable by a deterministic Turing machine in polynomial time [17].

A problem is *hard* for a complexity class if all other problems in that class can be reduced to it, and *complete* if it is hard and in the class itself. If a complete problem could be solved quickly, so could all other problems in its class. Just as **NP**-complete problems are believed to require a super-polynomial amount of search, **P**-complete problems are believed to be *inherently sequential*, so that the work needs to be done in step-by-step order and cannot be efficiently parallelized [6].

Non-linear CA's with certain algebraic properties [11, 12] can be predicted in $\mathcal{O}(\log t)$ or $\mathcal{O}(\log^2 t)$, placing them in the parallel complexity class **NC** of efficiently parallelizable problems. But CA Prediction is **P**-complete in general, since CA's exist that can perform universal computation [4, 9]. A number of other cellular automata and lattice systems have been shown to be **P**-complete, including diffusion-limited aggregation, single spin-flip Ising dynamics, fluid invasion, majority-voting CA's, lattice gases, and sandpiles [10, 13–15], variously in $d \geq 2$ or $d \geq 3$.

These systems are **P**-complete because the following problem can be reduced to them: given a Boolean circuit and the truth values of its inputs, is its output true or false? This CIRCUIT VALUE problem seems inherently sequential, since it is hard to imagine how one could calculate the output of an arbitrary circuit without going through it level-by-level.

The proofs in [10, 13–15] work by showing that we can build “wires” to carry truth values, and AND and OR gates to connect them, so that the future state of a particular site corresponds to the output of the circuit. In this paper, we do the same for a class of CA’s where one-dimensional structures called “ladders” can grow, turn, and block each other. This includes the “Life Without Death” rule, in which a cell turns on if exactly three of its neighbors are on (as in Conway’s Life), but then never turns off again.

Thus the prediction problem for such CA’s is **P**-complete. As corollaries to this, we find that the question of whether a finite initial seed grows to infinity is **P**-hard, and whether an initial condition with a periodic background evolves to a fixed point or not is undecidable.

2 Ladders

We will call a *ladder* a one-dimensional periodic structure, which grows in one direction in a straight horizontal or vertical line. Then define a *ladder CA* as one in which:

- There are finite seeds which give birth to horizontal or vertical ladders and nothing else.
- Ladders can be turned: there are finite structures with which they can collide, producing a new ladder 90° from the old one and nothing else.
- One ladder can block another: a horizontal ladder (say) can collide with a vertical one which is already there in such a way that it stops without generating anything else.

We assume that the CA rule is rotationally symmetric, so that ladders and turns work in all four lattice directions. If these collisions are sensitive to the ladders’ spatial or temporal phase, then we also require the following:

- Ladders can be shifted or delayed so that their phase can be adjusted as desired.

Two more behaviors follow from these requirements, as shown in figure 1:

- Ladders can be delayed by arbitrary times Δt, by executing a sequence of turns in a square of side $\mathcal{O}(\Delta t^{1/2})$.
- Ladders can end.

The blocking collision will be our main building block for Boolean circuits. If the presence or absence of a ladder corresponds to a wire being true or false, blocking acts as a logical gate as shown in figure 2, with inputs a and b and outputs b and $a \wedge \overline{b}$. (If necessary, we delay a so that it arrives after b does if b is present.)

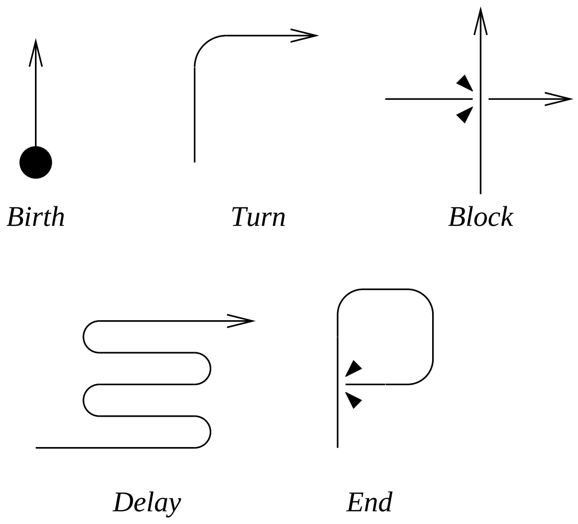

Fig. 1. Birth, turns, and blocking. With sequences of turns, we can also delay ladders, or end them by blocking them with themselves.

