σ Automata

Klaus Sutner
Carnegie Mellon University

Spring 2011

Outline

A Beer-Fueled Challenge

- Suppose a chessboard is wired up with lights and buttons.
- Each square has one light and one button.
- Pushing the button toggles the light.
- Alas, it also toggles the lights of the 4 edge-adjacent squares.

Starting from all lights off, can one turn all the lights on?

A Toy

Obviously by Induction . . .

$1 \times m$ is easy

$2 \times m$ still easy

$3 \times m$ already annoying

No real hope to do induction $n \times m \rightsquigarrow (n+1) \times m$

If you can't prove something, generalize.

Global Operation σ

Let $G = \langle V, E \rangle$ be some undirected graph.
$A \in \mathbf{2}^{V \times V}$ its adjacency matrix.
$\mathcal{C} = V \to \mathbf{2}$ the space of all configurations over G.
$\mathcal{C}$ is a boolean group under pointwise addition mod 2.
In fact, think of $\mathcal{C}$ as a vector space over $\mathbb{F}_2$.
Define the following global operator on configuration space:

$$\sigma : \mathcal{C} \longrightarrow \mathcal{C}$$
$$\sigma(X) = (A + I) \cdot X$$

Notation

The graph G is usually clear from context.

For emphasis we sometimes write

$$\sigma_G \quad \text{or} \quad \sigma(G)$$

The definition makes sense for any(locally finite directed graph, but we will focus on undirected graphs.

Forget Grids

Theorem (1989)

The all-ones configuration $\mathbf{1}$ *lies in the range of* σ.

Proof.

$\langle X, Y \rangle = |X \cap Y| \bmod 2$

σ is self-adjoint

rng(σ) is the orthogonal complement of ker(σ)

for X in the kernel the induced subgraph has only odd-degree vertices

every element of the kernel has even cardinality (handshake lemma)

$\mathbf{1}$ is orthogonal to the kernel

□

Aside: A Universal Configuration

So $\mathbf{1}$ is universal in the sense that it lies in the range of σ for any graph.

Theorem (P. Winkler, Y. Dodis (2001))

$\mathbf{1}$ *is essentially the only universal pattern.*

Chessboard Solution

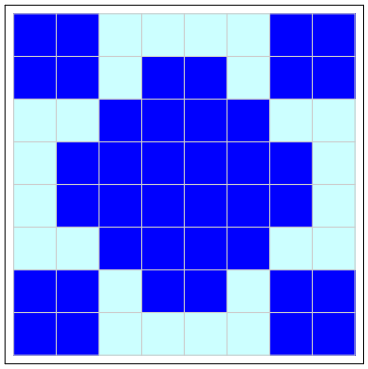

100 by 100

Discrete Dynamics

Think of $\langle \mathcal{C}, \sigma \rangle$ as a discrete dynamical system.

- reversibility
- predecessors
- (maximal) transients
- (maximal) cycle length
- cycle length distribution

We want computationally simple descriptions.

Cellular Automata

On the bi-infinite path P_∞ rule σ is known as elementary cellular automata number 150.

Studied in great detail Martin, Odlyzko and Wolfram 25 years ago.

ECA 150

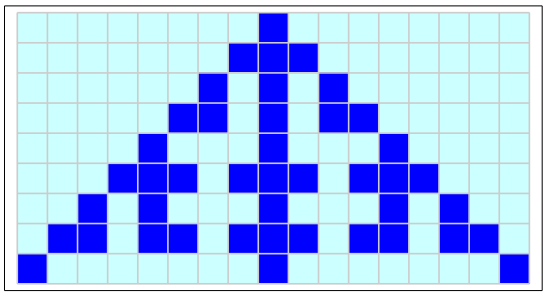

A Fractal

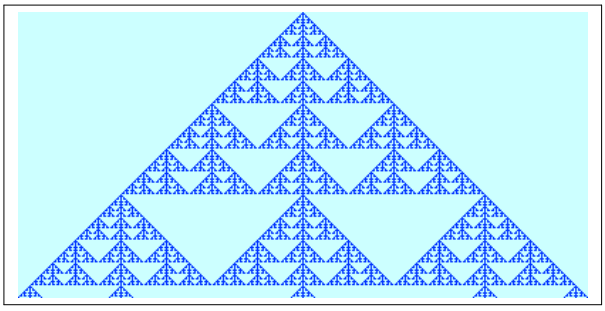

A Similar ECA

Elementary cellular automaton number 150 has a slightly simpler counterpart: ECA 90.

