

# Embedding Computation in One-Dimensional Automata by Phase Coding Solitons

KENNETH STEIGLITZ, FELLOW, IEEE, IRFAN KAMAL, AND ARTHUR WATSON

**Abstract**—We show that some kind of meaningful computation can be embedded in very simple, microscopically homogeneous, one-dimensional automata—filter automata with a parity next-state rule.

A systematic procedure is given for generating moving, periodic structures (“particles”). These particles exhibit soliton-like properties; that is, they often pass through one another with phase shifts. We then discuss ways to encode information in the phase of these particles.

Finally, the search for useful logical operations is reduced to a search for paths in certain graphs. As a demonstration of principle, we give the details of implementing a carry-ripple adder.

**Index Terms**—Cellular automata, parallel computation, solitons.

## I. INTRODUCTION

IN [1] a class of one-dimensional automata was described, called the *parity-rule filter automata* (parity-rule FA’s), which support particles with soliton-like properties. That is, although the operation of the automaton is nonlinear and irreversible, moving persistent structures pass through one another while preserving their identities. In this paper we will study the systematic generation and properties of these solitons. We will then show how they can be used to encode information and perform useful computation—using the carry-ripple adder to demonstrate the principle.

Embedding computation in a simple one-dimensional automaton has important practical advantages over the two-dimensional alternative. One-dimensional structures are easier to build and operate, they can be looped naturally, and no *a priori* decisions need be made about the size of the array. A one-dimensional automaton can be implemented in a highly parallel way much more easily than can a two-dimensional one, because the cell values can be streamed serially through many cascaded processors. In [2], for example, a VLSI implementation of a simple, fixed CA was described which performs more than  $10^8$  updates per second per chip. The concatenation of many identical processors then results in a highly concurrent machine with a fixed, regular structure.

The principal motivation of this work is the exploration of

Manuscript received December 9, 1985; revised September 30, 1986. This work was supported in part by NSF Grant ECS-8414674, U.S. Army Research—Durham Contract DAAG29-85-K-0191, DARPA Contract N00014-82-K-0549, and ONR Grant N00014-83-K-0275.

The authors are with the Department of Computer Science, Princeton University, Princeton, NJ 08544.  
IEEE Log Number 8716239.

microscopically homogeneous structures that support computation. Pseudoparticles then play a natural role in providing a medium for transmitting information from one place to another. As we will see, solitons are especially interesting pseudoparticles from this point of view, because they bear information in their carrier and envelope phase, and this information is changed when solitons collide.

Carter [3] discusses the possibility of implementing logical functions to build cellular automata (CA’s) using true (physical) solitons supported by chemical structures, while we discuss the opposite: implementing logical functions using the soliton-like structures that arise in certain automata. A further practical aspect of our work is the possibility that we can use true solitons in much the same way that we use those supported by automata.

## II. FILTER AUTOMATA

For the purposes of this paper we will restrict discussion to the special class of one-dimensional, binary-state, filter automata described in [1]. The site values will be denoted by  $a_i^t$ ,  $i = 1, \dots, n$ , where the superscript is the *time* variable  $t$ ,  $0 \leq t \leq +\infty$ , and the subscript is the *space* variable  $i$ ,  $-\infty \leq i \leq +\infty$ . The evolution of the automaton is determined by the fixed rule  $F$  of the form

$$a_i^{t+1} = F(a_{i-r}^{t+1}, a_{i-r+1}^{t+1}, \dots, a_{i-1}^{t+1}, a_i^t, \dots, a_{i+r}^t) \quad (2)$$

with

$$F(0, 0, \dots, 0) = 0.$$

The next state is thus computed using the newly updated values  $a_{i-r}^{t+1}, a_{i-r+1}^{t+1}, \dots, a_{i-1}^{t+1}$ , instead of  $a_{i-r}^t, a_{i-r+1}^t, \dots, a_{i-1}^t$ , as in the usual cellular automaton. This is analogous to the operation of an *infinite impulse response* digital filter, whereas a cellular automaton corresponds to a *finite impulse response* digital filter (see [4], for example).

It will be convenient to draw two-dimensional pictures that represent the evolution of an FA, with time increasing downward. In such a picture the FA “update window” looks like

$$\begin{array}{ccccccc} & & & & z & 1 & \cdots & r-1 & r \\ & & & & & & & & \\ -r & -r+1 & \cdots & -1 & z & & & & \end{array}$$

where  $z$  represents the “center” of the window.

