Automorphisms of transition graphs for linear cellular automata

EDWARD POWLEY*, SUSAN STEPNEY

Department of Computer Science, University of York, UK

The *transition graph* of a cellular automaton (CA) represents the CA's global dynamics. The *automorphisms* of the transition graph are its self-isomorphisms or "symmetries", and in a sense these are precisely the symmetries of the CA's dynamics.

Previously we have argued that studying how the total number of automorphisms varies with the number of cells on which the CA operates yields a partial classification of CA rules. One of the classes thus identified consists mainly of *linear* CAs; that is, CAs whose local rules are linear functions.

In this paper, we use the algebraic properties of linear CAs in general, and of one linear CA (elementary rule 90) in particular, to derive an expression for the number of automorphisms admitted by that CA. We use this expression to produce numerical results, and observe that the number of automorphisms as a function of the number of cells exhibits a correlation with a number-theoretic function, the *suborder function*.

1 INTRODUCTION

A *cellular automaton (CA)* consists of a finite nonempty set of *states*, a discrete lattice of *cells*, and a *local update rule* which maps deterministically the state of a cell and its neighbours at time t to the state of that cell at time $t+1$. A *configuration* of a CA is an assignment of a state to each cell. The local update rule extends to a *global map*, a function from configurations to configurations, in the natural way.

* email: `ed@cs.york.ac.uk`

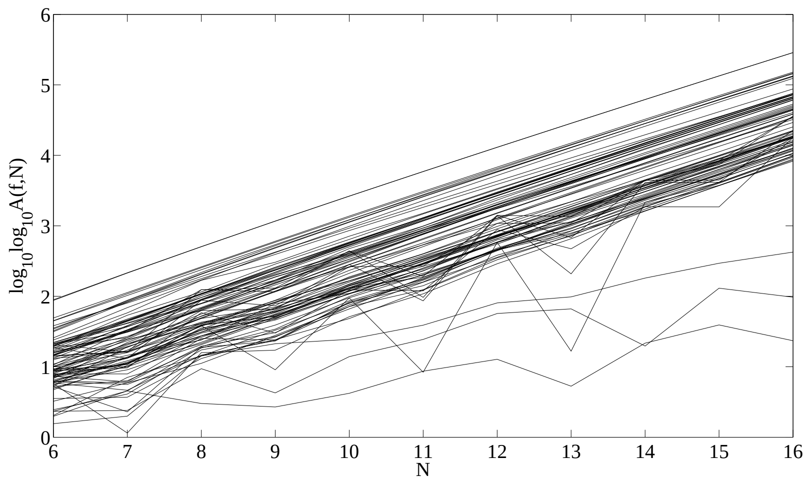

FIGURE 1
Plot of $\log_{10} \log_{10} A(f, N)$ (where $A(f, N)$ is the number of automorphisms) against N, for $6 \leq N \leq 16$ and for all 88 essentially different ECA rules. From [3].

The *transition graph* of a CA is a directed graph representation of the CA's global dynamics: the vertices of the transition graph correspond to the configurations of the CA, and the edges to the transitions between configurations according to the global map. An *automorphism* of a transition graph is an isomorphism of the graph onto itself. It can be shown that the group of automorphisms of the transition graph is precisely the group of bijections which commute with the CA's global map. In physics, a common notion of "symmetry" is of a transformation which, when applied to the initial state of a system, causes the subsequent behaviour of the system to undergo the same transformation but remain otherwise unchanged. Thus the automorphisms of the transition graph are, in a sense, the symmetries of the CA's global dynamics.

In [3] we investigate how the sizes of the automorphism groups vary with the number of cells N on which the CA operates, for the 88 essentially different *elementary CAs* (*ECAs*; see Section 2.2). Our numerical results are shown in Figure 1. We identify three distinct types of behaviour:

1. The relationship between N and the double logarithm of the number of automorphisms is approximately linear;

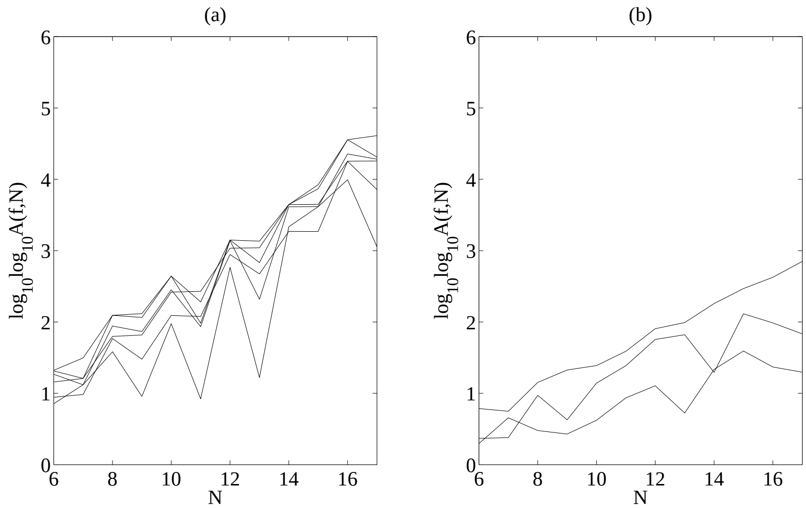

FIGURE 2
As Figure 1, but for $6 \leq N \leq 17$, and showing the (a) second and (b) third types of behaviour described in the text. From [3].