$$\sigma^- : \mathcal{C} \longrightarrow \mathcal{C}$$
$$\sigma^-(X) = A \cdot X$$

Uses open neighborhood instead of closed neighborhood.

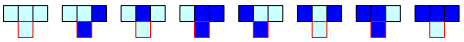

ECA 90

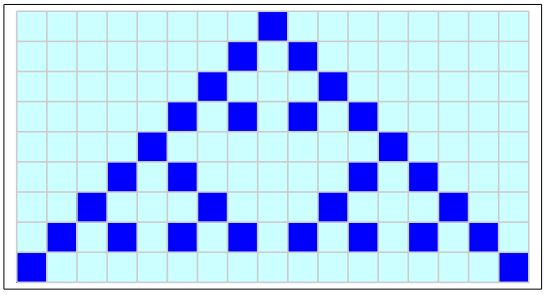

Sierpinski

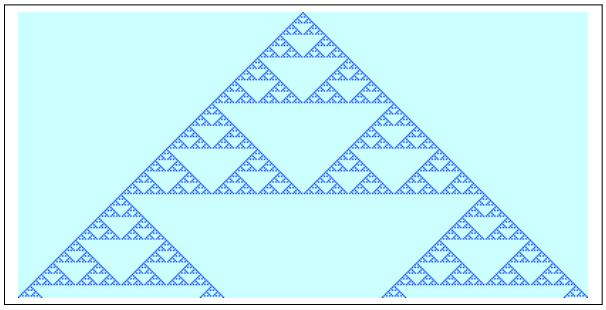

- σ Automata
2. Reversibility
- π Polynomials
- Irreversible Squares
- Orbits

Reversible Grids

So when is the $n \times m$ grid $P_{n,m}$ reversible under rule σ?

Note that the adjacency matrix for $P_{n,m}$ is a Kronecker matrix of size $nm \times nm$.

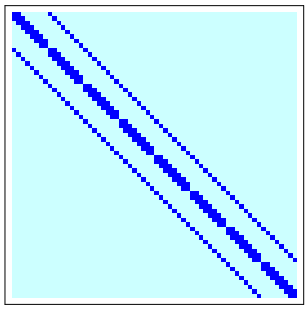

Reducing Dimension

The matrices are sparse and we can avoid large dimensions as follows.

Consider a configuration

$$X = (X_1, X_2, \ldots, X_n) \qquad X_i \in \mathbf{2}^m$$

Then X is in the kernel of σ iff

$$X_{i+1} = \sigma_m(X_i) + X_{i-1}$$

where σ_m is the operator on the path P_m.

Fibonacci Polynomials

But then the first row Z determines the rest of X.

Define the binary Fibonacci polynomials over $\mathbb{F}_2[x]$ as follows:

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

For example, $\pi_{51}(x)$ has the form

$$x^{50} + x^{48} + x^{44} + x^{42} + x^{40} + x^{34} + x^{32} + x^{12} + x^{10} + x^8 + x^2 + 1$$

