

Decomposition of Additive CA

Klaus Sutner

Carnegie Mellon University

Pittsburgh, PA 15213

Abstract

We consider finite additive cellular automata with fixed and periodic boundary conditions as endomorphisms over pattern spaces. A characterization of the nilpotent and regular parts of these endomorphisms is given in terms of their minimal polynomials. We determine the generalized eigenspace decomposition, and describe relevant cyclic subspaces in terms of symmetries. As an application, we calculate the lengths and frequencies of limit cycles in the transition diagram of the automaton.

1 Introduction

Questions relating to the evolution of variable size configurations on a finite cellular automaton are in general PSPACE-hard. Classification problems, such as the question whether all configurations evolve to a fixed point are even undecidable. Here we refer to the uniform version of the problem, where the local rule of the cellular automaton is fixed, but one considers grids of all finite sizes. The root for all these computational hardness properties is, of course, the fact that one-dimensional cellular automata are computationally universal, see [8, 12, 13]. At the other end of the spectrum lie additive cellular automata. Here the evolution of a configuration is predictable in the sense that it is not necessary to explicitly compute t steps in the evolution of a pattern X under some global rule ρ to determine $\rho^t(X)$. For example, if the global rule of the automaton is expressed by matrix multiplication, $\rho^t(X)$ can be computed in time polynomial in the size of the grid and $\log t$. Typical examples of such rules in the one-dimensional case are elementary rules 90 and 150, corresponding to the exclusive or of the two neighbors of a cell, and the exclusive or of the neighbors plus the center cell. Another minor variation concerns the type of boundary conditions. We will refer to all such rules generically as rule σ when details are irrelevant.

σ -automata were studied in great detail in [10] using binary polynomials as the main algebraic tool. The authors represent both the linear operator and the configurations as binary polynomials in some suitable quotient ring $\mathbb{F}_2[x]/(\tau)$. For example, rule 90 with cyclic boundary conditions on a grid of length n can be represented by multiplication with $x^{-1} + x$ modulo $\tau = x^n + 1$. A wealth of structural information about the transition diagram of the rule can then be obtained via the theory of finite fields. As it turns out, the properties of σ on a finite grid of size n depend strongly on number theoretic properties of n . For example, reversibility and in fact the co-rank of σ can be determined from simple divisibility properties of n . All these calculations can be carried out in time polynomial in $\log n$. The question arises which properties of the transition diagram can be determined efficiently, and in particular without recourse to matrix algebra. Note that since the number of configurations is 2^n one can not hope for time complexity polynomial in $\log n$ in general. As we will see, determining the cycle structure of the diagram involves the factorization of n , the

factorization of certain binary polynomials of degree n , and the computation of the period of their irreducible factors. Möbius inversion can be used to avoid part of the factorization of polynomials, but it seems unlikely that one can dispense with any of these computational tools.

In this paper we construe σ as an endomorphism on an n -dimensional vector space $\mathbf{2}^n$ over \mathbb{F}_2 , the Galois field with 2 elements. We refer to these vector spaces as *pattern spaces*. This allows us to exploit the self-adjointness of the global rule, viewed as an endomorphism. For example, the patterns that appear on limit cycles are precisely those that are orthogonal to the kernel of the endomorphism, which produces a decomposition of pattern space into invariant, orthogonal subspaces $V = K \oplus E$, see [11]. Since the minimal polynomials of the σ operators are known, see [14], we can push this decomposition further to obtain a detailed description of the elementary divisor spaces. As we will see, there are chains of natural subspaces in the elementary divisor spaces that correspond to the symmetries of the cellular automaton. Given the factorization of the minimal polynomial and the periods of the irreducible factors, one can easily determine the complete structure of the transition diagram. Using only polynomial arithmetic one can determine bases for the elementary divisor subspaces, as well as the relevant σ -cyclic subspaces, and the order of σ when restricted to these spaces. Thus, given a pattern X we can calculate the transient length and the length of the limit cycle in the orbit of X by calculating the representation of X with respect to these bases.

As one might suspect, the answers vary slightly depending on whether the center cell is included or excluded, corresponding to rules 150 and 90, respectively. There are also slight differences depending on whether cyclic or fixed boundary conditions are used. The common tool to handle these four types of systems is a version of binary Fibonacci or Chebyshev polynomials, first introduced in [11] for the purpose of analyzing two-dimensional σ -automata. Define the π -polynomials by the following second order homogeneous recurrence over $\mathbb{F}_2[x]$:

$$\begin{aligned}\pi_0 &= 0, \\ \pi_1 &= 1, \\ \pi_n &= x \cdot \pi_{n-1} + \pi_{n-2}.\end{aligned}\tag{1}$$

The π -polynomials can be computed easily either by using a logarithmic depth recursion, much as for the Fibonacci numbers, or by exploiting an explicit description of the coefficients of these polynomials. For example,

$$\pi_n(x) = \sum_i \binom{n+i}{2i+1} x^i \pmod{2}.$$

Using Lucas' theorem, the computation of the binomial coefficients modulo 2 can be handled by comparing the binary expansions of $n+i$ and $2i+1$.

The minimal polynomials of the σ operators can be expressed easily in terms of these polynomials. Moreover, the π -polynomials have a relatively simple multiplicative structure, and there is a uniform description of the factorization of these polynomials. Hence we can determine the elementary divisors of σ , and the corresponding decomposition of the pattern space. For the subspaces E so obtained, the order of the restriction $\sigma \upharpoonright E$ can be expressed as the period of the corresponding irreducible polynomial, multiplied by a power of 2 which is determined by the exponent of the corresponding irreducible factor of the minimal polynomial in question.

A good part of the discussion below is just the study of linear operators over pattern spaces, and uses only general tools from algebra; see for example [7, 6]. Background material on finite

fields, irreducible binary polynomials and shift-register sequences can be found in [4, 9, 2] and will be used without further comment. In order to give a more detailed analysis of σ -automata, one also has to consider the geometry of a pattern space together with σ . More precisely, we consider simulations, monomorphism of pattern spaces that commute with σ . In the case where the domain and codomain coincide we are dealing with automorphisms that commute with the shift. Of particular interest are geometric automorphisms: for fixed boundary conditions there is only one non-trivial such automorphism, namely reflection. For cyclic boundary conditions on the other hand, the geometric automorphism group is the dihedral group, and is generated by reflection and rotation. Invariance, or lack thereof, of various subspaces under these automorphisms is the key element in determining the structure of the transition diagram of σ in great detail.

We have limited our discussion here to the characteristic 2 case, though the results can be carried over, mutatis mutandis, to other prime fields. See [3] for a discussion of linear operators in this more general context.

This paper is organized as follows. In section 2 we introduce terminology and notation, and briefly recap some well-known results about the transition diagram of σ . We also provide the necessary background knowledge about π -polynomials. In section 3 we determine the elementary divisor decomposition of pattern space, and show how to compute the order of the restriction of σ to the generalized eigenspaces that make up the decomposition. The next section discusses simulations and shows how to exploit them to obtain a basis for the eigenspaces. The results are then applied in section 5 to obtain a complete description of the cycle lengths in the diagram. We will also describe the subspaces of the divisor spaces in terms of symmetries. In the last section we conclude by stating a few open problems relating to the analysis of σ -automata.

2 Minimal Polynomials and the Transition Diagram

Whenever necessary, we will indicate the boundary conditions by a subscript, and distinguish between rule 90 and 150 by a superscript. Thus, σ_c^- refers to rule 90 with cyclic boundary conditions and σ_z^+ refers to rule 150 with fixed boundary conditions. For emphasis we may indicate the size of the grid as in $\sigma_c^-(n)$. Correspondingly the pattern spaces together with the linear operators will be denoted by $\mathcal{C}_c^-(n) = \langle \mathbf{2}^n, \sigma_c^- \rangle$ and so forth. In order to apply the elementary divisor decomposition machinery to pattern spaces we need an explicit description of the minimal polynomials of the σ -operators. The following result and is established in [14].

Theorem 2.1 *The minimal polynomial of $\sigma_z^-(n)$ is π_{n+1} , and the minimal polynomial of $\sigma_c^-(n)$ is $\sqrt{x \pi_n}$ for n even, and $x\sqrt{\pi_n}$ for odd n .*

Thus, the minimal polynomial for $\sigma_z^-(n)$ has degree n , but the minimal polynomial of $\sigma_c^-(n)$ has degree $\lceil n/2 \rceil$. The minimal polynomials for the associated maps σ_z^+ and σ_c^+ are obtained simply by applying the involution $x \mapsto x + 1$ to these polynomials. We will denote the involution $x \mapsto x + 1$ on $\mathbb{F}_2[x]$ by a superscript $+$. Hence the minimal polynomial for $\sigma_z^+(n)$ is $\pi_{n+1}^+(x) = \pi_{n+1}(x + 1)$ and the minimal polynomial for $\sigma_c^+(n)$ is $(x + 1)\pi_{n/2}^+(x)$ or $(x + 1)\sqrt{\pi_n^+}$, depending on the parity of n .