2. The number of automorphisms seems to alternate between large and small values for even and odd N respectively (illustrated in Figure 2 (a));

3. The relationship is neither linear nor alternating, and the average rate of growth of the number of automorphisms is relatively slow (illustrated in Figure 2 (b)).

The majority of ECAs exhibit the first type of behaviour. Two of the three ECAs which exhibit the third type of behaviour are those identified by Wolfram [6] as being particularly well-suited to the generation of pseudo-random numbers, and all three exhibit relatively long-lived transient behaviour before reaching an attractor cycle. However, it is the second type of behaviour on which we focus in this paper.

A CA is *linear* if its local update rule is a linear function (a sum of constant multiples of its arguments). Of the six ECAs which exhibit the second type of behaviour described above, three are linear and a further two are linear plus a constant term. Furthermore, the only linear ECAs which do not exhibit this type of behaviour are the “trivial” ones: the identity ECA, the zero ECA, and the left- and right-shift ECAs. Thus we are justified in saying that this type of behaviour is characteristic of non-trivial linear rules, at least among the ECAs.

In [3], we describe a way of counting the automorphisms of a transition graph by iterating over its structure in a recursive fashion. The major drawback to that approach is that the entire transition graph must be generated, and the number of vertices in this graph is exponential with respect to N. In this paper (which extends and generalises some of the results presented in [2]), we use results of Martin et al [1] concerning the structure of the transition graph for a linear CA, and for one linear ECA (known as *rule 90*) in particular, to count automorphisms without the need to generate the entire graph. While the "brute force" method of [3] exhausts our computational resources for $N > 17$, the approach described herein only begins to do so for $N > 185$.

Having introduced in Section 2 the necessary definitions, in Section 3 we discuss how the property of linearity restricts the possible shape of the transition graph for a linear CA, and in Section 4 we consider how this helps us in counting the graph's automorphisms. Up until this point our discussion applies to all linear CAs, but in Section 5 we apply and extend the results for ECA rule 90 in particular, which allows us in Section 6 to produce and discuss numerical results.

2 PRELIMINARIES

2.1 Linear CAs

We restrict our attention to finite one-dimensional CAs, i.e. we take the lattice to be $\mathbb{Z}_N$ (the cyclic group of integers modulo N). This lattice has *periodic boundary condition*, in that we consider cell $N-1$ and cell 0 to be adjacent. The neighbourhood is specified in terms of its *radius* r, so that the neighbours of cell i are cells $i-r, \dots, i+r$. We further restrict our attention to CAs whose state set is $\mathbb{Z}_k$ for some $k \geq 2$. Thus the local update rule is a function $f : \mathbb{Z}_k^{2r+1} \to \mathbb{Z}_k$, which extends to a global map $F : \mathbb{Z}_k^N \to \mathbb{Z}_k^N$.

Such a CA is said to be *linear* if the local update rule is a linear function; that is, if there exist constants $\lambda_{-r}, \dots, \lambda_r \in \mathbb{Z}_k$ such that

$$f(x_{-r}, \dots, x_r) = \lambda_{-r} x_{-r} + \dots + \lambda_r x_r \ , \qquad (1)$$

where the operations of addition and multiplication are the usual operations of modular arithmetic in $\mathbb{Z}_k$.

Define the sum of two configurations u, v by

$$(u + v)[i] = u[i] + v[i] \qquad (2)$$

for all cells i, where $u[i]$ denotes the state of cell i in configuration u. Then linear CAs obey a principle of *additive superposition*: for all configurations u, v, we have

$$F(u + v) = F(u) + F(v)\,. \tag{3}$$

Clearly the analogous results hold for any number of repeated applications of F. Intuitively, the evolution of the CA from initial configuration $u + v$ is the sum of the evolutions from initial configurations u and v.

2.2 Elementary CAs

An *elementary CA (ECA)* has neighbourhood radius $r = 1$ and state set $\mathbb{Z}_2$. There are $2^{2^3} = 256$ possible local rules for an ECA. Wolfram [4] describes a way of assigning each of these rules a number between 0 and 255 inclusive; we use this numbering scheme throughout this paper.

Consider two rules to be equivalent if one can be obtained from the other by exchanging states 0 and 1, or by reflecting (reversing) the neighbourhood, or by performing both of these transformations in series. Then the space of local rules is partitioned into 88 equivalence classes [7]; thus we obtain the 88 *essentially different* ECA rules by choosing one rule from each class.

Of the 88 essentially different ECA rules, six are linear:

$$f(x_{-1}, x_0, x_1) = 0 \qquad \text{(rule 0)} \tag{4}$$

$$f(x_{-1}, x_0, x_1) = x_1 \qquad \text{(rule 170)} \tag{5}$$

$$f(x_{-1}, x_0, x_1) = x_0 \qquad \text{(rule 204)} \tag{6}$$

$$f(x_{-1}, x_0, x_1) = x_{-1} + x_0 \qquad \text{(rule 60)} \tag{7}$$

$$f(x_{-1}, x_0, x_1) = x_{-1} + x_1 \qquad \text{(rule 90)} \tag{8}$$

$$f(x_{-1}, x_0, x_1) = x_{-1} + x_0 + x_1 \qquad \text{(rule 150)} \tag{9}$$

In terms of global maps, rule 0 immediately maps every configuration to the homogeneous configuration of zeros, rule 170 shifts the entire configuration by one cell to the left, and rule 204 is the identity. The global dynamics of the other three rules are illustrated in Figure 3.