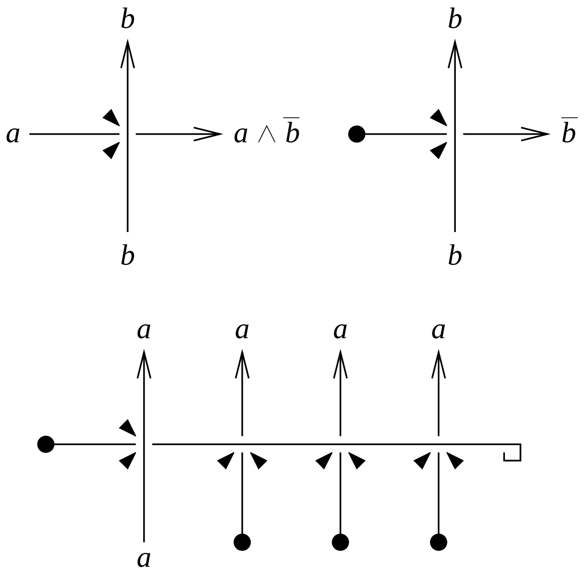

Fig. 2. The logical effect of a blocking collision. A ladder emerges on the right only if a is present and b is not, so that it corresponds to the truth value $a \wedge \overline{b}$. If a is always present, blocking negates b. We can also make multiple copies of a wire.

By setting a to TRUE, i.e. giving birth to a ladder which is always there, we can use blocking to negate wires. Then by combining blocking with negation, we can construct AND and OR gates, $a \wedge b = a \wedge \overline{\overline{b}}$ and $a \vee b = \overline{\overline{a} \wedge \overline{b}}$. We can also make a wire 'fan out' by making multiple copies of it, so that it can be used as input for more than one gate.

In two dimensions, since our ladders cannot cross each other, our circuits have to be planar. But if negation is available, we can simulate non-planar circuits by building a crossover gadget as shown in figure 3; the reader can check that with inputs a on the left and b on the right, the outputs are b on the left and a on the right as advertised. (In contrast, the majority-vote CA cannot express negation, and so can only express planar circuits in $d = 2$ [13].)

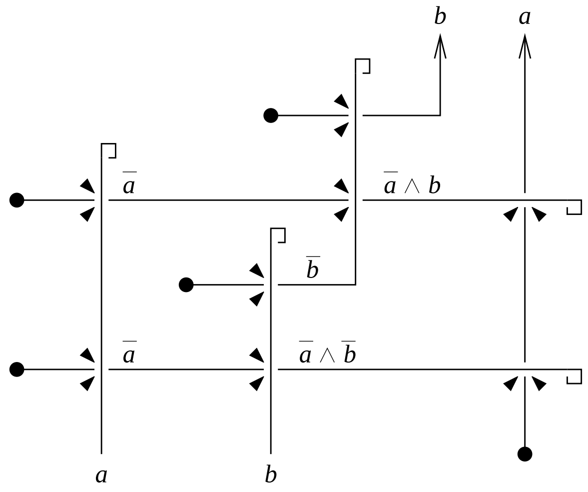

Fig. 3. A crossover gadget with 8 blocking gates; perhaps a clever reader can find a smaller one.

Because of this ability to do crossover, CIRCUIT VALUE for planar circuits is just as hard as the general case, and so is **P**-complete. So to complete our proof, we just need to verify that there is a polynomial relationship between the size of a planar circuit and the size and duration of the corresponding CA evolution.

Suppose our circuit has depth D and width W, i.e. it has D levels, each of which has at most W gates. We wish to simulate it with t time-steps of the CA in a rectangle of height h and width w (in pixels). We have to delay the births on the last level by $\mathcal{O}(t)$ so that they arrive at the blocking gates after the wires do, and smaller delays can synchronize the wires so that they arrive in the proper order.