It is also shown in [14] that the π -polynomials admit a decomposition into *critical factors* ρ_d as follows.

$$\pi_n(x) = x^{2^k-1} \prod_{d|m} \rho_d^{2^k}(x) = x^{2^k-1} \prod_{d|m} \rho_d(x^{2^k}) \quad (2)$$

where $n = 2^k \cdot m$, m odd. One can show that the degree of ρ_d is $\varphi(d)$ where φ denotes Euler's totient function. The polynomials ρ_d are products of squares of certain irreducible polynomials. In many cases, $\rho_d = \tau^2$ where τ is irreducible, but there are critical factors that are comprised of several irreducible polynomials. The first example is $\rho_{17} = (1 + x + x^4)^2 (1 + x + x^2 + x^3 + x^4)^2$. This will be a minor obstruction in our decomposition of pattern spaces later.

For our purposes here, it is more convenient to write the π -polynomials as

$$\pi_n(x) = x^a (x + 1)^b \prod_{i=1}^r \tau_i^c(x) \quad (3)$$

The second linear irreducible polynomial $x + 1$ is due to the critical factor ρ_3 , and only appears when n is a multiple of 3. The exponents are determined as follows. We write $D_2(n)$ for the largest power of 2 which divides n . To simplify the expressions below, let us adopt Knuth's convention to write $[\varphi]$ for the Boolean value, interpreted as 0 or 1, of any unary predicate φ , see [5]. Then

$$\begin{aligned} a &= D_2(n) - 1 \\ b &= [3 | n] \cdot 2 D_2(n) \\ c &= [n \neq 2^k, 3 \cdot 2^k] \cdot 2 D_2(n) \end{aligned}$$

The minimal polynomials for σ^+ are obtained simply by switching a and b , and applying the involution to the irreducible factors. While the exponents are easily determined from n , the irreducible factors τ_i^+ are somewhat more difficult to describe. Indeed, it is shown in [14] that every irreducible polynomial occurs as a factor of some π -polynomial. However, there appears to be no easy way to determine the least n for which a given irreducible polynomial τ divides π_n . Nonetheless, equation (2) in conjunction with Möbius inversion can be used to obtain the critical factors by plain polynomial arithmetic. More precisely, $\rho_m = \prod_{d|m} \pi_{m/d}^{\mu(d)}$, where m is odd and μ denotes the Möbius function. However, as already pointed out, a critical factor may be the product of the squares of 2 or more irreducible polynomials, so we still need a factoring algorithm to determine the elementary divisors of σ .

We will refer to the functional digraph of the global map of a σ -automaton as the *transition diagram* of the automaton. Thus, the vertex set of the diagram is a pattern space $V = \mathbf{2}^n$, and there is an edge from vertex v to vertex u if $\sigma(v) = u$. It is clear from the definition that the components of the transition diagram of σ are unicyclic. In particular, the co-orbit of $\mathbf{0}$ is a tree rooted at the fixed point $\mathbf{0}$, see [10, 11]. We denote this component K . It follows from the linearity of σ that the branching factor in K is the corank of σ . Moreover, from the results in [1] and [14], it follows that the tree is completely balanced (i.e., all leaves occur at the same level). The height of the tree is plainly the nilpotency index of σ , i.e., the least number k such that $\text{cork } \sigma^k = \text{cork } \sigma^{k+1}$. All the other trees in the diagram are isomorphic copies of the co-orbit of $\mathbf{0}$.

One can think of the whole diagram as a deterministic product automaton over a one-letter alphabet. The two factor machines are the co-orbit of $\mathbf{0}$, and the limit cycles. In terms of the linear operators, this corresponds to the standard decomposition of pattern space as $V = K \oplus E$ where $\sigma \upharpoonright K$ is nilpotent and $\sigma \upharpoonright E$ is an automorphism, see [7]. The two subspaces are sometimes

referred to as the Fitting-null and Fitting-one component, respectively. We will further decompose the regular part into natural subspaces $E = E_1 \oplus E_2$. For our purposes only those decompositions are of interest where the subspaces are σ -invariant. Recall that a subspace U of V is σ -invariant iff $\sigma(U) \subseteq U$, and σ -cyclic iff U is generated as a \mathbb{F}_2 -vector space by the σ -orbit of some $u \in U$. In other words, U is spanned by $\{\sigma^i(u) \mid 0 \leq i < d\}$ for some d . A better way of expressing these two conditions is to think of V as a $\mathbb{F}_2[x]$ module: the module operation is $f \cdot u = f(\sigma)(u)$. A σ -invariant \mathbb{F}_2 -subspace is none other than a $\mathbb{F}_2[x]$ -submodule. Likewise, a σ -cyclic \mathbb{F}_2 -subspace is a $\mathbb{F}_2[x]$ -cyclic submodule. A function $f : V \rightarrow V$ is *polynomially representable* in σ if there is a polynomial r such that $f(u) = r \cdot u$ for all $u \in V$.

We will show in the next section how to decompose the pattern space into a direct sum of σ -cyclic submodules. The dimension of a σ -cyclic submodule E is the degree of $\sigma \upharpoonright E$. Likewise, since these submodules are isomorphic as $\mathbb{F}_2[x]$ -modules to a quotient module of $\mathbb{F}_2[x]$ itself, we will be able to determine the order of the restriction of σ to these submodules. The geometric automorphism of reflection and rotation turn out to be polynomially representable, and will be helpful in analyzing the structure of the decomposition spaces.

3 The Elementary Divisor Decomposition

Since $\mathbb{F}_2[x]$ is a principal ideal domain, we can define the *order* of a pattern $u \in V$ to be the necessarily monic polynomial τ that generates the ideal of all annihilators of u : $(\tau) = \{r \in \mathbb{F}_2[x] \mid r \cdot u = 0\}$. Clearly, the order of any pattern divides the minimal polynomial of σ . For τ irreducible let V_τ be the subspace of all patterns whose order is a power of τ . Note that for two distinct irreducible polynomials τ and ρ we have $V_\tau \cap V_\rho = 0$, so we can obtain a decomposition as a direct sum. Hence, if $\tau_1, \tau_2, \dots, \tau_k$ are all the irreducible divisors of the minimal polynomial of σ , we have the natural decomposition

$$V = V_{\tau_1} \oplus V_{\tau_2} \oplus \dots \oplus V_{\tau_k}$$

where all the summands are in fact orthogonal due to the self-adjointness of σ .

If present, the subspace associated with $\tau = x$ represents the nilpotent part, and all the others represent the limit cycles, the digraph of the regular part. To get an actual count of the cycles of various lengths we need to consider subspaces of V_τ arising from the elementary divisors of σ . To this end, define the *generalized eigenspace* of f , for any polynomial $f \in \mathbb{F}_2[x]$, by

$$A(f) = \{u \in V \mid f \cdot u = 0\}.$$

It is easy to see that $A(f) = 0$ if, and only if, f and the minimal polynomial of σ are coprime. On the other hand, for coprime f and g we have $A(fg) = A(f) \oplus A(g)$, hence we will focus on eigenspaces of the form $A(\tau^e) \subseteq V_\tau$ where τ^e is a factor of the minimal polynomial of σ , τ irreducible. The dimension of $A(\tau^i)$ is $i \cdot \deg \tau$ for all $i \leq e$. The eigenspaces $A(\tau^e)$ may fail to be irreducible in the sense that they may admit further decomposition into σ -invariant components, see the discussion for cyclic boundary conditions below. However, the σ -cyclic subspaces determined there are themselves irreducible.

As a first step towards the elementary divisors decomposition, let us dispense with the nilpotent part K in $V = K \oplus E$. The nilpotency index is the maximum k such that x^k divides the minimal polynomial. Hence, the dimension of K as a \mathbb{F}_2 vector space is the co-rank of σ^k . The other

summand has the form $E = \bigoplus E_{i,j}$ where each of the subspaces is σ -cyclic and the minimal polynomials of $\sigma \upharpoonright E_{i,j}$ have the form $\tau_i^{e_{i,j}}$, $i = 1, \dots, s$ and $j = 1, \dots, k_i$. Here τ_i is a sequence of distinct irreducible polynomials. We may safely assume $e_{i,j} \geq e_{i,j+1}$, in which case the minimal polynomial of σ is the product $\tau_1^{e_{1,1}} \tau_2^{e_{2,1}} \dots \tau_s^{e_{s,1}}$. As we will see, there is a slight difference between the decompositions for fixed and periodic boundary conditions.

3.1 Operator σ^-

For fixed boundary conditions there is exactly one irreducible eigenspace for each one of the factors of the minimal polynomial.

Theorem 3.1 *Let $\pi_{n+1} = x^a \prod_{i=1}^r \tau_i^b$ be the minimal polynomial of $\sigma_z^-(n)$. Then the elementary divisor decomposition of V is $K \oplus E_1 \oplus \dots \oplus E_r$ where K has dimension a and E_i has dimension $b \deg \tau_i$. All the summands are pairwise orthogonal. Moreover, the summands are irreducible, and there are no other irreducible subspaces.*

Proof. The minimal polynomial of the nilpotent part of any linear operator is of the form x^e . Since the minimal polynomial of $\sigma_z^- = \nu + \alpha$ is the product of the respective minimal polynomials, it must be x^a .