We will also be discussing only one kind of FA; the next-





However, suppose that all cells are known except the  $r$  cell at the upper right. Then its contents are obviously determined by the remaining cells in all but the two cases where everything in the window except perhaps the lower  $z$  cell is filled with a 0:

1) If the lower  $z$  contains a 1, then there is no way to fill in the upper  $r$  cell and be consistent with the update rule. When this situation is reached in the algorithm, it means that no particle with the current left edge exists, so another is tried.

2) If the lower  $z$  contains a 0, then the upper  $r$  cell can be filled with either a 0 or a 1 and still be consistent with the update rule. In the algorithm it is filled with a 0. To fill it with a 1 would create a second left particle boundary transition, which a single particle cannot contain.

Now it becomes clear how the algorithm proceeds. There are only  $r^p$  possible left ends to a period- $p$  particle, because of the "speed of light,"  $r - 1$ , and the fact that the left edge of a particle cannot move to the right. Furthermore, such initial configurations can be generated easily in lexicographic order, complete with  $r$  0's to the left of each initial 1. We begin, then, with  $p - 1$  (staggered) stacked "windows," the top of each one residing on its own orbital row (excluding the bottom row), each having one of the initial 1's in its upper  $(r - 1)$ st position. Then, as described above, we proceed upward from the bottom window, filling in the second cell of the particle at each orbital position. When we have filled in the second cell of the top orbital position, we copy the contents to the second cell of the bottom orbital position, to enforce the period- $p$  condition. Then we advance all of the windows one step, and proceed for succeeding iterations just as in the first. This main loop terminates when one of the following three conditions is satisfied.

1) A gap of  $2r + 3$  consecutive 0's is found in the particle, in which case it is recorded and the next initial configuration is considered.

2) A state is reached where there is no possible way to fill in one of the cells, in which case the next initial configuration is considered.

3) The maximum particle length being searched is exceeded, in which case the next initial configuration is considered.

Thus, an outline of the algorithm is

```

for each of the  $r^p$  left edges
  while ( a. max particle length not exceeded and
         b. dead end not reached and
         c. particle not isolated )
    begin
      update windows;
      record particle if appropriate
    end.
```

The complexity of the algorithm is easily seen from the above outline. There are  $r^p$  steps, each consisting of a maximum of  $L$  ( $=$  maximum desired particle length) iterations of a constant time loop. Therefore, the "practical" complexity of the algorithm is best stated as  $O(r^p L)$ , especially since we are usually concerned with particles smaller than some given value of  $L$ ; for instance  $L = 32$  gives all particles which are codable by an integer on a VAX, or  $L$

$= 80$  gives all particles whose pictures fit on a CRT screen.

However, we can use the algorithm to generate all particles of period  $p$  regardless of size. To see this, consider the state of the construction procedure at a given instant. Note that the current state, and thus the entire future of the procedure, is completely determined by the contents of the current windows. Thus, an exact repetition of the contents of all of the window cells at some time during the procedure is equivalent to nontermination of the main loop (if the maximum length condition is dropped), since then the same repetition would occur forever. This is clearly equivalent to the nonexistence of a period- $p$  particle with the left end which is currently being considered. Now there are no more than  $2(r + 1)(p - 1)$  window cells, and since each can have only two possible values, there are at most  $4^{(r+1)(p-1)}$  distinct contents. Thus, there must be some repetition of contents some time within the first  $4^{(r+1)(p-1)}$  iterations of the main loop, assuming that no particle was found before then. At that time we may conclude that no particle is forthcoming, and go on to the next initial left edge. Thus, taking  $L = 4^{(r+1)(p-1)}$ , we have an algorithm for generating all particles of period  $p$  for a given  $r$ , with the time complexity of  $O(r^p 4^{(r+1)(p-1)})$ .

As a bonus, the argument above establishes the following fact.