These delays can be made as in figure 1, by folding sequences of turns into square 'egg timers' of side $\mathcal{O}(t^{1/2})$; each level then has width $w = \mathcal{O}(Wt^{1/2})$ and height $\Delta h = h(D+1) - h(D) = \mathcal{O}(t^{1/2})$. In the worst case wires at each level have to travel a horizontal distance $\mathcal{O}(w)$ to meet each other, as well as $\mathcal{O}(\Delta h)$ to get through the level; for wide circuits $\mathcal{O}(w)$ dominates, and we have

$$t(D+1) \sim t(D) + w(D)$$
$$h(D+1) \sim h(D) + t(D)^{1/2}$$
$$w(D) \sim W t(D)^{1/2}$$

giving the leading behavior for large D and W

$$t(D) \sim (WD)^2$$
$$h(D) \sim WD^2$$
$$w(D) \sim W^2 D$$

so that the time and area of the CA grow as $t \sim n^2$ and $hw \sim n^3$ respectively for circuits with $WD = n$ gates. This is rather conservative; in a 'nice' planar circuit, such as one simulating a one-dimensional cellular automaton, wires only have to travel a distance comparable to that of a single gate with egg timers to meet their neighbors, so $t(D+1) \sim t(D) + t(D)^{1/2}$ and we get a leading behavior of

$$t(D) \sim D^2$$
$$h(D) \sim D^2$$
$$w(D) \sim WD$$

which for a square circuit in which $W = D = n^{1/2}$ gives time $t \sim n$ and area $hw \sim n^2$.

In any case, the relationship is polynomial, and we have proved:

Proposition 1. FINITE-TIME CA PREDICTION *is* **P**-*complete for ladder CA's.*

Suppose we want to know whether a given finite seed will grow to infinity. If we let the circuit's output ladder grow and kill off all the others, a circuit will grow to infinity if and only if its output is true. Thus we have the corollary:

Corollary 2. INFINITE GROWTH FROM FINITE SEED *is* **P**-*hard for ladder CA's.*

This problem would be **P**-complete if it were in **P**. However, if there are structures that allow new ladders to be created (such as a 'ladder gun' analogous to the glider gun in Conway's Life) it seems more likely that it is undecidable.

We can also set up an infinite circuit, with a finite seed surrounded by a periodic initial background. Such a circuit can simulate a computationally universal one-dimensional CA that simulates a Turing machine (e.g. [9]). If the Turing machine head disappears upon halting, then the ladder CA ends in a fixed point if and only if the Turing machine halts. Thus

Corollary 3. *The set of finite seeds and periodic backgrounds on an infinite lattice for which a ladder CA will end in a fixed point is recursively enumerable but not recursive, and so* LONG-TERM CA PREDICTION WITH PERIODIC BACKGROUND *is undecidable for ladder CA's.*

3 Life without Death

Let us now show that the ladder CA design principles of the previous section can be implemented using one of the simplest imaginable two-dimensional rules: Life without Death (LwoD).[1]

In this growth model an empty site becomes permanently occupied if 3 of its 8 nearest neighbors in the square lattice are occupied. The remarkable behavior of LwoD is described in several recipes from the Primordial Soup Kitchen [7]; essential features have been noted over the years by Packard and Wolfram [16] and various members of the Life List Internet group devoted to Conway's Game of Life.

From suitable initial seeds one encounters complex dendritic crystal patterns. The dynamics evolve as a pseudo-stochastic mix of three ingredients: (i) chaotic lava, (ii) horizontal and vertical ladders which advance at speed $1/3$ by a weaving motion and seem to outrun the surrounding lava, and (iii) parasitic chutes which emerge from the lava but can only shoot along the edges of ladders at speed $2/3$. The resulting interactions — for instance, when a chute reaches the end of a ladder it nips it in the bud, and causes a lava eruption — give rise to remarkable self-organization in which very large regions of the growing crystal can be cloned exactly in time. As one indication of LwoD's sensitive dependence on initial conditions, given any finite initial configuration A, one can produce configurations B and C, each consisting of at most 100 cells, such that $A \cup B$ grows forever, but $A \cup C$ reaches a fixed point [5]. Apart from minor variants, we know of no other simple CA rule with this property.