2.3 Transition graphs

Consider a CA with state set $\mathbb{Z}_k$, and global map F on N cells. The *transition graph** of this CA is the directed graph with vertex set $\mathbb{Z}_k^N$ and edge set

$$\mathcal{E} = \left\{ (s, F(s)) \;:\; s \in \mathbb{Z}_k^N \right\}\,. \tag{10}$$

* Some authors use the term *state transition graph* or *state transition diagram*; however, given that we use "state" to mean the local condition of a single cell (as opposed to "configuration" meaning the global condition of the entire lattice), this term would be confusing.

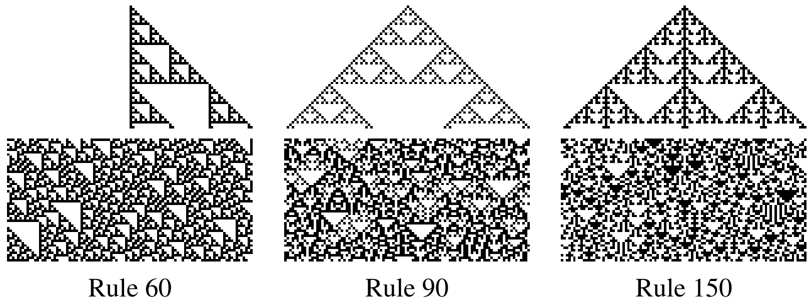

FIGURE 3
Space-time diagrams for the three non-trivial linear ECAs, on an initial configuration consisting of a single cell in state 1 (top) and an initial configuration in which each cell is randomly assigned state 0 or 1 (bottom).

In other words, the vertices correspond to configurations, and the edges show which configurations are mapped to which by the global map. In terms of dynamical systems, the transition graph represents the *phase space* of the CA, with paths in the graph corresponding to trajectories through the phase space.

Examples of transition graphs are shown in Figures 4 and 5.

Since CAs are deterministic, every vertex has out-degree 1. Therefore the graph consists of a number of disjoint cycles, with a (possibly single-vertex) tree rooted at each vertex in each cycle; the edges in the trees are directed away from the leaves and towards the root (i.e. towards the cycle). The disjoint connected components of this graph, each consisting of a single cycle and its trees, are the basins of attraction of the CA, and so we call them the *basins* of the graph.

An *automorphism* of a transition graph is an isomorphism of the graph onto itself; that is, a bijection α on the vertex set such that there is an edge from vertex u to vertex v if and only if there is an edge from $\alpha(u)$ to $\alpha(v)$. Automorphisms are "symmetries" in the sense that they are transformations of the graph which leave the overall shape of the graph unchanged.

Denote by $A(f, N)$ the number of automorphisms of the transition graph for the CA with local rule f on N cells. For the ECAs, we replace f with the rule number, so for example $A(90, N)$ is the number of automorphisms for rule 90 on N cells.

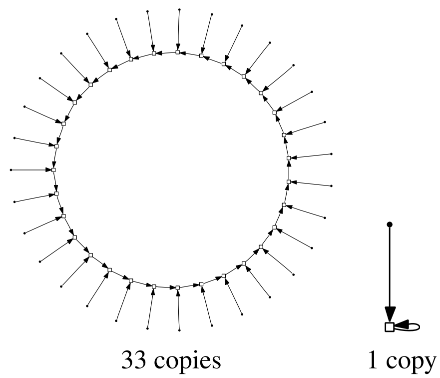

FIGURE 4
Transition graph for ECA rule 90 on 11 cells.

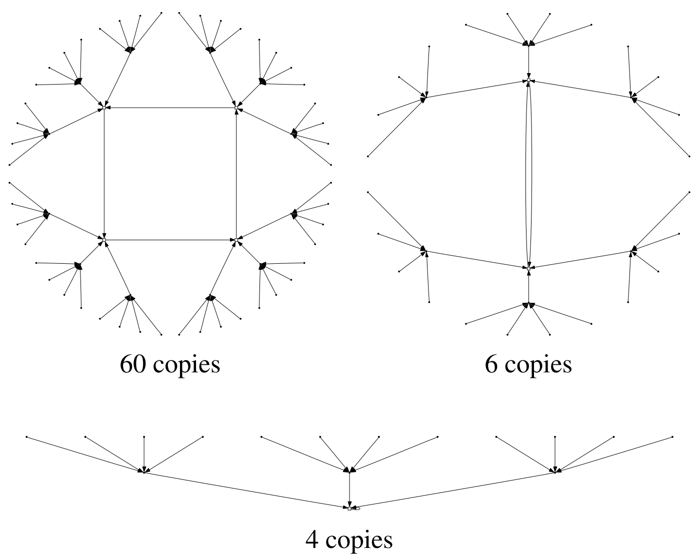

FIGURE 5
Transition graph for ECA rule 90 on 12 cells.