Now consider one of the prime powers τ^b in the factorization of the minimal polynomial where τ is an irreducible polynomial and $b = 2 \cdot \eta_2(n+1)$.

There is an associated elementary divisor E such that the minimal polynomial of $\sigma \upharpoonright E$ is τ^b . E can be generated by choosing an arbitrary element g of V of order τ^b , and letting $E = \mathbb{F}_2[x] \cdot g$ as a cyclic $\mathbb{F}_2[x]$ -module. From the preceding remarks, E is σ -cyclic. Furthermore, the natural map from $\mathbb{F}_2[x]$ to E as a $\mathbb{F}_2[x]$ -module homomorphism has as kernel the ideal (τ^b) . Note that the \mathbb{F}_2 -dimension of E is none other than the least d such that there are coefficients c_i such that $\sum_{0 \leq i \leq d} c_i \sigma^i(g) = 0$. This is the same as asserting that τ^b must divide $\sum c_i x^i$.

Now note that V has dimension n , which is $a + b \sum \deg \tau_i$. It follows that there is exactly one elementary divisor for each irreducible factor of the minimal polynomial. Orthogonality follows from the fact σ is self-adjoint. Lastly, irreducibility is equivalent to being σ -cyclic for eigenspaces. It can be shown that the number I_j of irreducible subspaces of dimension $j \cdot d$ of an eigenspace $E = A(\tau^e)$, $d = \deg \tau$, is determined by

$$I_j = 1/d(\text{rk } \tau^{j+1}(\alpha) + \text{rk } \tau^{j-1}(\alpha) - 2 \text{rk } \tau^j(\alpha)), \quad (4)$$

where $\alpha = \sigma \upharpoonright E$, see [6]. Hence $N_e = 1$, but $N_j = 0$ for all $j < e$. □

From the argument in the proof we can see that E is isomorphic to $\mathbb{F}_2[x]/(\tau^b)$ as a $\mathbb{F}_2[x]$ -module. Note that in the theorem, $a = \eta_2(n+1) - 1$ and $b = \eta_2(n+1) + 1$. Hence, given the factorization of π_{n+1} , we can easily compute the dimensions of the σ -cyclic subspaces. Also note that a is the nilpotency index of the the linear map σ_z^- , it coincides with the dimension of the co-orbit of $\mathbf{0}$ since the co-rank of σ is 1 in the irreversible case.

The argument for cyclic boundary conditions is similar, but now the minimal polynomial is of lower degree. Using again equation (4), it turns out that there are 2 irreducible subspaces of $A(\tau^e)$ of dimension $e \cdot \deg \tau$, and no irreducible spaces of smaller dimension. The relationship between the two irreducible spaces can be described in terms of the actual geometry of the pattern space. To this end consider the two rotations $R, L : V \rightarrow V$ and the reflection $S : V \rightarrow V$. More precisely, $S(e_i) = e_{n-i+1}$ where e_1, \dots, e_n is the standard basis, $R(u)(i) = u(i+1)$ where index

$n + 1$ is interpreted as 1, and $L = R^{-1} = R^{n-1}$. The maps are clearly linear; as a matter of fact, $\sigma_c^- = R + L$. All three maps commute with σ_c^- , they are examples of auto-simulations, see section 4.2 below. Since R and S do not commute, it follows by the the theorem on bicommutants that neither S nor R are polynomially representable, see [6]. Hence, irreducible subspaces may well fail to be invariant under S and R . Indeed, we will show that for $A(\tau^e) = E_1 \oplus E_2$ we have $E_2 = R(E_1)$.

Theorem 3.2 *Let $x^a \prod_{i=1}^r \tau_i^b$ be the minimal polynomial of $\sigma_c^-(n)$. Then the elementary divisor decomposition of V is $K \oplus E_1 \oplus E_1' \oplus E_2 \dots \oplus E_r \oplus E_r'$ where K has dimension a and E_i and E_i' both have dimension $b \deg \tau_i$. The summands are pairwise orthogonal, except for E_i and E_i' . Moreover, the summands are irreducible, and there are no other irreducible subspaces.*

Proof. It only remains to verify that indeed for any generator u_0 of E_i we have $R(u_0) \in E_i'$, for all $i = 1, \dots, r$. Since $A(\tau_i^b) = E_i \oplus E_i'$ is invariant under R it suffices to show that $R(u_0)$ is not in E_i .

Assume for the sake of a contradiction that $R(u_0) \in E_i$, whence $R(u_0) = r \cdot u_0$ for some $r \in \mathbb{F}_2[x]$. But then $R(u) = r \cdot u$ for all $u \in E_i$. The linear map $r(\sigma) = r(L + R)$ is symmetric in R and L , so that $R(u) = L(u)$ for all $u \in E_i$. But the only patterns satisfying this equation lie in the kernel of σ , and we have the desired contradiction. \square

We can obtain a basis for the whole eigenspace by rotations of the generator rather than by applications of σ_c^- .

Corollary 3.1 *Assume the notation from the last theorem. The generalized eigenspace $A(\tau^b)$ has a basis of the form $\{R^i(g) \mid i < 2b \deg \tau\}$.*

Proof. As in the last proof, let g be any element of order τ^b . Since $A(\tau^b)$ is closed under R , the set $B = \{R^i(g) \mid i < 2d\}$, $d = b \deg \tau$, is certainly contained in the space, and it suffices to show that B is \mathbb{F}_2 -linearly independent. An easy induction shows that as \mathbb{F}_2 spaces

$$\text{span}(\sigma^i(g), \sigma^i(R(g)) \mid i < k) = \text{span}(L^i(g), R^j(g) \mid i < k, j \leq k).$$

Now assume that B is linearly dependent. Since generators are invariant under rotations, we have $\sum_{-d < i \leq d} c_i R^i(g)$ for some coefficients c_i . This contradicts the independence of $\{\sigma^i(g), \sigma^i(R(g)) \mid i < d\}$. \square

By a similar argument, $\{R^i(g) \mid i \leq 2d\}$ is linearly dependent. Using equation (1) it is easy to see that

$$R^i(u) = \pi_{i-1} \cdot u + \pi_i \cdot R(u)$$

for any u in pattern space and any integer i (assuming the obvious extension of the recurrence for the π -polynomials to all integers). For $u \in A(\tau^b)$ the π -polynomials are effectively computed modulo τ^b , which produces the appropriate values for d .

Note that since τ^c is the minimal polynomial of $\alpha = \sigma \upharpoonright E$, we can find a basis for E for which α is represented by the companion matrix of τ^c :

$$\begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & & & & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ a_0 & a_1 & a_2 & \dots & a_{d-1} \end{pmatrix}$$

where $\tau^c = \sum_{i < d} a_i x^i + x^d$. Thus, iterations of α can be expressed as a shift-register sequence. The characteristic and minimal polynomial of this matrix is easily seen to be τ^c . Of course, the natural geometry of pattern space is destroyed by the necessary base transformations.

Example

As an example, consider $\langle \mathbf{2}^{69}, \sigma_c^- \rangle$. The minimal polynomial here has the form

$$\begin{aligned} & x(1+x)(1+x^2+x^3+x^4+x^8+x^{10}+x^{11}) \\ & (1+x^5+x^6+x^9+x^{10}+x^{11}+x^{12}+x^{13}+x^{14}+x^{15}+x^{16}+x^{21}+x^{22}) \end{aligned}$$

Correspondingly, there are $1 + 2 \cdot 3 = 7$ spaces in the elementary divisor decomposition. The first is the kernel of σ_z^- and has dimension 1. Omitting the isomorphic copies obtained by rotation, the spaces have dimensions 1, 11 and 22, respectively. As we see shortly, the order of σ on these spaces is 1, 2047, and 4194303. Figure 1 shows bases for the decomposition spaces, again with the isomorphic copies omitted. For the sake of clarity, we have inserted blank rows between the basis vectors for each subspace.




Figure 1: Bases for decomposition spaces of $\langle \mathbf{2}^{69}, \sigma_c^- \rangle$

The picture was generated by performing row reduction on a basis obtained from a generator by iterating σ . Note that the basis for the 11-dimensional space is obtained by embedding reflections of the identity matrix. We will return to this issue shortly.

3.2 Operator σ^+

It is clear that the arguments from the last section carry over, mutatis mutandis, to σ^+ . Indeed, since $\sigma^+(u) = \sigma^-(u) + u$, the σ^+ -cyclic spaces are precisely the σ^- -cyclic spaces and have the same

generators. However, the minimal polynomials for σ^+ are slightly more complicated, and one has to be careful to account properly for irreducible factors of the form $x + 1$. For fixed boundary conditions the minimal polynomial of σ_z^+ on a grid of size n is

$$\pi_{n+1}^+(x) = \pi_{n+1}(x + 1) = x^b (x + 1)^a \prod_{i=1}^r \tau_i^c \quad (5)$$

where the τ_i are irreducible polynomials of degree at least 2, and the exponents $a, b, c \geq 0$ are as in equation (3). As corollaries to theorems 3.1 and 3.2 we obtain eigenspace decompositions for σ^+ .