Our objective here is to harness some of the complex phenomenology of LwoD in order to construct the defining elementary collisions of a ladder CA: births, turns, blocks and phase shifts. In figures 4–9 below we present a series of still frames which document the necessary interactions. We invite the reader to download WinCA [3], or any of several other interactive 2d CA simulators [8], to see the dynamics in action and confirm the constructions.

To begin, figure 4 shows one seed of size 8 which gives birth to a ladder. (The pentomino piece actually spawns a two-sided ladder, and then interaction with the other three cells kills its left branch.) Note that the ladder's horizontal spatial periodicity is 4. In fact, the head cycles through 12 distinct patterns in order to weave its periodic pattern, the last 6 of which are vertical flips of the first 6.

Recall that a separate killing mechanism is not needed in a ladder CA since ladders can block themselves after 3 turns. Nevertheless, as further evidence of LwoD's sensitive dependence, figure 5 shows that its ladders can be stopped by a single well-placed cell!

Next, figures 6 and 7 show two static configurations which cause an LwoD ladder to execute a 90° turn to the right, while otherwise causing only a local disturbance. The turn in figure 6 is instructive, since it involves typical instabilities which must be controlled after a collision. The actual ladder turn is caused

[1] Pronounced "El-wod." It is related to the Hubbard model in physics.

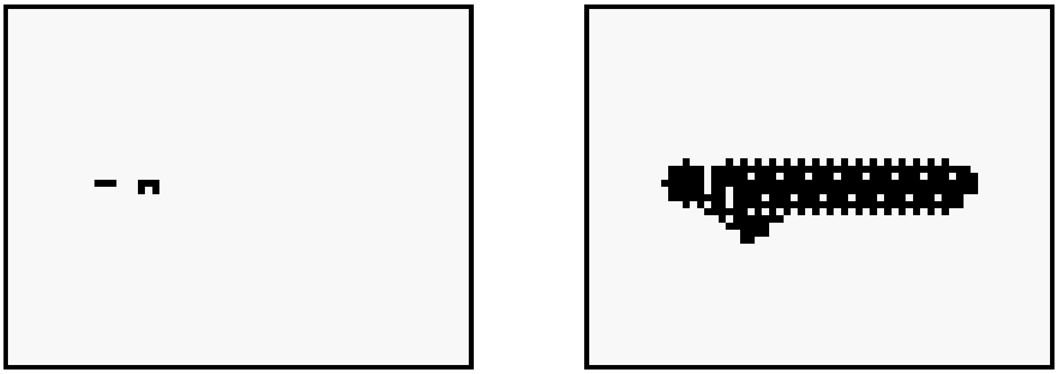

Fig. 4. Birth of a ladder from an initial configuration of 8 cells.

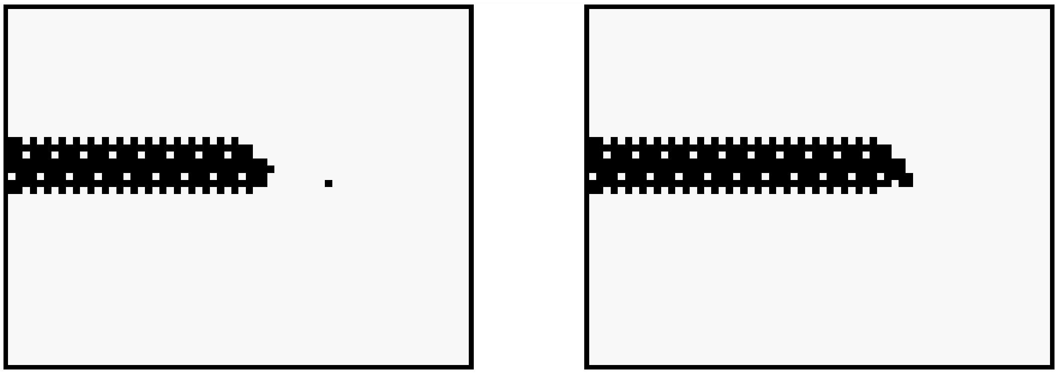

Fig. 5. A ladder killed by a single cell.

by the single isolated cell, but without further engineering another ladder would turn to the left and a chute would race back along the lower edge of the incoming horizontal ladder. The 29-cell blob blocks the former effect, and the row of 3 cells stops the latter while causing only a small transient lava flow. Figure 7 shows a turn which is both more elegant and mysterious: four carefully placed cells induce a brief burst of lava which leads to a right ladder turn without any undesirable side-effects.