It follows almost immediately from the definitions that α is an automorphism of the transition graph if and only if

$$\alpha \circ F = F \circ \alpha \,, \tag{11}$$

where F is the CA's global map. If

$$s_0, s_1, s_2, s_3, \ldots \tag{12}$$

is a sequence of configurations visited by a CA on successive time steps, it follows that applying α to the initial configuration yields the sequence

$$\alpha(s_0), \alpha(s_1), \alpha(s_2), \alpha(s_3), \ldots \,. \tag{13}$$

In other words, α is applied to each configuration in the sequence, but the sequence is otherwise unchanged. Thus, borrowing terminology from physics, α can be considered a "symmetry" of the CA's global dynamics.

A *predecessor* of a configuration s is a configuration p such that $F(p) = s$. When constructing the transition graph, it is often preferable to begin at a configuration on a cycle and "work backwards", iteratively finding predecessors until that particular basin has been exhausted, than to compute the adjacency matrix by iterating through the entire configuration space. Wuensche and Lesser [8] describe the process in full.

A configuration is *reachable* if it has at least one predecessor, and *unreachable* otherwise. Unreachable configurations are sometimes referred to as *Garden of Eden* configurations. The unreachable configurations correspond to the leaf vertices in the trees of the transition graph.

3 TRANSITION GRAPHS FOR LINEAR CAS

This section presents some results which restrict the possible forms of the transition graph for a linear CA.

Lemma 1 ([1, Lemma 3.3])**.** *In the transition graph for a linear CA, the trees rooted at vertices in cycles form a single isomorphism class.*

Figures 4 and 5 illustrate this result. Note that this result applies to *all* trees in a given transition graph, whether rooted at vertices in the same cycle or in different cycles. An immediate consequence of this result is that two basins are isomorphic if and only if their cycles have the same length.

Now that we know the trees are all identical, the following result gives us some insight into their structure:

Lemma 2. *All reachable configurations of a linear CA have the same number of predecessors.*

Proof. The zero configuration is always reachable (since $F(0) = 0$), so it suffices to show that any configuration has the same number of predecessors as the zero configuration.

By additive superposition, for all configurations u, v we have $F(u) = F(v)$ if and only if $u = v+q$ for some predecessor q of the zero configuration [1, Lemma 3.2]. Let w be a reachable configuration, and choose a predecessor v of w. Then the predecessors of w are the configurations

$$\{u \;:\; F(u) = F(v) = w\} = \{v + q \;:\; F(q) = 0\} \; . \tag{14}$$

Note that $v + q = v + q'$ implies $q = q'$, so each value of q yields a unique element of this set. In other words, the predecessors of w are in one-one correspondence with the possible values of q, which themselves are the predecessors of the zero configuration. □

Suppose that the zero configuration has p predecessors. If we restrict our attention to a tree in the transition graph, we see that all non-root non-leaf vertices have in-degree p. However, one of the predecessors of the root vertex is the preceding vertex in its cycle, so within the tree the root vertex has in-degree $p - 1$. This can be seen in Figure 5, where $p = 4$.

What are the cycle lengths? We do not know of a general answer to this question, but we do have the following result:

Lemma 3 ([1, Lemma 3.4])**.** *For a linear CA on N cells, let Π_N be the length of the cycle reached from an initial configuration in which a single cell has state 1 and all other cells have state 0. Then Π_N is the maximal cycle length; furthermore, all cycle lengths are factors of Π_N.*

Note that the initial configuration described is not necessarily on the cycle itself; indeed, there are linear CAs in which such a configuration is unreachable.

4 COUNTING AUTOMORPHISMS

In [3] we give recursive expressions for the number of automorphisms of a transition graph for a CA in general, in terms of the sizes of isomorphism classes of certain families of subgraph. The results in this section can be considered as special cases of the results in [3], although we prove them here from first principles.

Theorem 4. *Consider a transition graph for a linear CA, in which the distinct cycle lengths are $l_1, \ldots, l_k$, and there are m_i cycles of length l_i. Suppose that each of the trees in this transition graph has A_T automorphisms. Then the transition graph has*

$$A_G = \prod_{i=1}^{k} \left(m_i! \times l_i^{m_i} \times A_T^{l_i m_i} \right) \tag{15}$$

automorphisms.

Proof. An automorphism of the transition graph is a permutation of the basins which preserves their isomorphism classes, composed with an automorphism on each of the basins. Since all trees rooted at cycle vertices are isomorphic, two basins are isomorphic if and only if they have the same cycle length; thus the isomorphism classes have sizes $m_1, \ldots, m_k$, and the total number of permutations is

$$P_B = \prod_{i=1}^{k} m_i! \,. \tag{16}$$

An automorphism of a basin is a "rotation" of the basin's cycle which preserves isomorphism classes of the trees rooted on the cycle, composed with an automorphism on each of the trees. Since all trees are isomorphic, any of the l_i rotations will suffice, and so the number of automorphisms of a basin whose cycle length is l_i is

$$A_B(l_i) = l_i A_T^{l_i} \,. \tag{17}$$

Thus the total number of automorphisms of the transition graph is

$$A_G = P_B \times \prod_{i=1}^{k} A_B(l_i)^{m_i} \tag{18}$$

$$= \prod_{i=1}^{k} \left(m_i! \times \left(l_i A_T^{l_i} \right)^{m_i} \right) \tag{19}$$

$$= \prod_{i=1}^{k} \left(m_i! \times l_i^{m_i} \times A_T^{l_i m_i} \right) \tag{20}$$

as required. □

Note that the overall exponent of A_T in A_G is $\sum_{i=1}^{k} l_i m_i$, which is the number of configurations which appear in cycles.