Corollary 3.2 *Consider the minimal polynomial of $\sigma_z^+(n)$ as in equation (5). Then the elementary divisor decomposition of V is $K \oplus E_0 \oplus E_1 \dots \oplus E_r$ where K has dimension b , E_0 has dimension a , and $E_i, i > 0$, has dimension $b \deg \tau_i$. All the summands are pairwise orthogonal. They are also irreducible, and there are no other irreducible subspaces.*

Corollary 3.3 *As in equation (5), let $x^b (x + 1)^a \prod_{i=1}^r \tau_i^c$ be the minimal polynomial of $\sigma_c^+(n)$. The elementary divisor decomposition of V is $E_0 \oplus E_1 \oplus E'_1 \oplus E_2 \dots \oplus E_r \oplus E'_r$ where E_0 has dimension b , and E_i and E'_i both have dimensions a for the spaces associated with $x + 1$, and $c \deg \tau_i$ otherwise. The summands are pairwise orthogonal, except for E_i and E'_i . They are also irreducible, and there are no other irreducible subspaces.*

3.3 The Order of σ

We can now calculate the order of $\sigma \upharpoonright E$ for a generalized eigenspace E in the decomposition. For any element u in the regular part of the decomposition, define the *period* of u , in symbols $\text{per } u$, to be the least $i > 0$ such that $\sigma^i(u) = u$. Recall that the *period* of an irreducible polynomial τ is the least $p > 0$ such that $x^p + 1 = 0 \pmod{\tau}$. Thus, the period of a polynomial is the order of any of its roots in the multiplicative subgroup of a splitting field, and in particular divides $2^d - 1$ where d is the degree of τ . The period can be calculated by factoring $2^d - 1 = p_1^{e_1} \dots p_k^{e_k}$ and then finding the least exponents c_i such that $p = p_1^{c_1} \dots p_k^{c_k}$ satisfies $x^p + 1 = 0 \pmod{\tau}$. This requires essentially only a fast polynomial exponentiation algorithm with a fixed modulus τ , see e.g. [2]. Incidentally, for small-degree polynomials a brute force calculation based on shift-register sequences seems to be more efficient. We write $C_m \in \mathbb{F}_2[x]$ for the m th cyclotomic polynomial.

Lemma 3.1 *Consider a generalized eigenspace $E = A(\tau^e)$ where $\tau \neq x$ is irreducible. Let e' be the least power of 2 larger or equal to e and let $\alpha = \sigma \upharpoonright E$. Then the order of α is $e' \cdot p$ where p is the period of τ . Moreover, the period is determined by the condition $\tau \mid C_p$.*

Proof. Recall that $E = \mathbb{F}_2[x] \cdot u_0$ where u_0 is some generator of order τ^e . Hence, as an $\mathbb{F}_2[x]$ -module, E is isomorphic to the quotient $\mathbb{F}_2[x]/(\tau^e)$. The endomorphism σ here corresponds to multiplication by x , so we need to determine the least j such that $x^j = 1 \pmod{\tau^e}$. In other words, we need to find the period of τ .

In the special case $e = 1$ we are dealing with a Galois field, and it is clear that the order of x in the multiplicative subgroup is just m where $\tau \mid C_m$. For the general case, first note that x indeed lies in the multiplicative subgroup: since b is a power of 2 we can apply the Frobenius homomorphism

to obtain $x^{e'm} + 1 = (x^{e'})^m + 1 = 0 \pmod{\tau(x^{e'})}$. By the same argument, $C_{2^k r} = C_r$. But then the order of x cannot be less than $e' \cdot m$, for $x^r + 1$ can be divisible by τ only if $m = r$. \square

The exponents of the irreducible terms in the minimal polynomials of our σ -operators are all powers of two, with the exception of $x + 1$ for σ_z^+ . All the generators of E have as period the order of $\sigma \upharpoonright E$. However, there are non-generators that also have maximal period.

Lemma 3.2 *Let $E = A(\tau^{2^k})$ be a generalized eigenspace and let $p = 2^k \text{ per } \tau$ be the order of σ on E , where $k > 0$. Then for all $u \in E$, $u \in A(\tau^{2^{k-1}})$ if, and only if, $\text{per } u < p$.*

Proof. The period of any element in $A(\tau^{2^i})$ is at most $2^i \text{ per } \tau$, so the implication from left to right is obvious. Now suppose $q = \text{per } u < p$. Since q must divide p , we have $q = 2^j r$ where $j \leq k$ and r is odd. Now let $f = \gcd(x^q + 1, \tau^{2^k})$. Then $f \cdot u = 0$. To see this, note $f = a(x^q + 1) + b\tau^{2^k}$ for some cofactors $a, b \in \mathbb{F}_2[x]$. But for $r < p$ the GCD is 1, and for $r = p$ it is τ^{2^i} for some $i < k$. Hence $u \in A(\tau^{2^i})$ for some $i < k$, and we are done. \square

Note that the generators of smaller eigenspaces may well have the same period as the generators of the larger space, it is only when the exponent of the corresponding irreducible term reaches a smaller power of 2 that the periods decrease. The number of generators in $A(\tau^e)$ is easy to determine since, as a $\mathbb{F}_2[x]$ module, the space is isomorphic to $\mathbb{F}_2[x]/(\tau^e)$. Thus, there are $(2^d - 1)2^{(e-1)d}$ generators in $A(\tau^e)$, corresponding to the units in $\mathbb{F}_2[x]/(\tau^e)$.

Combining the last lemma with the results from the previous section, we can now give a uniform description of the transient length and the period of all four σ operations.

Theorem 3.3 *Let $p = x^a(x+1)^b \prod \tau_i^c$ be the minimal polynomial of σ . Then the nilpotent part of σ has nilpotency index a , and the regular part has order the least common multiple of the periods of τ_i , multiplied by c .*

Example As an example consider σ_z^- on a grid of size $n = 50$. Since n is even, σ_z^- is reversible. π_{51} factors as the square of the following irreducible terms:

$$1 + x, 1 + x + x^4, 1 + x + x^2 + x^3 + x^4, 1 + x^2 + x^3 + x^4 + x^8, 1 + x^2 + x^3 + x^7 + x^8$$

Thus, there are 5 subspaces in the decomposition, with dimensions 2, 8, 8, 16, 16, respectively. The restrictions of σ_z^- to these spaces have orders 2, 30, 10, 510, 510. Hence, the order of σ_z^- is 510.

4 Symmetries and Eigenspaces

Assume the eigenspace decomposition $V = K \oplus \bigoplus E_i$ from the last section. Given a pattern u on a cycle in the diagram, the cycle is contained in the σ -cyclic subspace $U = \mathbb{F}_2[x] \cdot u$. The decomposition has the property that any σ -invariant subspace U can be recovered from the components in the various eigenspaces: $U = \bigoplus (U \cap E_i)$. Hence it suffices to analyze the σ -invariant subspaces of the eigenspaces $A(\tau^e)$ where e is the appropriate power of 2. In the fixed boundary case, the eigenspaces are irreducible, hence the only σ -invariant subspaces are of the form $A(\tau^j)$ where $0 \leq j \leq e$. In the cyclic case we have to consider two such chains of subspaces. Now let $P_i : V \rightarrow E_i$ be the canonical projection. As far as the length of the cycle generated by $u_i = P_i(u)$ is concerned, it follows from

lemma 3.2 that only exponents of τ that are powers of 2 are relevant. Hence, for fixed boundary conditions, we have to study the chain

$$E_{i,k} \subseteq E_{i,k-1} \subseteq \dots \subseteq E_{i,1} \subseteq E_{i,0} = E_i \quad (6)$$

where $E_{j,i} = A_z(\tau_i^{2^{k-j+1}}, n)$. By lemma 3.2, the orbit of u_i is a cycle of length 2^{k-j+1} per τ_i where j is maximal such that $u_i \in E_{i,j}$. The length of the orbit of $u = \sum u_i$ is thus the least common multiple of these local periods. Likewise, one can easily count the number of cycles of any given length. We will return to this topic in section 5.

The purpose of this section is to show that the subspaces $E_{i,j}$ are closely connected to the geometry of the pattern spaces. Moreover, we will give efficient methods to generate bases for all these spaces.