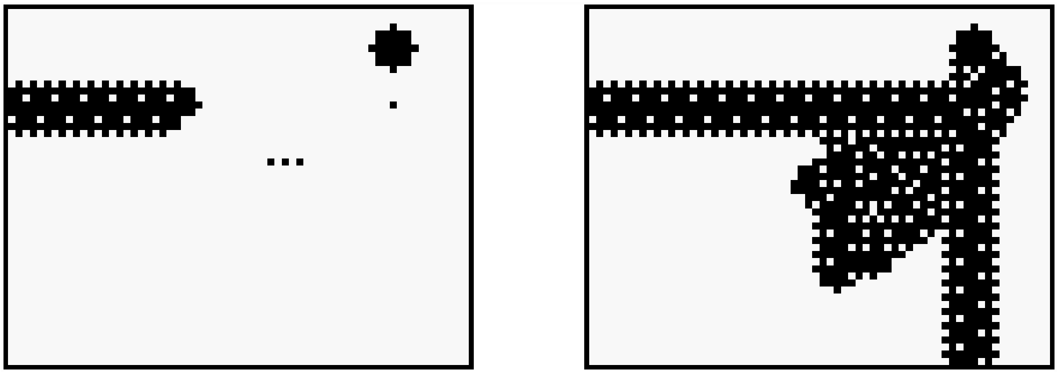

Fig. 6. A turn.

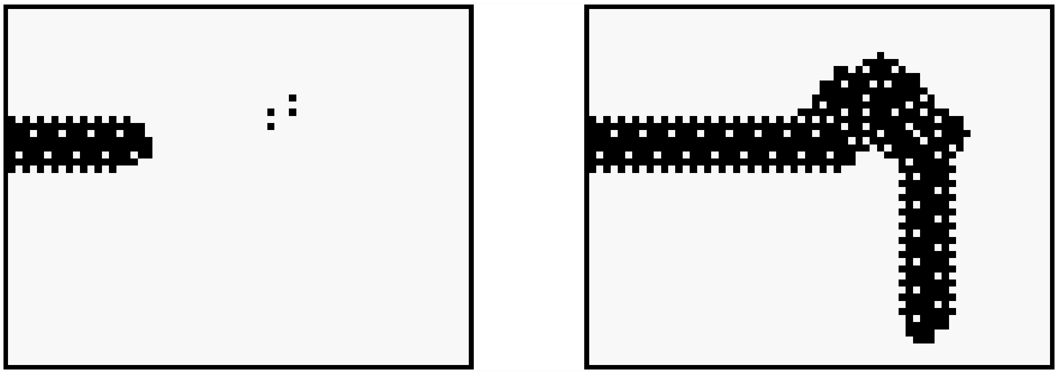

Fig. 7. Another turn.

Blocking in LwoD raises the most delicate issues of the construction since a horizontal ladder crashing into the side of a vertical one creates uncontrolled chutes and lava unless the collision has the right timing. So let us denote the 8 possible phases of such an interaction as (h, v), where $h \in \mathbb{Z}_4$ and $v \in \mathbb{Z}_2$ are the horizontal and vertical spatial phases. We represent h by the shape of the horizontal ladder's head the first time it advances to a new column, while v is 0 or 1 depending on whether the top row of the horizontal ladder lines up with an occupied or unoccupied cell on the vertical ladder's left edge.

Figure 8 shows the two clean blocking collisions which arise. If we label the lower one as (0,0), the upper one is (2,1) since the head of its horizontal ladder has an identical shape but is two cells further to the left, while it has opposite parity with respect to the boundary of the vertical ladder. (There are also two 'dirty' collisions, which create small transient eruptions of lava, but these are not needed for our purposes.) Note that the final fates of the two collisions are mirror images; this is a consequence of the head's flip symmetry after half a cycle. To complete our verification of the ladder CA axioms, we need to show that any ladder headed for the side of another already present can be phase-shifted to guarantee a clean block.