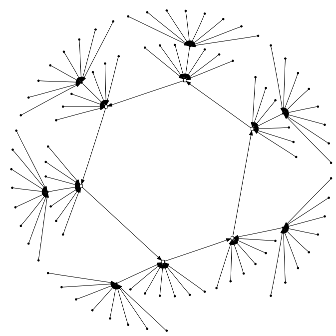

FIGURE 6
A basin in the transition graph for the linear CA with state set $\mathbb{Z}_2$ and local rule

$$f(x_{-3}, \dots, x_3) = x_{-3} + x_0 + x_2 + x_3\,, \tag{21}$$

on $N = 12$ cells. The entire transition graph comprises 40 isomorphic copies of this basin, five basins with cycle length 3, and one basin with cycle length 1. The trees are unbalanced. For example, the configurations $u = 111111111111$ and $v = 001100110011$ are predecessors of the zero configuration 000000000000, and neither u nor v is on the cycle. Configuration u is reachable (one of its predecessors is 000100010001), but v is unreachable, as are the other five non-zero predecessors of the zero configuration.

Say that a tree is *balanced* if all leaf vertices are the same distance from the root. Trees in transition graphs for linear CAs are not always balanced, although empirical investigation shows counterexamples to be rare enough to make balanced trees worth studying. The trees shown in Figures 4 and 5 are balanced; Figure 6 shows an example of a transition graph with unbalanced trees.

The *depth* of a balanced tree is the distance from the root to a leaf. The trees shown in Figure 4 have depth 1, and those shown in Figure 5 have depth 2.

Theorem 5. *Consider a balanced tree of depth d in which the root vertex has in-degree $p-1$, and all other non-leaf vertices have in-degree p. Such a tree has*

$$A_T(d,p) = \frac{p!^{p^{d-1}}}{p} \tag{22}$$

automorphisms.

Proof. If $d = 0$, the tree consists of a single vertex (the root vertex). Clearly this vertex has zero in-degree, and so we must have $p = 1$. (Indeed, the converse of this argument also holds, so that $d = 0$ if and only if $p = 1$.) This "tree" has only a single automorphism (the identity), and

$$A_T(0, 1) = \frac{1!^{1^{-1}}}{1} = 1 \tag{23}$$

as required.

Suppose $d > 0$, and proceed by induction on d. If $d = 1$, then the tree has $p - 1$ leaf vertices, each an immediate predecessor of the root vertex. The automorphisms of this tree are precisely the permutations of the leaves, of which there are $(p - 1)!$. Substituting $d = 1$ into Equation 22 yields

$$A_T(1, p) = \frac{p!^{p^0}}{p} = \frac{p!}{p} = (p - 1)!\,, \tag{24}$$

as required.

Now let $d > 1$, and assume, as an inductive hypothesis, that

$$A_T(d - 1, p) = \frac{p!^{p^{d-2}}}{p}\,. \tag{25}$$

Denote by T_d the tree, as described in the statement of the theorem, of depth d. Now T_d can be obtained from T_{d-1} by taking each leaf v in T_{d-1}, and adding p new leaves as immediate predecessors of v (so that v is no longer a leaf but the root of a subtree of depth 1).

An automorphism of T_d is an automorphism of T_{d-1} composed with an automorphism of each of the new subtrees added in transforming T_{d-1} into T_d. There are $(p - 1)p^{d-2}$ such subtrees, each having $p!$ automorphisms, and so T_d has

$$A_T(d, p) = A_T(d - 1, p) \times p!^{(p-1)p^{d-2}} \tag{26}$$

$$= \frac{p!^{p^{d-2}} \times p!^{(p-1)p^{d-2}}}{p} \tag{27}$$

$$= \frac{p!^{(1+p-1)p^{d-2}}}{p} \tag{28}$$

$$= \frac{p!^{p^{d-1}}}{p} \tag{29}$$

automorphisms, as required. □

5 EXAMPLE: ELEMENTARY RULE 90

Recall that rule 90 is the linear ECA with local rule

$$f(x_{-1}, x_0, x_1) = x_{-1} + x_1 \, . \tag{30}$$

The properties of rule 90 are studied extensively by Martin et al [1]; in this section, we use these properties to derive an expression for the numbers of automorphisms in rule 90's transition graphs.

The results in Section 3 describe some of the properties of the transition graph for a linear CA; for rule 90, we can go further than this, and can completely describe the trees in the transition graphs:

Lemma 6. *Let N be odd. The transition graph for rule 90 on N cells has the following properties:*

1. *All trees consist of a single edge [1, Theorem 3.3], and thus are balanced and have depth $d = 1$;*
2. *Reachable configurations have two predecessors, so $p = 2$ [1, Theorem 3.2];*
3. *The number of vertices in cycles is 2^{N-1}, i.e. precisely half of all vertices [1, Corollary to Theorem 3.3].*

Lemma 7. *Let N be even, and let $D_2(N)$ be the largest power of 2 which divides N:*

$$D_2(N) = \max \left\{ 2^j \ : \ 2^j | N \right\} \, . \tag{31}$$

The transition graph for rule 90 on N cells has the following properties:

1. *All trees rooted at vertices in cycles are balanced and have depth $d = \frac{1}{2} D_2(N)$ [1, Theorem 3.4];*
2. *Reachable configurations have four predecessors, so $p = 4$ [1, Theorem 3.2];*
3. *The number of vertices in cycles is $2^{N - D_2(N)}$ [1, Corollary to Theorem 3.2].*

The only elements missing for a complete description of the transition graph are the cycle lengths and their multiplicities. As far as we are aware, there is no simple expression for these; however, Martin et al [1] describe an algorithm which allows them to be computed. We shall not discuss the details of this algorithm, but it is essential in producing the numerical results described in Section 6.