*Theorem 1:* The number of particles with period  $p$  for any radius- $r$  filter automaton is no greater than  $r^p$ .  $\square$

#### V. CODING INFORMATION IN PARTICLE PHASE

A particle can be viewed as carrying information in the following way. Assume that we establish an absolute origin in space ( $x = 0$ ) and time ( $t = 0$ ), and a reference particle that starts in the state corresponding to its canonical code with its left end at  $x = 0$ . Given any instance of a particle, run it backwards in time to  $t = 0$  and measure its orbital state  $s$  and the position  $x$  of its left end. The orbital state so obtained will be called the *orbital phase* of the particle, and will take on values  $s = 0, 1, \dots, p - 1$ , where  $p$  is the period, and  $s = 0$  corresponds to the canonical code of the particle. The position  $x$  so obtained is called the *translational phase*. Thus, when no collision takes place, the orbital and translational phases of a particle both remain 0. We might wish to keep track of the translational phase only modulo the particle's displacement  $d$ , in which case the particle has  $p \times d$  distinct phase states.

We will choose to do things a little differently, however, because we will be studying the effects of collisions between pairs of particles. To go further we need to discuss collisions in some detail. First, we note that if two particles start close enough together, it may happen that they interact in a way that is impossible when they start far apart. In such cases, we say the collision is *improper*; otherwise we say it is *proper*. We will restrict our attention to proper collisions, because we will always provide an initial spacing large enough to ensure a typical collision. Next, two collisions will be said to be the *same* if they can be put in concordance by a shift in space and time. Otherwise, they will be *different*. Finally, we need the following fact, which is proved in [1].

*Theorem 2:* Let two particles have displacements  $d_1, d_2$ , periods  $p_1, p_2$ , and speeds  $d_1/p_1 < d_2/p_2$ , so that particle 2

can hit 1. Let the difference in speeds be  $\Delta s = d_2/p_2 - d_1/p_1$ . Then the number of different proper collisions is no larger than

$$DET = p_1 p_2 \cdot \Delta s = p_1 d_2 - p_2 d_1$$

and displacing one particle by DET results in the same collision, as long as it remains proper.

We call DET the *determinant* of the collision, it being the  $2 \times 2$  determinant with rows  $p_1 p_2$  and  $d_1 d_2$ .  $\square$

It is now clear that if we have particles with only two different displacement-period pairs, then it is the relative positions of particles modulo DET that count, insofar as the results of collisions are concerned. For this reason, we will measure translational phase modulo DET, and a particle in such a collision system will have  $p \times DET$  phase states.

In what follows we will place particles in the initial array of the automaton at predetermined positions and in predetermined orbital states. Collisions will then take place, and the final results of those collisions will arrive at some arbitrary fixed position far to the left, which will be thought of as the "output" position. We assume that an observer can measure the time of arrival and orbital state of each particle that arrives at the output position, and thereby determine its orbital and translational phase.

### VI. COLLISION TABLES

We now describe a table that determines the results of pairwise collisions between two given particles. The particular form of the table is not the only one possible, but is one that will be useful to us in the carry-ripple example. The table is constructed empirically, by simply simulating the parity-rule automaton.

A typical table (Table I) is given below for collisions in the  $r = 5$  parity-rule filter automaton between the code 155 particle with displacement-period pair (22/8), and the faster code-11 particle with  $d/p = (9/3)$ . There are three sections in the table, each corresponding to a possible initial orbital state of the fast particle; the codes in the orbit of the fast particle are given in the headings, namely 11, 25, 37. Within each section there is one row for DET consecutive spacings, starting with a value large enough to ensure that all collisions are proper; in this case  $3 \cdot r = 15$ . The numerical value of the spacing is the number of cells between the right end of the slow particle and the left end of the fast. In each row we record the code of the resulting fast particle, and if it is in the orbit of the original fast particle (that is, if the identity of the fast particle is unchanged by the collision), we follow this with the changes in its orbital and translational phases ( $\Delta\phi_o$  and  $\Delta\phi_t$ ). The same information for the slow particle follows in that row.

Because there can be only DET distinct collisions, it should be clear that the second and third sections of the table are derivable from the first, given the displacement-period pairs of the particles and the widths of the orbital codes. We will see that this redundant form of the table is convenient when searching for useful logical operations, as described in the next section.