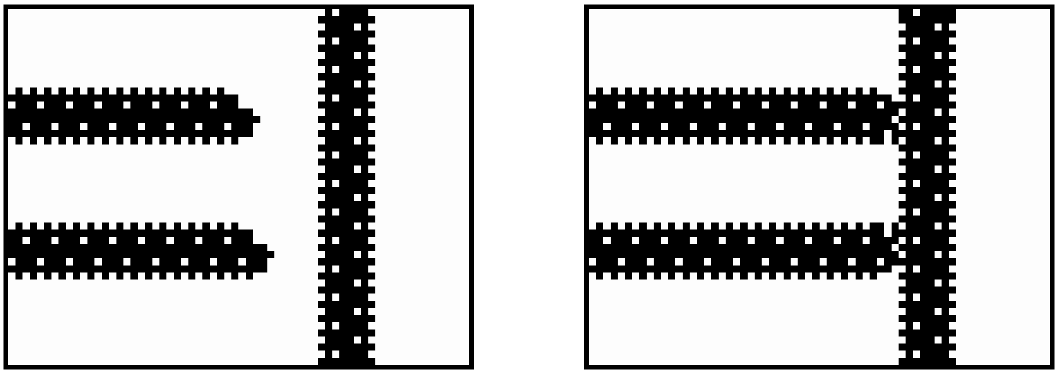

Fig. 8. The two clean blocking collisions, with phases $(h, v) = (0, 0)$ and $(2, 1)$.

Our solution combines the turns of figures 6 and 7 into a two-turn perturbation induced by the static debris of the middle frame in figure 9. The top and middle frames show the same horizontal ladder approaching the same vertical one, while the bottom frame shows the result of collision with the debris. The net effect is an advance of the same head one cell horizontally and a change of vertical parity, shifting (h, v) by $(-1, 1) = (3, 1)$. This action partitions the group $\mathbb{Z}_4 \times \mathbb{Z}_2$ into a subgroup consisting of phases of the form (even, even) and (odd, odd), and another coset comprised of the remaining phases. Since one of the clean blocks in figure 8 belongs to the subgroup, and the other to the complementary coset, it follows that an arbitrary ladder can be phase-shifted to achieve a clean block by applying between 0 and 3 of these shifts. This completes the demonstration.

Let us conclude with a few additional remarks. First, the inevitable question: are there other ladder CA rules? A number of minor variants of Life without

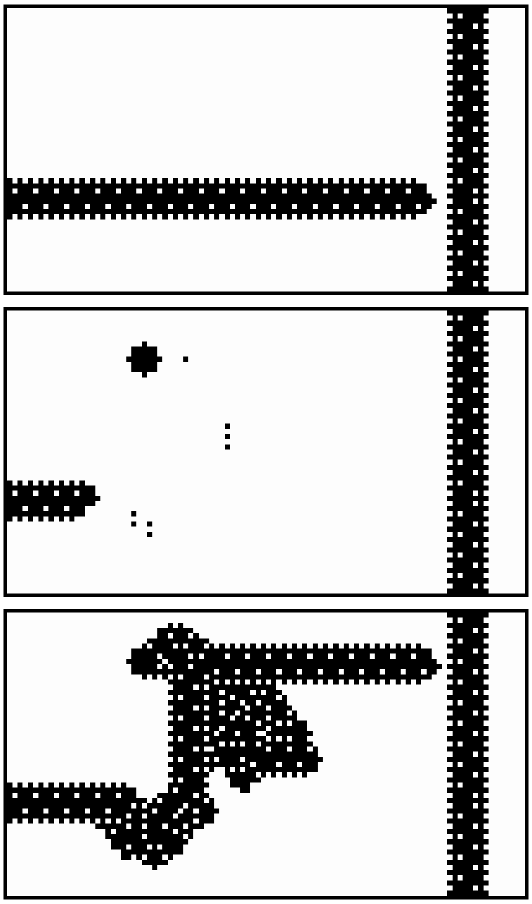

Fig. 9. A $(-1,1)$ phase-shift.