The following theorem is obtained from Theorems 4 and 5 in the cases described in Lemmas 6 and 7:

Theorem 8. *Suppose that, on some number of cells N, rule 90 has cycles of distinct lengths $l_1, \ldots, l_k$, with m_i cycles of length l_i. Let*

$$A_T^c = \begin{cases} 1 & \text{if } N \text{ is odd} \\ 24^{2^{N-2}}/4^{2^{N-D_2(N)}} & \text{if } N \text{ is even}, \end{cases} \tag{32}$$

where $D_2(N)$ is as defined in Equation 31. Then the transition graph for rule 90 on N cells has

$$A_G = \left(\prod_{i=1}^{k} m_i! \times l_i^{m_i}\right) \times A_T^c \tag{33}$$

automorphisms.

Proof. By Theorem 4, it suffices to show that

$$A_T^c = \prod_{i=1}^{k} A_T(d,p)^{l_i m_i}, \tag{34}$$

where $A_T(d,p)$ is as defined in Theorem 5, and d and p are chosen appropriately. Note that

$$c = \sum_{i=1}^{k} l_i m_i \tag{35}$$

is the total number of vertices on cycles, and so

$$\prod_{i=1}^{k} A_T(d,p)^{l_i m_i} = A_T(d,p)^c. \tag{36}$$

If N is odd, Lemma 6 gives $d = 1$ and $p = 2$. Now

$$A_T(1,2) = \frac{2!^{2^0}}{2} = 1. \tag{37}$$

Indeed, it is easy to see that a tree consisting of a single edge admits only a single automorphism, namely the identity. Thus we have $A_T(1,2)^c = 1$ as required, regardless of the value of c.

If N is even, Lemma 7 gives $d = \frac{1}{2}D_2(N)$ and $p = 4$, so

$$A_T(d,p) = \frac{4!^{4^{\frac{1}{2}D_2(N)-1}}}{4} \tag{38}$$

$$= \frac{24^{2^{D_2(N)-2}}}{4} \,. \tag{39}$$

Lemma 7 also gives $c = 2^{N-D_2(N)}$, so

$$A_T(d,p)^c = \frac{24^{2^{D_2(N)-2} \times 2^{N-D_2(N)}}}{4^{2^{N-D_2(N)}}} \tag{40}$$

$$= \frac{24^{2^{N-2}}}{4^{2^{N-D_2(N)}}} \tag{41}$$

as required. □

6 NUMERICAL RESULTS FOR RULE 90

Theorem 8 yields the number of automorphisms for rule 90 on N cells, if the cycle lengths l_i and multiplicities m_i are known. Martin et al [1] give an algorithm for computing the l_is and m_is. This allows us to compute numbers of automorphisms for any value of N, although naturally some values of N require more computation than others.

Some results are shown in Figure 7. Observe that the relationship between the double logarithm of the number of automorphisms and N is approximately linear when N is even; fewer automorphisms occur when N is odd, and there seems to be a lower bound achieved on some (but not all) prime values of N.

Martin et al [1] define the *suborder function* of 2 modulo N, denoted $\mathrm{sord}_N(2)$, by

$$\mathrm{sord}_N(2) = \begin{cases} \min\left\{ j \;:\; 2^j \equiv \pm 1 \mod N \right\} & \text{if } N \text{ is odd} \\ 0 & \text{if } N \text{ is even.} \end{cases} \tag{42}$$

The suborder function is plotted in Figure 8. The suborder function $\mathrm{sord}_N(2)$ has an upper bound of $\frac{1}{2}(N-1)$, achieved on some (but not all) prime values of N.

Comparing Figure 8 with Figure 7, a correlation is apparent: values of N which yield many automorphisms give small values of $\mathrm{sord}_N(2)$, and vice versa. Indeed, it can be verified numerically that

$$\log_{10}\log_{10} A(90,N) \approx 0.30N - 0.28\,\mathrm{sord}_N(2) - 0.04 \,. \tag{43}$$

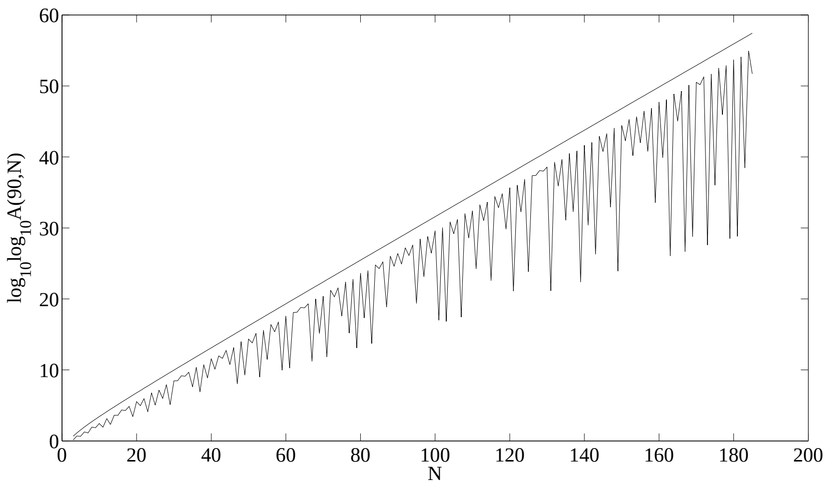

FIGURE 7
Plot of $\log_{10}\log_{10} A(90, N)$ (lower line) against N, for $3 \le N \le 185$. For comparison, $\log_{10}\log_{10} A(204, N) = 2^N!$ is also plotted (upper line).