4.1 Simulations

To describe the symmetries of pattern spaces, and to obtain generators for eigenspaces, it is convenient to consider structural maps between pattern spaces. To this end, consider two automata $\langle \mathbf{2}^m, \sigma \rangle$ and $\langle \mathbf{2}^n, \sigma' \rangle$, not necessarily of the same type. A monomorphism $f : \mathbf{2}^m \rightarrow \mathbf{2}^n$ that commutes with the σ operators, $f \circ \sigma = \sigma' \circ f$, will be called a *simulation*. There are several natural simulations between σ -automata based on symmetries and repetition of patterns:

$$\begin{aligned} \langle \mathbf{2}^m, \sigma_z^- \rangle &\longrightarrow \langle \mathbf{2}^{r(m+1)-1}, \sigma_z^- \rangle \\ \langle \mathbf{2}^m, \sigma_z^- \rangle &\longrightarrow \langle \mathbf{2}^{r(m+1)}, \sigma_c^- \rangle \\ \langle \mathbf{2}^m, \sigma_c^- \rangle &\longrightarrow \langle \mathbf{2}^{r m}, \sigma_c^- \rangle \end{aligned}$$

By abuse of notation, we will refer to all these maps as rep for repetition. It will always be clear from the domain and codomain which simulation we are referring to. For example, for $r = 2$, the first embedding is $\text{rep}(u) = (u, 0, S(u))$, the second $\text{rep}(u) = (u, 0, S(u), 0)$, and the last is duplication $\text{rep}(u) = (u, u)$. The images of these embeddings have rotational or reflectional symmetries in the target spaces. Simulations are obviously closed under composition, and we have, for example, the commutative diagram

$$\begin{array}{ccc} \mathcal{C}_z^-(n) & \xrightarrow{\text{rep}} & \mathcal{C}_z^-(2n+1) \\ \downarrow \text{rep} & & \downarrow \text{rep} \\ \mathcal{C}_c^-(2n+2) & \xrightarrow{\text{rep}} & \mathcal{C}_c^-(4n+4) \end{array}$$

The translation from fixed to cyclic boundary conditions will be used later. Note that in any simulation, the minimal polynomial of the simulated automaton divides the minimal polynomial of the simulating one. As a matter of fact, for the first two simulations above, the irreducible factors in both polynomials are the same, only the exponents change. Since the repetition maps preserve eigenspaces they can be used to transfer these spaces to higher dimensional pattern spaces:

$$\begin{aligned} A_z(\tau^e, m) &\xrightarrow{\text{rep}} A_z(\tau^e, r(m+1)-1) \\ A_c(\tau^e, m) &\xrightarrow{\text{rep}} A_c(\tau^e, r m) \end{aligned}$$

where $r \geq 1$ is arbitrary. Since the dimension of the eigenspaces depends only on the degree of τ and e , these maps are isomorphisms. In particular, the maps preserve the order of a pattern. To obtain patterns with higher order we need one other type of simulation, *squaring*. As the name indicates, this simulation is motivated by the fact that patterns can be construed as polynomials, a technique used frequently in [10]. To be more explicit, there is a natural \mathbb{F}_2 -vector space isomorphism $\text{tp} : \mathbf{2}^m \rightarrow \mathbb{F}_2^{(m)}[x]$, with inverse fp , from m -dimensional pattern space to the set of polynomials of degree less than m . The action of σ_c^- on this space can be expressed by multiplication with $x^{-1} + x = x^{m-1} + x$ in the quotient ring $\mathbb{F}_2[x]/(x^m + 1)$.

Now define the *squaring* operation $\text{sq} : \mathbf{2}^n \rightarrow \mathbf{2}^{2n}$ by $\text{sq}(u) = \text{fp}(\text{tp}(u)^2)$. In terms of the standard basis, this means $\text{sq}(e_i) = e_{2(i-1)+1}$. By abuse of notation, we write $\text{sq} : \mathbf{2}^n \rightarrow \mathbf{2}^{2n+1}$, $\text{sq}(e_i) = e_{2i}$, also for the fixed boundary case. In either case, sq is a monomorphism, and indeed a simulation. By iteration we obtain embeddings

$$\begin{aligned} \langle \mathbf{2}^m, \sigma_z^- \rangle &\xrightarrow{\text{sq}} \langle \mathbf{2}^{2^k(m+1)-1}, (\sigma_z^-)^{2^k} \rangle \\ \langle \mathbf{2}^m, \sigma_c^- \rangle &\xrightarrow{\text{sq}} \langle \mathbf{2}^{2^k m}, (\sigma_c^-)^{2^k} \rangle \end{aligned}$$

Hence we have to following proposition.

Proposition 4.1 *Let g be a generator for the eigenspace $A_z(\tau^e, m)$. Then $\text{sq}(u)$ generates $A_z(\tau^{2^e}, 2m+1)$. Likewise, for any generator g of $A_c(\tau^e, m)$, $\text{sq}(u)$ generates $A_c(\tau^{2^e}, 2m)$.*

The next two sections explain how to obtain the initial generator.

4.2 Symmetries

Simulations from a σ -automaton to itself will be referred to as *auto-simulations*. The auto-simulations naturally form a group that acts on the pattern space. Those auto-simulations for which the subspaces in our decomposition are invariant can already be expressed as a polynomial in σ .

Lemma 4.1 *Let F be an auto-simulation of a σ -automaton such that all the eigenspaces E in the decomposition are F -invariant. Then F is polynomially representable.*

Proof. First consider a σ -cyclic decomposition space $E \subseteq A(\tau^e)$ and let U be the set of generators of E . Fix a generator $u_0 \in U$ and, for any $u \in U$, let $r_u \in \mathbb{F}_2[x]$ such that $u = r_u \cdot u_0$. Clearly, the map $v \mapsto r_u \cdot v$ is an auto-simulation, and any auto-simulation on E can be represented in this form. It follows that the number of auto-simulations is $|U|$.

For an auto-simulation $F : V \rightarrow V$ on the whole space, note that the projections $P_i : V \rightarrow E_i$ are polynomially representable, $P_i(u) = r_i \cdot u$. We have just seen that the restrictions $F_i = F \upharpoonright E_i$ are polynomially representable. But then $F = \sum F_i \circ P_i$ is also polynomially representable. \square

Thus, for fixed boundary conditions all auto-simulations are polynomially representable. As we will see shortly, the reflection $S \upharpoonright E$ can even be represented by a polynomial of the form x^t , i.e., by iteration of σ . However, for cyclic boundary conditions the reflection S and rotation R fail to be polynomially representable.

A *symmetry* is an auto-simulation that respects the geometry of the automaton. For fixed boundary conditions there is only one symmetry, namely reflection. For cyclic boundary conditions we also have rotations. Given a group G of symmetries and a pattern u , we write $G_x =$

$\{f \in G \mid f(u) = u\}$ for the *stabilizer* of u . G_x is a subgroup of G , and we can measure the degree of symmetry of a pattern by the size of this subgroup. We will see that the eigenspaces in the elementary divisor decomposition have important subspaces determined by higher degrees of symmetry.

The next two lemmata describe the relationship between cyclic subspaces of the eigenspaces $A(\tau^e)$ and $A(\tau^{2e})$ in terms of involutions on the space. The first one makes use of the self-adjointness of σ .

Lemma 4.2 *Let E be a subspace of V , $F : E \rightarrow E$ a linear self-adjoint map of nilpotency index 2. Let $E_1 \subseteq E$ be the kernel of f . Then $2 \dim E_1 = \dim E$ and E_1 is self-orthogonal: $E_1^\perp = E_1$.*

Proof. Let n be the dimension of the pattern space V . Since $f^2(E) = 0$ and since f is self-adjoint we must have $\text{im } f \subseteq \ker f = (\text{im } f)^\perp$. But then $\dim E = \dim \text{im } f + \dim (\text{im } f)^\perp = 2 \dim \text{im } f$, and our first claim follows. Also note that n must be even. But then $\text{im } f = \ker f$, and the second claim follows. \square

Lemma 4.3 *Let $E \subseteq A(p^2)$ be a σ -cyclic subspace where $p = \tau^e$, $\tau \neq x$, is a power of an irreducible factor of the minimal polynomial. Let $E_1 = E \cap A(p)$ and denote the order of $\sigma \upharpoonright E_1$ by t . Consider an involution $F : E \rightarrow E$ that commutes with σ .*

Then F is polynomially representable, $F(u) = f \cdot u$. Moreover, $f = x^t \pmod{\tau}$. If F is not the identity, the E_1 is the set of fixed points of F .

Proof. Let u_0 be a generator of E . Then there is a $f \in \mathbb{F}_2[x]$ such that $F(u_0) = f \cdot u$. Then $F(u) = F(h \cdot u_0) = h \cdot F(u_0) = hf \cdot u_0 = f \cdot u$.

By the definition of p , $x^t + 1 = 0 \pmod{p}$. Since F is an involution, we have $f^2 + 1 = 0 \pmod{p^2}$. But p is a power of an irreducible, so $f + 1 = 0 \pmod{p}$ and our claim follows.

Note that we must have $f + 1 = gp$ where $g \in \mathbb{F}_2[x]$. Since F is not the identity, g and p a coprime. Now consider $u \in E_1$. Then $p \cdot u = 0$ and therefore $0 = gp \cdot u = (f + 1) \cdot u = F(u) + u$, and u is a fixed point.