This table has all the information we need to predict the result of any collision between a fast and slow particle, given

TABLE I  
AN EXAMPLE OF A STATE-TRANSITION TABLE FOR THE COLLISION OF TWO PARTICLES

start		155	22	8	11	9	3
spacing	code	$\Delta\phi_o$	$\Delta\phi_t$	code	$\Delta\phi_o$	$\Delta\phi_t$	
15	11	0	2	167	-	-	
16	11	2	3	155	2	2	
17	11	1	2	155	0	0	
18	11	2	3	155	2	2	
19	11	2	3	155	2	2	
20	659	-	-	3	-	-	
next		155	25				
15	11	2	4	155	2	2	
16	659	-	-	3	-	-	
17	11	0	2	167	-	-	
18	11	2	4	155	2	2	
19	11	1	1	155	0	0	
20	11	2	4	155	2	2	
next		155	37				
15	11	2	5	155	2	2	
16	11	2	5	155	2	2	
17	659	-	-	3	-	-	
18	11	0	2	167	-	-	
19	11	2	5	155	2	2	
20	11	1	3	155	0	0	
end							

Each of the three sections corresponds to a different initial orbital state of the fast particle (namely 11, 25, and 37); each row within each section corresponds to an initial spacing;  $\Delta\phi_o$  corresponds to the shift in orbital phase;  $\Delta\phi_t$  corresponds to the shift in translational phase, measured positive to the left for the slow particle, and positive to the right for the fast. When the identity of a particle changes, the phase shifts are undefined.

that in the initial state the slow particle is in its canonical state. As an example, suppose that the fast particle (11) is in orbital state 25 and its left end is positioned 101 cells to the right of the right end of the slow particle, which is in its canonical state of 155. The progression of states is periodic modulo the determinant, 6, so we need to enter the table at the row corresponding to a spacing of  $17 = 101 \text{ mod } 6$ . The entry in the code-25 section of the table then shows that the fast particle emerges with its identity intact, and with an orbital phase shift of 0 and a translational phase shift of 2, while the slow particle is changed to one with canonical code 167.

### VII. SEARCHING FOR LOGICAL OPERATIONS

There are many ways we might try to encode and process information using solitons, and we will describe here only one, our goal being to demonstrate the principle.

In the basic scheme we will study, the fast particle is in one of two states, identified with "0" and "1," and it passes through a number of composite structures, each composed of one or more slow particles in their canonical states but with variable spacing, and with identical displacement-period pairs. (See Fig. 2.) These composite structures will be called *particle bundles*.

Because the slow particles are all in their canonical states, the results of all the collisions can be predicted from the collision tables described in the previous section, and in fact the "0" and "1" states of the fast particle can be associated with two particular rows of the table. The problem is to choose




Fig. 2. The collision scheme studied here. A single fast particle passes through several slow particles in their canonical states, comprising a particle bundle.

the two rows so that the effects of collisions correspond to the logical operations we wish to perform.

We can automate the search for the slow particle bundle as follows. Create a directed graph  $G = (V, A)$  with one node for each pair of rows in the collision table. Choose a fast particle and a set of slow particles with a given displacement-period pair. Suppose passage of the fast particle through a particular slow particle, at a particular spacing offset relative to an arbitrary reference (see Fig. 2), maps row  $i$  to row  $i'$  and row  $j$  to row  $j'$ . Then create an arc that goes from the node corresponding to row-pair  $ij$  to the node corresponding to  $i'j'$ , and label this arc with the code of the slow particle and the spacing offset that produced this result. Do this for all possible collisions between the given fast particle and the given set of slow particles.

If now we want to map the fast particle state "0" to "X" and "1" to "Y," where  $X$  and  $Y$  are each 0 or 1, we need to find a path in the graph  $G$  that goes from a node of the form  $AB$  to a node of the form  $CD$ , where  $C = A$  if  $X = 0$ ,  $C = B$  if  $X = 1$ , and similarly for  $D$  and  $Y$ . (For example, if we want the complementation operation, we search for a path from a node of the form  $AB$  to a node of the form  $BA$ .) Each arc on the path represents a slow particle at a particular spacing offset in the bundle that we need to pass through to effect the operation.

The problem of implementing a particular set of logical operations thus reduces to that of finding a corresponding set of paths in a directed graph. The graph has  $(\text{DET} \cdot p_2)^2$  nodes, where  $p_2$  is the period of the fast particle; and  $\text{DET}$  arcs leaving each node for each slow particle allowed in a bundle. We should use breadth-first search to implement the path search so as to find paths with the minimum number of arcs.