Coefficients of Fibonacci Polynomials

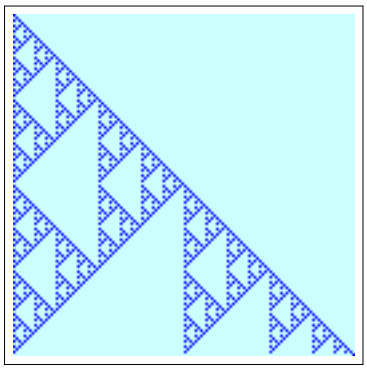

Expanding Patterns

Given row $Z \in \mathbf{2}^m$ define

$$X = (\pi_1(\sigma_m)(Z), \pi_2(\sigma_m)(Z), \ldots, \pi_n(\sigma_m)(Z)),$$

Then X is in the kernel of $\sigma_{n,m}$ iff

$$\pi_{n+1}(\sigma_m)(Z) = \mathbf{0}$$

It follows that the corank of $\sigma_{n,m}$ is at most $\min(n, m)$.

We can now compute kernel patterns and coranks relatively easily.

400 by 400

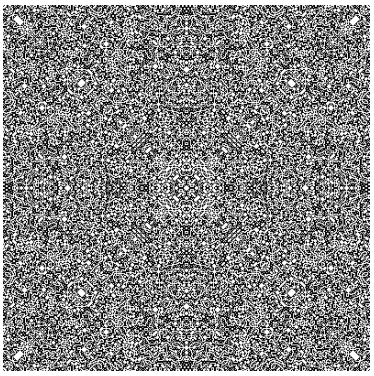

A kernel pattern for σ on the 400×400 grid.

Reversibility for σ

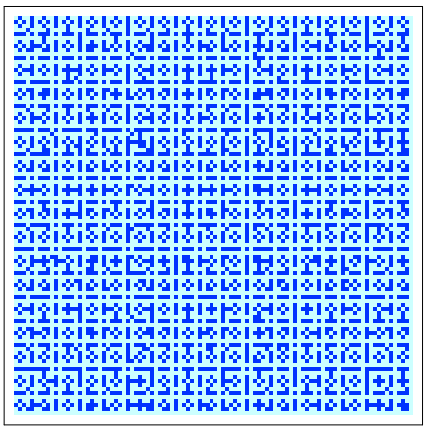

Reversibility for σ^-

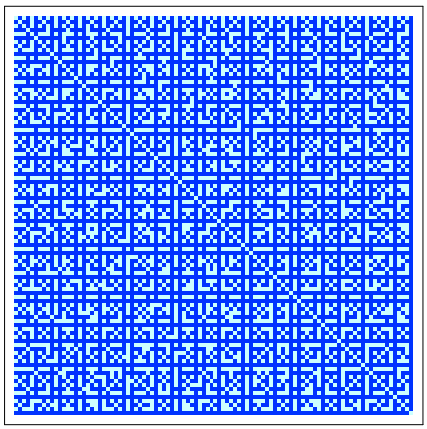

σ^-: Integer Division

The last picture should look eminently familiar: it's the GCD modulo 2.

A little check shows that the corank of σ^- on $P_{n,m}$ is

$$\gcd(n+1, m+1) - 1.$$

Computationally this is an ideal answer: all we need is integer division, can be handled in time polynomial in $\log n + \log m$.

σ^- is easy . . .

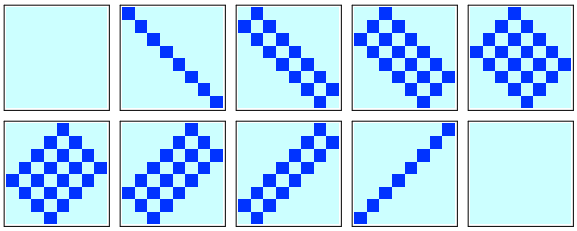

The canonical matrix representations of $\pi_i(\sigma_8^-)$, $0 \leq i < 10$.

So how much more difficult can $\sigma = \sigma^- + I$ be?

σ not so much

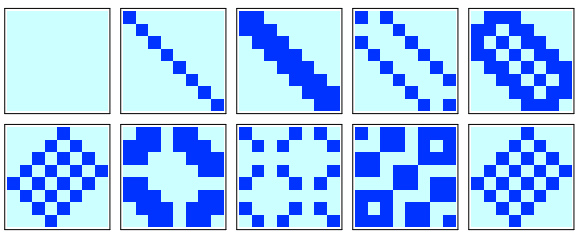

The same picture for σ_8, apparently much more complicated.