On the other hand, let u be a fixed point. Then again $0 = gp \cdot u$. Consider cofactors $a, b \in \mathbb{F}_2[x]$ such that $ag + bp = 1$. Then $0 = agp \cdot u = (1 + bp)p \cdot u = p \cdot u + bp^2 \cdot u = pu$. Thus, $u \in E_1$, and we are done. \square

As an application of these lemmata consider the eigenspace $E = A_z(\tau^{2^{k+1}}, n)$. The order of σ_z^- on E is $t = 2^{k+1}$ per τ , and for all $u \in E$ we have $\sigma_z^{t/2}(u) = S(u)$. The subspace $E_1 = A_z(\tau^{2^k}, n)$ is clearly self-orthogonal and consists of the fixed points of E under S . We will see shortly that this is no coincidence.

4.3 Fixed Boundary Conditions

While S is the only symmetry of the whole space, for $n + 1 = 2^k \cdot m$ there are other symmetries on subspaces that are relevant to our analysis. To be more precise, for any function $f : \mathbf{2}^m \rightarrow \mathbf{2}^m$ define an extension $\text{rep}(f) : \mathbf{2}^{2m+1} \rightarrow \mathbf{2}^{2m+1}$ by $\text{rep}(f)(u) = (f(u_{1:m}), u_{m+1}, f(u_{m+2:2m+1}))$. Here $u_{i:j}$ denotes the projection of u to the space spanned by e_i, e_{i+1}, \dots, e_j . Given $n + 1 = 2^k \cdot (m + 1)$ where $m + 1$ is odd, let $S_i = \text{rep}(S(2^{k-i}(m + 1) - 1))$ where $S(r) : \mathbf{2}^r \rightarrow \mathbf{2}^r$ is plain reflection. This leads to the following description of the eigenspaces.

Theorem 4.1 *Let $n + 1 = 2^k \cdot (m + 1)$ where $m + 1$ is odd, and consider an irreducible factor τ of the minimal polynomial of $\sigma_z(n)$. There is a strictly increasing chain of σ -cyclic subspaces*

$$E_{k+1} \subseteq E_k \subseteq \dots \subseteq E_1 \subseteq E_0 = A_z(\tau^{2^{k+1}}, n)$$

where $E_i = A_z(\tau^{2^{k-i+1}}, n)$, $i = 0, \dots, k + 1$. The dimension of E_i is $2^{k-i+1} \deg \tau$. Moreover, E_{i+1} is the set of fixed points of S_i in E_i .

Proof. The proof is a straightforward induction on i , using lemma 4.3. □

Generators for the σ -cyclic spaces in the chain can be found by determining the null spaces of the matrices $\tau^e(\sigma)$. However, we can use simulations to bypass linear algebra. Suppose $n + 1 = 2^k \cdot (m + 1)$ where $m + 1$ is odd. The minimal polynomials for m and n have the form

$$\tau_1^2 \tau_2^2 \dots \tau_r^2 \quad \text{and} \quad x^{2^k-1} \tau_1^{2^{k+1}} \tau_2^{2^{k+1}} \dots \tau_r^{2^{k+1}}.$$

Hence, σ_z^- is reversible on $\mathbf{2}^m$, but has a nilpotent part of dimension $2^k - 1$ on $\mathbf{2}^n$. A natural basis is obtained by applying the repetition map to the identity matrix of size $2^k - 1$.

For the regular part of σ_z we proceed as follows. By embedding a smaller pattern space via a repetition map, we can obtain patterns with nested symmetry in the eigenspaces: $E_{i,j} = E_i \cap \text{rep}(\mathbf{2}^{2^{k-j}(m+1)-1})$. Computationally it is more efficient to construct a generator for $E_{i,j}$ directly. To this end, let $m' = 2^e(m + 1) - 1$, and consider the following diagram.

$$\begin{array}{ccc} A_z(\tau^2, m) & \xrightarrow{\text{sq}} & A_z(\tau^{2^{k+1}}, n) \\ \text{sq} \downarrow & & \uparrow \\ A_z(\tau^{2^{i+1}}, m') & \xrightarrow{\text{rep}} & A_z(\tau^{2^{i+1}}, n) \end{array} \quad (7)$$

Hence, we only need a generator for $A_z(\tau^2, m)$, all others can be obtained by squaring and repetition. To obtain the first generator, we can exploit the simulation from $\mathcal{C}_z(m)$ to $\mathcal{C}_c(2m + 2)$. We will see in section 4.4 below how to determine a generator for cyclic boundary conditions, using cyclotomic polynomials. Note that the minimal polynomial of $\sigma_c^-(m')$ is the minimal polynomial of $\sigma_z^-(m)$ multiplied by x , so we have precisely the same irreducible factors other than x . Now consider the following diagram.

$$\begin{array}{ccc} \mathcal{C}_z^-(m) & \xrightarrow{\text{rep}} & \mathcal{C}_c^-(2m + 2) \\ \uparrow & & \uparrow \\ A_z(\tau^2, m) & \xrightarrow{\text{rep}} & A_c(\tau^2, n) \end{array} \quad (8)$$

The eigenspace on the left has dimension $2 \deg \tau$, the space on the right $4 \deg \tau$. Hence, the image of the left eigenspace is one of the two isomorphic spaces on the right, and we can retrieve a generator for the first by applying the inverse of the simulation map, after an appropriate rotation of the pattern.

In some simple cases one can describe generators more directly. For example, consider $n + 1 = 2^k \cdot m$ and let p be an odd prime dividing m . Set $l = o_p(m)$. Since $\pi_{p^l} = \rho(p) \rho(p^2) \dots \rho(p^l)$ there is a chain of σ -cyclic subspaces in the decomposition corresponding to the irreducible polynomials in these critical factors. The degree of $\rho(p^i)$ is $(p - 1)p^{i-1}$, and we have embeddings

$$\begin{aligned} \mathbf{2}^{p-1} &= A_z(\rho(p), p-1) \xrightarrow{\text{rep}} A_c(\rho(p^e), p^e-1) \\ \mathbf{2}^{p-1} &= A_z(\rho(p), p-1) \xrightarrow{\text{sq}} A_c(\rho^{2^e}(p), 2^e p-1) \end{aligned}$$

It is easy to see that $A_z(\rho(p^i), p^i-1) \subseteq \mathbf{2}^{p^i-1}$ is generated as a σ -cyclic module by the standard basis vector e_{p^i} . Hence we have

$$A_z(\pi_{p^k}, 2^k-1) = \bigoplus \text{rep}(A_z(\rho(p^i), 2^{p^i}-1)).$$

The last space can be embedded into $\mathbf{2}^n$ via the squaring map.

4.4 Cyclic Boundary Conditions

For cyclic boundary conditions the group of symmetries is the dihedral group

$D = \mathcal{D}_n$, generated by reflection S and rotation R . Thus, the standard presentation for \mathcal{D} is $R^n = S^2 = 1$, $RS = SR^{-1}$. We can measure the degree of symmetry of a pattern $u \in \mathbf{2}^n$ by the stabilizer \mathcal{D}_u . It will be convenient to identify \mathcal{D} with the semidirect product $\mathbb{Z}_n \times_{\theta} \mathbb{Z}_2$ where $\theta(1)(x) = -x$. Since the stabilizer is a subgroup of \mathcal{D} it has the form $H \times_{\theta} 1$ or $H \times_{\theta} \mathbb{Z}_2$ where H is a cyclic subgroup of \mathbb{Z}_n . As we will see, there always are generators whose stabilizer is of the second form. Our goal is to establish the following theorem.

Theorem 4.2 *Let $n = 2^k \cdot m$ where m is odd, and consider an irreducible factor τ of the minimal polynomial of $\sigma_c(n)$. There is a strictly increasing chain of pairs of σ -cyclic subspaces*

$$E_k \oplus E'_k \subseteq E_{k-1} \oplus E'_{k-1} \subseteq \dots \subseteq E_0 \oplus E'_0$$

where $E_i \oplus E'_i = A_c(\tau^{2^i}, n)$, $i = 0, \dots, k$. Then there are generators in $u \in E_i$, and $v \in E_{i+1}$ such that $\mathcal{D}_v/\mathcal{D}_u \cong \mathbb{Z}_2$.

Suppose $n = 2^k \cdot m$ where m is odd and $k \geq 0$. The minimal polynomials for σ_c^- for m and n here have the forms

$$x \tau_1 \tau_2 \dots \tau_r \quad \text{and} \quad x^{2^{k-1}} \tau_1^{2^k} \tau_2^{2^k} \dots \tau_r^{2^k},$$

respectively. The nilpotent part for $\mathbf{2}^m$ is simply the kernel of σ_c^- , which has dimension 1 and its only non-trivial member is $\mathbf{1}$. For $\mathbf{2}^n$, the nilpotent part K consists of two σ -cyclic subspaces, generated by $\text{sq}(\mathbf{1})$ and $R(\text{sq}(\mathbf{1}))$. Needless to say, $\text{sq}(\mathbf{1}) = \sum_{0 \leq i < m} e_{i \cdot 2^k + 1}$. Since by lemma 3.1 a basis can also be obtained by rotation, the natural basis for K is just the Kronecker product of the identity matrix of size 2^k , and the all-ones vector.