To make these ideas more concrete, we describe the embedding of a simple carry-ripple adder.

#### VIII. THE CARRY-RIPPLE ADDER

To implement a carry-ripple adder, we will encode each pair of addend bits to one of three different particle bundles:  $w$  for the case when the addend bits are both 0,  $x$  when one is 0 and the other 1, and  $y$  when they are both 1. (See Fig. 3 for an example.) The resulting sequence of particle bundles is then transmitted (at the slow speed of the particles used to construct the particle bundles) to the left, most significant bit-pairs first. They are separated by a distance that is 0 mod  $\text{DET}$ , so that the collisions with the fast particle will be in accordance with the collision table; but a distance large enough to ensure proper collisions. The fast particle, which represents the carry bit, is then sent in from the right, at an initial spacing that is equal mod  $\text{DET}$  to the initial spacing chosen for the collision table (in the example of Table I, a spacing of 15).

Given this particular structure, it is now easy to write down the transition rules that the collisions between the fast particle and the slow-particle bundles need to satisfy. Letting  $c$  be the state of the carry bit, we require

$$\begin{aligned} wc &\rightarrow 0w' \\ xc &\rightarrow cx' \\ yc &\rightarrow 1y'. \end{aligned}$$

In these rules,  $wc$  represents the configuration with the particle bundle  $w$  to the left of the carry bit  $c$ , the arrow indicates the result of the collision, and  $0w$  means that the fast particle is moved to its "0" state, and emerges to the left of the particle bundle  $w$ .

For this scheme to operate properly as an adder, we must be able to decode the resulting slow-particle bundles, and that means that the result  $w'$  when  $c = "0"$  must be different from  $w'$  when  $c = "1"$ , and similarly for  $x'$  and  $y'$ . We call this the *distinguishability* criterion. Finally we require that all the particles taking part in collisions have their speed unchanged by the collisions, so that the results of the collisions can be easily interpreted.

The graph-search problem that results from this formulation of the carry-ripple adder is not difficult, and many solutions have been obtained. One such solution for  $r = 5$ , and two slow particles in each slow-particle bundle, is the following:

fast particle: code 11,  $d/p = 9/3$   
 fast particle state "0" = orbital state 0, offset 0  
 fast particle state "1" = orbital state 0, offset 2  
 slow particle  $d/p = 22/8$ ,  $\text{DET} = 6$   
 slow-particle bundle  $w$ : code 155 (offset 2) code 155 (offset 4)  
 slow-particle bundle  $x$ : code 165 (offset 2) code 207 (offset 3)  
 slow-particle bundle  $y$ : code 155 (offset 0) code 165 (offset 2).

Fig. 4 shows a typical collision in this adder, the collision  $yc$  when  $c$  is in state "0."

Another solution was found for  $r = 3$ , with three slow particles in each slow-particle bundle

fast particle: code 7,  $d/p = 5/3$   
 fast particle state "0" = orbital state 0, offset 0  
 fast particle state "1" = orbital state 0, offset 1  
 slow particle  $d/p = 14/10$ ,  $\text{DET} = 8$   
 slow-particle bundle  $w$ : code 919 (offset 3) code 919 (offset 0) code 919 (offset 5)  
 slow-particle bundle  $x$ : code 919 (offset 3) code 919 (offset 0) code 731 (offset 2)  
 slow-particle bundle  $y$ : code 731 (offset 7) code 919 (offset 0) code 731 (offset 2)

#### IX. DISCUSSION

We have demonstrated that some kind of meaningful computation can be embedded in a very simple, microscopically homogeneous, one-dimensional automaton. This embedding differs in one important aspect from the ones used to show the universality of two-dimensional cellular automata. In

		addend codes							
		B	C	A	B	B	A	C	
	addend 1	0	1	0	1	0	0	1	
	addend 2	1	1	0	0	1	0	1	
end carry	↓	1	0	0	0	1	1	1	0
		sum							




Fig. 3. An example of the form of a carry-ripple adder. The fast particle travels through the slow-particle bundles, propagating the carry bit.