$m = 10$

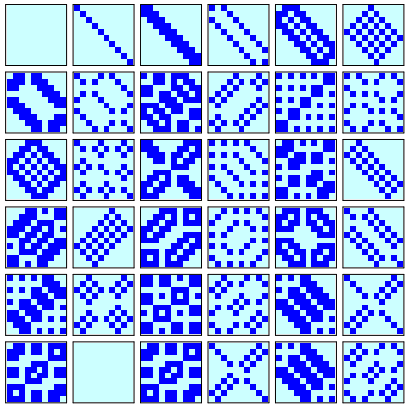

σ: Polynomial Division

Theorem (2000)

The dimension of the kernel of σ on $P_{n,m}$ is

$$\deg \gcd(\pi_{n+1}(x), \pi_{m+1}(x+1)).$$

Note the involution $x \mapsto x+1$ in the second term.

Algorithmically, time is now polynomial in n, not $\log n$.

Easy Divisibility Properties

Proposition

$$\pi_p = \pi_{q+1}\pi_{p-q} + \pi_q\pi_{p-q-1} \quad \textit{where } p \geq q+1$$
$$\gcd(\pi_n, \pi_m) = \pi_{\gcd(n,m)}$$
$$\pi_{2^k n} = x^{2^k-1}\,\pi_n^{2^k}$$
$$\pi_{2n+1} = \pi_{n+1}^2 + \pi_n^2$$
$$\pi_{2^k-1}\pi_{2^k+1} = (x^{2^k-1}+1)^2$$

The Chebyshev Angle

We can also think of these polynomials as the binary version of Chebyshev polynomial of the second kind:

$$\pi_n(x) = U_{n1}(x/2) \bmod 2$$

Slightly useful since we can lift results about Chebyshev polynomials. For example,

$$\pi_n(x) = \frac{\alpha^n + \beta^n}{\alpha + \beta}$$

where α and β are roots of the characteristic polynomial $t^2 + xt + 1$ (in the algebraic closure of of the rational function field).

Factorization

Theorem (2000)

Let $n = 2^k \cdot m$, m odd. Then

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

Here the critical factors $\rho_d(x)$ are products of irreducibles.

In particular for odd n we have

$$\pi_n(x) = \prod_{d|n} \rho_d(x)$$

Rank of Apparition

Given a sequence (a_n) the rank of apparition of m to be the least n such that

$$m \text{ divides } a_n$$

if it exists. The rank of apparition is particularly important for m prime.

In our case, we are interested in the least n such that some irreducible $\tau \in \mathbb{F}_2[x]$ divides π_n.

Dates back at least to work by M. Ward and L. K. Durst; the generalized Lucas-Lehmer test is closely connected to divisibility questions about Lucas sequences.

Rank of Apparition for Fibonacci Polynomials

Proposition

Every irreducible polynomial divides some π_n.

So the rank of apparition is always defined. Indeed the critical factors have the form

$$\rho_d = \prod(\tau^2 \mid \tau \text{ irreducible }, \text{rap}(\tau) = d)$$

So we need to understand ranks of apparition.

Pinning down RAP

Theorem (2000)

Let $\tau \in \mathbb{F}_2[x]$ be an irreducible polynomial of degree d.

Then τ's rank of apparition k divides $2^d \pm 1$.

In fact, the rap divides $2^d - 1$ if, and only if, the linear term in τ vanishes. In either case, d is the suborder of 2 in the multiplicative group $\mathbb{Z}_k^*$.

Here are the frequencies of raps of all 99 irreducibles of degree 10.

rap	count	±
25	1	+
41	2	+
93	3	−
205	8	+
341	15	−
1023	30	−
1025	40	+

Application: Total Irreversibility

The corank of σ on $P_{n,m}$ is at most $\min(n, m)$. Call the automaton totally irreversible if it attains this bound.