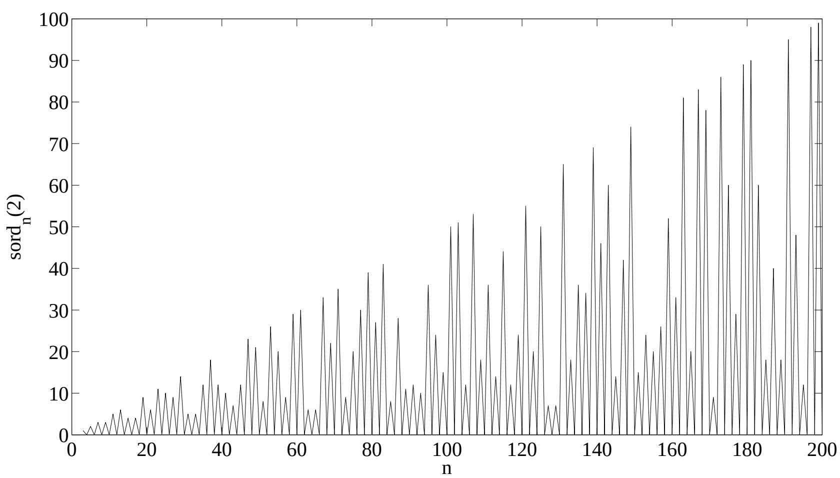

FIGURE 8
Plot of the suborder function of 2 modulo n.

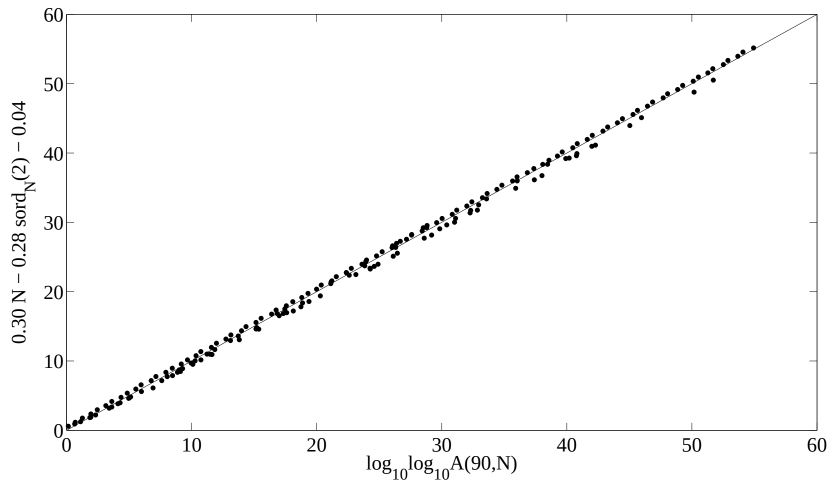

FIGURE 9
Plot of the two sides of Equation 43. The diagonal "$y = x$" line is plotted for comparison.

The error in this approximation is illustrated in Figures 9 and 10.

How does the suborder function relate to rule 90? Let Π_N be the length of the cycle reached from the initial configuration with a single cell in state 1, as defined in Lemma 3. For rule 90, we have the following results:

Lemma 9. *1. If N is a power of 2, then $\Pi_N = 1$ [1, Lemma 3.5];*

2. If N is even but not a power of 2, then $\Pi_N = 2\Pi_{N/2}$ [1, Lemma 3.6];

3. If N is odd, then Π_N is a factor of $2^{\mathrm{sord}_N(2)} - 1$ [1, Theorem 3.5].

Furthermore, Martin et al [1] observe that $\Pi_N = 2^{\mathrm{sord}_N(2)} - 1$ for the majority of odd N, the first few exceptions being

$$N = 37, 95, 101, 141, 197, 199, 203, \dots . \tag{44}$$

Theorem 8 states that, if N is odd, the number of automorphisms is given by

$$A_G = \left(\prod_{i=1}^{k} m_i! \times l_i^{m_i} \right) . \tag{45}$$

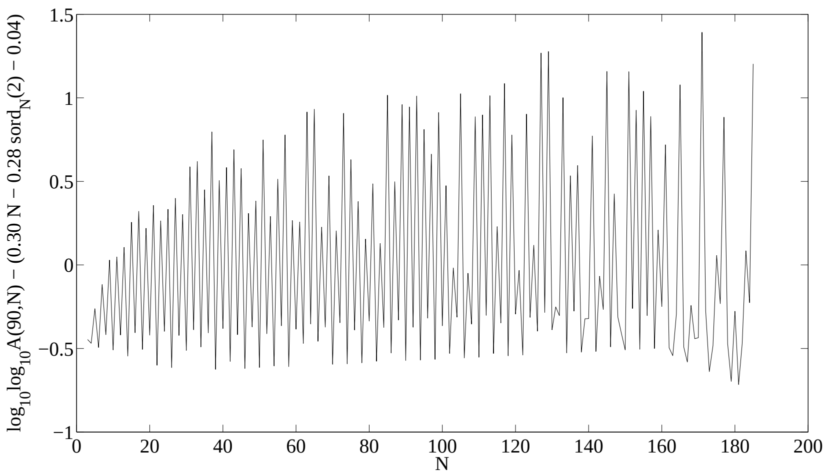

FIGURE 10
Plot of the difference between the two sides of Equation 43 against N.