Fig. 4. The collision implementing  $yc \rightarrow 1y'$  when  $c$  is "0"—the carry bit goes high. The particle codes are, from left to right, 155 (22/8), 165 (22/8), 11 (9/3), and the offsets 0, 2, 0. Note that the identity of the 155 is changed, to code 167 (22/8). (The picture is circularly wrapped to fit on the page.)

the latter, many cells must be put in their proper initial states before computation can proceed, and bits are stored by the presence or absence of particles [5], [6]. In the one-dimensional structure discussed here, all the information necessary for computation to take place can enter in bit-serial form, with all the initial states zero. The question of whether filter automata are universal, however, remains open.

#### REFERENCES

- [1] J. K. Park, K. Steiglitz, and W. P. Thurston, "Soliton-like behavior in automata," *Physica D*, vol. 19D, pp. 423-432. Reprinted in *Theory and Applications of Cellular Automata*, S. Wolfram, Ed. Hong Kong: World Scientific, 1986, pp. 333-342.
- [2] K. Steiglitz and R. R. Morita, "A multi-processor cellular automaton chip," in *Proc. 1985 IEEE Int. Conf. Acoust., Speech, Signal Processing*, Tampa, FL, Mar. 1985.
- [3] F. L. Carter, "The molecular device computer: Point of departure for large scale cellular automata," in *Cellular Automata*, D. Farmer, T. Toffoli, and S. Wolfram, Eds. Amsterdam, The Netherlands: North-Holland Physics, 1984, pp. 175-194.
- [4] A. V. Oppenheim and R. W. Schaefer, *Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1975.
- [5] E. R. Berlekamp, J. H. Conway, and R. K. Guy, "What is life?" in *Winning Ways for Your Mathematical Plays, Vol. 2: Games in Particular*. New York: Academic, 1982, ch. 25.
- [6] F. Nourai and R. S. Kashef, "A universal four-state cellular computer," *IEEE Trans. Comput.*, vol. C-24, pp. 766-776, Aug. 1975.
- [7] S. Wolfram, "Glider gun guidelines," unpublished manuscript, Mar. 1985. (Available from *Scientif. Amer.*)
- [8] ———, "Universality and complexity in cellular automata," *Physica D*, vol. 10D, pp. 1-35, 1984. Reprinted in *Theory and Applications of Cellular Automata*, S. Wolfram, Ed. Hong Kong: World Scientific, 1986, pp. 91-125.
- [9] C. H. Goldberg, manuscript in preparation.




**Kenneth Steiglitz** (S'57-M'64-SM'79-F'81) was born in Weehawken, NJ, on January 30, 1939. He received the B.E.E., M.E.E., and Eng.Sc.D. degrees from New York University, New York, NY, in 1959, 1960, and 1963, respectively.

Since September 1963 he has been at Princeton University, Princeton, NJ, where he is now Professor of Computer Science, teaching and conducting research on VLSI design and implementation of signal processing, optimization algorithms, and the foundations of computing. He is the author of

*Introduction to Discrete Systems* (New York: Wiley, 1974), and coauthor, with C. H. Papadimitriou, of *Combinatorial Optimization: Algorithms and Complexity* (Englewood Cliffs, NJ: Prentice-Hall, 1982).

Dr. Steiglitz is a member of the VLSI Committee of the IEEE ASSP Society, is serving his second term as member of the Administrative Committee, and has also served on the Digital Signal Processing Committee, and as Awards Chairman of that Society. He is an Associate Editor of the journal *Networks*, and is a former Associate Editor of the *Journal of the Association for Computing Machinery*. A member of Eta Kappa Nu, Tau Beta Pi, and Sigma Xi, he received the Technical Achievement Award of the ASSP Society in 1981, the ASSP Society Award in 1986, and the IEEE Centennial Medal in 1984.



**Irfan Kamal**, a citizen of Bangladesh, was born on July 1, 1965. He received the B.S.E. degree (with Honors) in electrical engineering and computer science from Princeton University, Princeton, NJ, in 1986.

His interests include novel computer architectures, graph theory, and the design and analysis of learning systems. At present he is on leave from the field, working as an analyst at Merrill Lynch.

Mr. Kamal has been a member of Tau Beta Pi since 1985.



**Arthur Watson** was born in New York on November 25, 1964. He received the A.B. degree in mathematics from Princeton University, Princeton, NJ, in 1986.

He is currently working towards the Ph.D. degree in computer science at Princeton University. His current interests include algorithm and automata theory and computer games.