For the regular part, we can construct a chain of pairs of σ -cyclic subspaces for each of the irreducible terms τ as follows.

$$\begin{array}{ccc} A_c(\tau, m) & \xrightarrow{\text{sq}} & A_c(\tau^{2^k}, n) \\ \text{sq} \downarrow & & \uparrow \\ A_c(\tau^{2^i}, m) & \xrightarrow{\text{rep}} & A_c(\tau^{2^i}, n) \end{array} \quad (9)$$

It remains to find a way to calculate generators for the first generalized eigenspace $A_c(\tau, m)$ where $\rho_m = \tau^2$. Using the translation to and from polynomials of degree less than m , one can see that a generator for $A_c(\tau, m)$ can be obtained as

$$u = \text{fp}((x^m + 1) / \gcd(x^m + 1, \tau(x^{m+1} + x))).$$

We write C_d for the d th cyclotomic polynomial over \mathbb{F}_2 . Thus, $C_d = (x - a_1)(x - a_2) \dots (x - a_k)$ where the a_i are all the primitive d th roots of unity in some suitable splitting field. In terms of these polynomials, the generator has the form

$$u_0 = \text{fp}\left(\prod_{e|n, e \neq n} C_e\right).$$

For the sake of this argument, let us call a univariate polynomial of degree d symmetric if its coefficient list is symmetric: $c_i = c_{d-i}$ for all $i = 0, \dots, d$. It follows from the definition that cyclotomic polynomials are symmetric, and it is easy to see that products of symmetric polynomials are again symmetric. In terms of the standard presentation of the dihedral group, the stabilizer of $u = \text{fp}(C_m) \in \mathbf{2}^{\mathcal{F}(m)}$ has the form $\{1, S\}$ if m is composite, and is equal to \mathcal{D} otherwise. It follows that the stabilizer of generator u_0 above is isomorphic to $H \times_{\theta} \mathbb{Z}_2$.

Proof. (of theorem 4.2)

The proof is again by induction on i , using 4.3. As we have seen in theorem 3.2 and corollary 3.1, the eigenspaces $A_c(\tau^{2^i}, n)$ are not σ -cyclic, but consist of two isomorphic copies, where the isomorphism is the rotation R .

Now consider an irreducible factor τ in the minimal polynomial, say, $\tau | \rho_d$ where $d | m$. We have just shown that there is a generator u of $A_c(\tau, d)$ whose stabilizer is of the form $1 \times_{\theta} \mathbb{Z}_2$. The image $\text{rep}(u)$ in $A_c(\tau, m)$ has stabilizer isomorphic to $\mathbb{Z}_{m/d} \times_{\theta} \mathbb{Z}_2$, and the image in $A_c(\tau, n)$ has stabilizer isomorphic to $\mathbb{Z}_{2^k m/d} \times_{\theta} \mathbb{Z}_2$. Our claim follows. \square

On occasion, one can exploit properties of the cyclotomic polynomials to streamline the polynomial arithmetic in the computation of a generator. For any integer $m \geq 2$ with prime decomposition $m = \prod p_i^{e_i}$ let $m' = p_1 p_2 \dots p_r$. Then

$$C_m(x) = \prod_{d|m} (x^d + 1)^{\mu(m/d)} \tag{10}$$

$$C_m(x) = C_{m'}(x^{m/m'}) \tag{11}$$

where μ is again the Möbius function. It follows that for any prime p , $C_p(x) = \sum_{i < p} x^i = \frac{x^p + 1}{x + 1}$. Hence, if $m = p^k$, we have a generator $\text{fp}(x^{p^{k-1}} + 1)$. When m is the product of two distinct primes p and q there is a generator

$$(x^p + 1)(x^q + 1)/(x + 1) = (x^p + 1) \sum_{i_q} x^i.$$

The second form shows how to determine the corresponding pattern without polynomial arithmetic. Using equation (2) similar expressions can be derived for other cases.

5 The Cycle Structure

Given the decomposition of pattern space from the last section, we can now determine the regular part of the diagram of σ completely. More precisely, we can determine all possible cycle lengths, as well as the total number of cycles of each length. From the previous results, it is easy to compute bases for the various cycle spaces.

Given the elementary divisor decomposition $V = K \oplus E_0 \oplus E_1 \dots \oplus E_r$ and the order e_i of $\sigma \upharpoonright E_i$ for each component, it is now easy to determine the cycle structure. A cycle that lies in one of the subspaces E_i will be called *pure*, and *compound* otherwise. In the next section we will see how to determine the lengths and numbers of all pure cycles. All the compound cycles have as length the least common multiple of the lengths of some of the pure cycles, and the count is determined by the product of the counts of the corresponding pure cycles, and a nested gcd/lcm of the cycle lengths. For example, if we have three pure cycles in separate subspaces of lengths c_1, c_2, c_3 and counts b_1, b_2, b_3 , respectively, they will contribute $\gcd(c_1, \text{lcm}(c_2, c_3)) \cdot b_1 b_2 b_3$ many cycles of length $\text{lcm}(c_1, c_2, c_3)$.

Now consider an eigenspace $E = A(\tau^e)$. Given the period p of τ and d the degree of τ it is straightforward to calculate the number and lengths of all pure cycles in E . More precisely, E contains 1 cycle of length 1, p cycles of length $(2^d - 1)/p$, and

$$\frac{2^{d \cdot 2^i} - 2^{d \cdot 2^{i-1}}}{p \cdot 2^i}$$

cycles of length $p \cdot 2^i$, up to length $p \cdot e$. The number of pure cycles in a subspace thus increases at a doubly exponential rate with the length of the cycle.

5.1 Fixed Boundary Conditions

From section 4 we can now easily determine the complete cycle structure of an σ -automaton. For fixed boundary conditions, the subspaces with higher degrees of symmetry produce to shorter cycles in the diagram.

Example Consider again the σ_z^- -automaton on $n = 50$ cells. As we have seen, the minimal polynomial has a factorization of the form $\tau_1^2 \tau_2^2 \dots \tau_5^2$. Correspondingly, there are 5 elementary divisor subspaces E_1, \dots, E_5 , each with a subspace of symmetric patterns. The last two irreducible factors have the same period, the cycle structure of E_4 and E_5 is therefore identical. The lengths and counts for the pure cycles are shown in the next table.

E_1	E_2	E_3	E_4, E_5
1 1	1 1	1 1	1 1
1 1	15 1	5 3	255 1
2 1	30 8	10 24	510 128

For all cycles we obtain the following count. Since σ_z^- is reversible for $n = 50$, the cycles account for the whole pattern space.

length	count
1	2
2	1
5	6
10	99
15	32
30	8688
255	131584
510	2207646809856

Thus, the probability for a randomly selected pattern to not lie on a cycle of length 510 is about $3 \cdot 10^{-8}$.

Recall that $\tau^+(\sigma^+) = \tau(\sigma)$, so that the situation for σ_z^+ is strictly analogous. However, there is a minor complication due to irreducible factors of the form $(x+1)^b$. The period of $x+1$ is trivially 1, so we are dealing with cycles of length 2^i . However, the exponent b here is of the form $2^k - 1$, so the largest of the subspaces does not have dimension a power of 2, and the calculation of the frequencies of pure cycles has to be adjusted correspondingly.

Example Consider the σ_z^+ -automaton on $n = 39$ cells. The minimal polynomial is $(1+x)^7(1+x+x^2)^{16}$, and leads to the following frequencies for pure cycles.

E_1		E_2	
1	1	1	1
1	1	3	1
2	1	6	2
4	3	12	20
8	14	24	2720
-	-	48	89477120

Note that a full-dimensional subspace would have $(2^8 - 2^4)/8 = 30$ cycles of length 8. The frequencies for all cycles are as follows.

length	count
1	2
2	1
3	2
4	3
6	9
8	14
12	335
24	349350
48	11453071360

5.2 Cyclic Boundary Conditions

For cyclic boundary conditions, we have to slightly adjust the arguments from the last section. Since each eigenspace $E = A(\tau^e)$ consists of a direct sum of two σ -cyclic spaces $E_1 \oplus E_2$, which are

isomorphic as $\mathbb{F}_2[x]$ -modules, all the pure cycles now appear in two subspaces, and the maximum cycle length in each is $e \cdot \text{per}(\tau)$.

Recall that $\text{mp}_c^-(n) = x\pi_{n/2}$ or $\text{mp}_c^-(n) = x\sqrt{\pi_n}$ depending on whether n is even or odd, respectively. Consider an irreducible $\tau \neq x$ dividing the minimal polynomial, say $\text{mp}_c^-(n) = \tau^{2^k} \cdot q$ where τ does not divide q , and let $E = A(\tau^{2^k}) \subseteq \mathbf{2}^n$. Since $\tau^{2^{k+1}}$ is a factor of π_n we have the nullspace of $U_k = \ker \tau^{2^k}(\sigma_z^-) \subseteq \mathbf{2}^{n-1}$, and this space consists precisely of the symmetric patterns.