The l_is are factors of Π_N, which in turn is a factor of $2^{\mathrm{sord}_N(2)} - 1$. The m_is have no immediately apparent relationship with $\mathrm{sord}_N(2)$. Thus it is somewhat mysterious that, of all the quantities which determine the cycle lengths and multiplicities, it should be $\mathrm{sord}_N(2)$ which, along with N, ends up dominating the expression for the number of automorphisms.

7 CONCLUSION

We have argued that automorphisms of transition graphs are, in at least two different senses, "symmetries" of a CA. Our original aim with this work was to investigate whether the "amount" of symmetry in a CA is in any way correlated with its qualitative dynamics (or its "Wolfram class" [5]). As discussed in [3], the partial correlation we have found is not so much with the amount of symmetry itself, but with the way in which this amount varies with the number of cells.

Upon first inspection, this variation seems to be more erratic for the linear CAs than for the nonlinear CAs: in Figure 1, it is the linear ECAs which display the most dramatic "zig-zag" patterns. Intuitively, it seems that the non-trivial linear CAs are much more sensitive to changes in the number of cells. Perhaps this is not so surprising: if N is a power of 2, it can be shown that the only attractor for rule 90 on N cells is a fixed point, namely the zero configuration. So there are cases where a small change in the number of cells significantly changes the long-term qualitative dynamics: rule 90 on 128 cells is in Wolfram's class 1 (its long-term behaviour is homogeneous), whereas on 127 or 129 cells it is in class 3 (its long-term behaviour is "chaotic").

Of course, the range of the data shown in Figure 1 is insufficient to draw any real conclusions; it is only when the numerical results are extended to much larger values of N (Figure 7) that a pattern becomes apparent. The numerical results also show the argument above regarding power-of-2 values of N to be somewhat misleading: in fact, it is around the powers of 2 that we observe the *smallest* fluctuations in numbers of automorphisms.

The data plotted in Figure 2 (a) suggest that the numbers of automorphisms for other linear CAs exhibit similar behaviour to that for rule 90. To test this hypothesis, it is desirable to generalise the results of Section 5 to other linear CAs. In particular, to apply Theorems 4 and 5 (which take care of the common, though not universal, case where the trees are balanced), four quantities must be known:

1. The number of predecessors of a reachable configuration;
2. The depth of a tree (or equivalently, the maximum transient length);
3. The set of cycle lengths;
4. The multiplicity of each cycle length.

While we have not yet done so, it seems plausible that such results can be found for individual linear CAs via similar techniques to those used by Martin et al [1] for rule 90. Whether results can be found which apply to *all* linear CAs remains to be seen.

This work is based entirely on two facts about transition graphs for linear CAs: all the trees are isomorphic, and every non-leaf vertex has the same in-degree. These properties certainly do not hold in general, so the approach described here is not applicable to nonlinear CAs. To extend those numerical results to larger N, a completely different approach is needed.

ACKNOWLEDGEMENT

We thank the anonymous referees for their helpful comments. This work is funded by an EPSRC DTA award.

REFERENCES

[1] Olivier Martin, Andrew M. Odlyzko, and Stephen Wolfram. (1984). Algebraic properties of cellular automata. *Communications in Mathematical Physics*, 93:219–258.

[2] Edward Powley and Susan Stepney. (2008). Automorphisms of transition graphs for a linear cellular automaton. In Andrew Adamatzky, Ramon Alonso-Sanz, Anna Lawniczak, Genaro Juarez Martinez, Kenichi Morita, and Thomas Worsch, editors, *proceedings of Automata 2008: Theory and Applications of Cellular Automata*, pages 55–68. Luniver Press.

[3] Edward Powley and Susan Stepney. (2008). Automorphisms of transition graphs for elementary cellular automata. In *proceedings of Automata 2007: 13th International Workshop on Cellular Automata*. To appear in Journal of Cellular Automata.

[4] Stephen Wolfram. (1983). Statistical mechanics of cellular automata. *Reviews of Modern Physics*, 55(3):601–644.

[5] Stephen Wolfram. (1984). Universality and complexity in cellular automata. *Physica D*, 10:1–35.

[6] Stephen Wolfram. (1986). Random sequence generation by cellular automata. *Advances in Applied Mathematics*, 7:123–169.

[7] Stephen Wolfram. (1986). Tables of cellular automaton properties. In *Theory and Applications of Cellular Automata*, pages 485–557. World Scientific.

[8] Andrew Wuensche and Mike Lesser. (1992). *The Global Dynamics of Cellular Automata: An Atlas of Basin of Attraction Fields of One-Dimensional Cellular Automata*. Addison Wesley.