Death admit ladders with similar but slightly different architectures. Four alternatives on the 8-cell Moore neighborhood[2] we know of have the same birth rule, but preclude survival for certain population counts: (i) 8, (ii) 7, (iii) 7 or 8, and (iv) 1 or 3. These slight differences no doubt destabilize many, if not all, of the elementary interactions just described, so it remains an open question whether any of these variants satisfies our axioms. We invite the reader with an insatiable appetite for experimental mathematics to investigate these or other candidates. Another interesting rule (see [7], recipe48.html) has birth with 2 or 5 occupied neighbors, and certain survival. In that case there are diagonal ladders, but we have not investigated the possibilities for turns and other complex collisions. Clearly a ladder CA can be designed to have precisely the desired properties by including sufficiently many states and/or sites in the rule table; the appeal of LwoD and its cousins lies in their mathematical simplicity and their ability to emulate circuits on their own terms.

Finally, we pose an intriguing open question about Life without Death: is there any initial seed that fills the lattice linearly in time with a positive asymptotic density? Experiments from large lattice balls suggest there is, but we know of no way to prove this other than to construct a 'ladder gun' or a space-filler such as [1]. This would require far greater ingenuity than the above constructions, but we think there is a good chance that one is possible.

[2] No relation.

Acknowledgments. We thank Janko Gravner for help with the LwoD figures and Dean Hickerson for expertise on the behavior of LwoD's relatives. We also thank Rascal, Scurry, Oscar and Spootie for companionship and inspiration. This work was supported in part by NSF grant ASC-9503162.

References

1. D. Bell, *Max*, http://www.cs.jhu.edu/ callahan/patterns/max.html
2. A. Condon, "A theory of strict **P**-completeness." STACS 1992, in *Lecture Notes in Computer Science* **577** (1992) 33–44.
3. R. Fisch and D. Griffeath, *WinCA: a cellular automaton modeling environment.* Windows 3.x/95 software. version 1.0, 1996. Available from [7].
4. E. Fredkin and T. Toffoli, "Conservative logic." *Int. J. Theor. Phys.* **21** (1982) 219–253.
5. J. Gravner and D. Griffeath, *Cellular automaton growth on* $\mathbb{Z}^2$*: theorems, examples, and open problems.* In preparation.
6. R. Greenlaw, H.J. Hoover, and W.L. Ruzzo, *Limits to Parallel Computation:* **P**-*Completeness Theory.* Oxford University Press, 1995.
7. D. Griffeath. *Primordial Soup Kitchen,* http://math.wisc.edu/ griffeat/kitchen.html
8. Jörg Heitkötter, *Zooland*, http://www.krl.caltech.edu/ brown/alife/zooland/
9. K. Lindgren and M.G. Nordahl, "Universal Computation in Simple One-Dimensional Cellular Automata." *Complex Systems* **4** (1990) 299–318.
10. J. Machta and R. Greenlaw, "The computational complexity of generating random fractals." *J. Stat. Phys.* **82** (1996) 1299
11. C. Moore, "Quasi-linear cellular automata." *Physica D* **103** (1997) 100-132, *Proceedings of the International Workshop on Lattice Dynamics.*
12. C. Moore, "Non-Abelian cellular automata," Santa Fe Institute Working Paper 95-09-081, or "Predicting non-linear cellular automata quickly by decomposing them into linear ones", submitted to *Physica* **D**.
13. C. Moore, "Majority-Vote Cellular Automata, Ising Dynamics, and **P**-Completeness." Santa Fe Institute Working Paper 96-08-060, to appear in *J. Stat. Phys.*
14. C. Moore and M.G. Nordahl, "Lattice Gas Prediction is **P**-complete." Santa Fe Institute Working Paper 97-04-034, available as comp-gas/9704001.
15. C. Moore, "The computational complexity of sandpiles." In preparation.
16. N. Packard and S. Wolfram, "Two-dimensional cellular automata." *J. Stat. Phys.* **38** (1985) 901–946, and private communication.
17. C.H. Papadimitriou, *Computational Complexity.* Addison-Wesley, 1994.