The natural embedding $\eta : \mathbf{2}^{n-1} \rightarrow \mathbf{2}^n$, $u \mapsto u0$ (or any symmetry thereof) is a partial simulation: η fails to commute with the σ operators on the whole space, but for $u \in U_k$ we have $\sigma_c^-(\eta(u)) = \eta(\sigma_z^-(u))$. Hence, $\eta(U_k)$ is a $d \cdot e$ dimensional subspace of E , and indeed we can decompose E as $\eta(U_k) \oplus R(\eta(U_k))$. Likewise, for each subspace $U_i = \ker \tau^{2^i}(\sigma_z^-)$, $i = 0, \dots, k$, in $\mathbf{2}^{n-1}$ we obtain two subspaces in $\mathbf{2}^n$. Since only powers of 2 can arise as the order of restrictions of σ_c^- to σ -cyclic subspaces, this characterizes all such subspaces.

6 Conclusion

We have show that the structure of the transition diagram of σ is entirely determined by the periods of the irreducible factors of the minimal polynomial of σ on the pattern space in question. The minimal polynomial, in turn, can be described in terms of π -polynomials, a binary version of the Fibonacci polynomials. The factorization theorem for these polynomials affords a uniform description in terms of certain irreducible factors and determines the structure of the elementary divisor decomposition.

For most values of n , the periods of the irreducible factors of π_n are divisors of the period of the irreducible factor of highest degree. Hence, in most transition diagrams, the maximal cycle length is attained by a pure cycle. However, this is not always the case: for $n \leq 200$ the exceptions are 38, 54, 77, 109, 110, 155, 174, 182. We do not currently understand the nature of the exceptions to this rule. Likewise, we do not know if the cycle structure can be determined using only polynomial arithmetic but without recourse to factorization.

Another open problem concerns the relationship between the diagram of σ_z^- and σ_z^+ . Let us disregard irreducible factors of degree 1, which contribute only to the nilpotent part, and to pure cycles of length a power of 2, respectively. Then the irreducible factors of the corresponding minimal polynomial for a grid of size n are simply obtained by applying the involution $x \mapsto x + 1$, which preserves the degree of the irreducible factors, and thus the dimension of the associated subspaces. However, the involution changes the period of an irreducible polynomial in a rather complicated fashion. For example, there are irreducibles of degree 10 and period 1023 whose images under the involution have period 33, 93, 341 and 1023, respectively. Likewise there are degree 10 irreducibles with period 341 whose images have period 11, 93, 341 and 1023, respectively. We are not aware of a simple characterization of the period of $\tau(x + 1)$ given the period of τ . This problem is similar to the question of computing the π -depth of an irreducible polynomial τ , i.e., the least number n such that τ divides π_n . It can be shown that for any irreducible of depth d the degree must be the suborder of 2 in the multiplicative subgroup \mathbb{Z}_d^* , but again the involution seems to change the depth in a rather complicated fashion. For the relevance of this problem to the study of two-dimension σ -automata see [14].

Given a specific pattern X one can use the elementary divisor decomposition and the associated subspaces as in the last section to determine the transient length as well as the period of X by

computing the representation of X with respect to the bases of the subspaces. While one can determine the bases using only polynomial arithmetic and factorization, and the order of σ on these spaces by computing the period of an irreducible polynomial, the last step seems to require linear algebra. Specifically, one has to solve a system of linear equations over \mathbb{F}_2 . We do not know whether one can avoid this last step.

i	ρ_i	periods
2	x	1
3	$1 + x$	1
5	$1 + x + x^2$	3
7	$1 + x^2 + x^3$	7
9	$1 + x + x^3$	7
11	$1 + x + x^2 + x^4 + x^5$	31
13	$1 + x + x^4 + x^5 + x^6$	63
15	$1 + x^3 + x^4$	15
17	$1 + x + x^4, 1 + x + x^2 + x^3 + x^4$	15, 5
19	$1 + x + x^4 + x^5 + x^6 + x^8 + x^9$	511
21	$1 + x^5 + x^6$	63
23	$1 + x^2 + x^3 + x^4 + x^8 + x^{10} + x^{11}$	2047
25	$1 + x + x^2 + x^3 + x^5 + x^6 + x^{10}$	1023
27	$1 + x + x^5 + x^7 + x^9$	511
29	$1 + x + x^8 + x^9 + x^{12} + x^{13} + x^{14}$	16383
31	$1 + x^2 + x^5, 1 + x^3 + x^5, 1 + x^2 + x^3 + x^4 + x^5$	31, 31, 31
33	$1 + x + x^2 + x^3 + x^5, 1 + x + x^3 + x^4 + x^5$	31, 31
35	$1 + x^5 + x^6 + x^7 + x^9 + x^{11} + x^{12}$	4095
37	$1 + x + x^2 + x^8 + x^9 + x^{10} + x^{12} + x^{13} + x^{16} + x^{17} + x^{18}$	87381
39	$1 + x^3 + x^4 + x^5 + x^6 + x^7 + x^8 + x^{11} + x^{12}$	1365
41	$1 + x + x^3 + x^4 + x^7 + x^8 + x^{10}, 1 + x + x^4 + x^9 + x^{10}$	341, 1023
43	$1 + x + x^7, 1 + x + x^2 + x^3 + x^4 + x^5 + x^7,$	127, 127,
43	$1 + x + x^2 + x^4 + x^5 + x^6 + x^7$	127
45	$1 + x^3 + x^4 + x^5 + x^7 + x^9 + x^{12}$	4095
47	$1 + x^4 + x^6 + x^7 + x^8 + x^{16} + x^{20} + x^{22} + x^{23}$	8388607
49	$1 + x^2 + x^3 + x^7 + x^9 + x^{10} + x^{11} + x^{14} + x^{17} + x^{19} + x^{21}$	2097151
51	$1 + x^2 + x^3 + x^4 + x^8, 1 + x^2 + x^3 + x^7 + x^8$	255, 255
53	$1 + x + x^2 + x^4 + x^5 + x^{16} + x^{17} + x^{18} + x^{20} + x^{21} + x^{24} + x^{25} + x^{26}$	67108863
55	$1 + x^3 + x^4 + x^5 + x^6 + x^8 + x^{11} + x^{12} + x^{14} + x^{15} + x^{17} + x^{19} + x^{20}$	349525
57	$1 + x + x^2 + x^3 + x^5 + x^6 + x^9, 1 + x + x^3 + x^4 + x^6 + x^8 + x^9$	511, 511
59	$1 + x + x^2 + x^{16} + x^{17} + x^{18} + x^{24} + x^{25} + x^{26} + x^{28} + x^{29}$	536870911
61	$1 + x + x^{16} + x^{17} + x^{24} + x^{25} + x^{28} + x^{29} + x^{30}$	1073741823
63	$1 + x^3 + x^6, 1 + x^2 + x^3 + x^5 + x^6, 1 + x^2 + x^4 + x^5 + x^6$	9, 63, 21
65	$1 + x + x^6, 1 + x + x^2 + x^4 + x^6,$	63, 21,
65	$1 + x + x^3 + x^4 + x^6, 1 + x + x^2 + x^5 + x^6$	63, 63
67	$1 + x + x^{16} + x^{17} + x^{24} + x^{25} + x^{28} + x^{29} + x^{30} + x^{32} + x^{33}$	8589934591
69	$1 + x^5 + x^6 + x^9 + x^{10} + x^{11} + x^{12} + x^{13} + x^{14} + x^{15} + x^{16} + x^{21} + x^{22}$	4194303

References

- [1] R. Barua. Additive cellular automata and matrices over finite fields. Technical Report 17/91, Indian Statistical Institute, Calcutta, India, 1991.
- [2] E. R. Berlekamp. *Algebraic Coding Theory*. McGraw-Hill, 1968.
- [3] W. Chin, B. Cortzen, and J. Goldman. Linear cellular automata with boundary conditions. *Linear Algebra and its Applications*, 322:193–206, 2001.
- [4] S. W. Golomb. *Shift Register Sequences*. Aegean Park Press, 1982.
- [5] R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, 1988.
- [6] W. Greub. *Linear Algebra*. Springer-Verlag, 1981.
- [7] T. W. Hungerford. *Algebra*. Springer-Verlag, 1974.
- [8] K. Culik II, L. P. Hurd, and S. Yu. Computation theoretic aspects of cellular automata. *Physica D*, 45:357–378, 1989.
- [9] R. Lidl and H. Niederreiter. *Finite Fields*, volume 20 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1984.
- [10] O. Martin, A. M. Odlyzko, and S. Wolfram. Algebraic properties of cellular automata. *Commun. Math. Phys.*, 93:219–258, 1984.
- [11] K. Sutner. On σ -automata. *Complex Systems*, 2(1):1–28, 1988.
- [12] K. Sutner. Classifying circular cellular automata. *Physica D*, 45(1–3):386–395, 1990.
- [13] K. Sutner. The complexity of finite cellular automata. *Journal of Computer and Systems Sciences*, 50(1):87–97, 1995.
- [14] K. Sutner. σ -automata and Chebyshev polynomials. *Theoretical Computer Science*, (230):49–73, 2000.