Theorem (P. Sarkar 1996)

$n = 4$ is the only totally irreversible square.

The proof is quite hard.

Application: Reversibility Test

Theorem (2000)

For any fixed $m \geq 1$ there are positive integers $t_1, \ldots, t_r$ such that rule σ on $P_{n,m}$ is reversible if, and only if, none of the t_i divides $n + 1$.

The test can be handled in time polynomial in $\log n$, but computation of the t_i appears to require

- the factorization of $\pi_{m+1}(x)$,
- computation of the rank of apparition of the corresponding irreducible factors.

The Source of All Evil

. . . is the involution $x \mapsto x + 1$.

$$\pi_n(x) = \ldots \tau(x) \ldots$$
$$\pi_n(x+1) = \ldots \tau(x+1) \ldots$$

Clearly $\tau(x + 1)$ is again irreducible and has the same degree as $\tau(x)$.

But what is the rank of apparition of $\tau(x + 1)$?

τ	τ^+	rap	irreducible polynomial	τ	τ^+	rap	irreducible polynomial
1	17	257	$1+x+x^3+x^4+x^8$	16	7	257	$1+x+x^3+x^7+x^8$
2	12	51	$1+x^2+x^3+x^4+x^8$	17	1	51	$1+x^2+x^3+x^7+x^8$
3	15	257	$1+x+x^3+x^5+x^8$	18	19	257	$1+x+x^2+x^3+x^4+x^7+x^8$
4	29	255	$1+x^2+x^3+x^5+x^8$	19	18	257	$1+x+x^5+x^7+x^8$
5	13	255	$1+x^3+x^4+x^5+x^8$	20	21	85	$1+x^3+x^5+x^7+x^8$
6	10	257	$1+x+x^2+x^3+x^4+x^5+x^8$	21	20	255	$1+x^4+x^5+x^7+x^8$
7	16	255	$1+x^2+x^3+x^6+x^8$	22	22	255	$1+x^2+x^3+x^4+x^5+x^7+x^8$
8	23	257	$1+x+x^2+x^3+x^4+x^6+x^8$	23	8	257	$1+x+x^6+x^7+x^8$
9	9	257	$1+x+x^5+x^6+x^8$	24	14	257	$1+x+x^2+x^3+x^6+x^7+x^8$
10	6	255	$1+x^2+x^5+x^6+x^8$	25	30	257	$1+x+x^2+x^4+x^6+x^7+x^8$
11	28	255	$1+x^3+x^5+x^6+x^8$	26	27	85	$1+x^2+x^3+x^4+x^6+x^7+x^8$
12	2	85	$1+x^4+x^5+x^6+x^8$	27	26	257	$1+x+x^2+x^5+x^6+x^7+x^8$
13	5	257	$1+x+x^2+x^4+x^5+x^6+x^8$	28	11	257	$1+x+x^4+x^5+x^6+x^7+x^8$
14	24	257	$1+x+x^3+x^4+x^5+x^6+x^8$	29	4	85	$1+x^2+x^4+x^5+x^6+x^7+x^8$
15	3	257	$1+x+x^2+x^7+x^8$	30	25	255	$1+x^3+x^4+x^5+x^6+x^7+x^8$

Ranks of apparition for all 30 irreducibles of degree 8. τ^+ is the image of τ under the involution $x \mapsto x+1$.

Involution Diagram

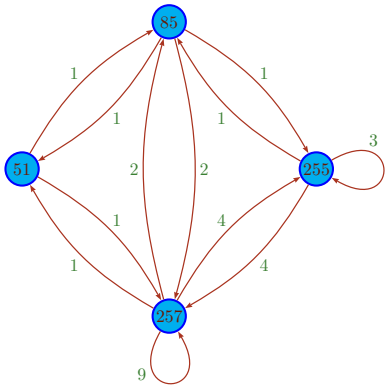

- σ Automata
- Reversibility
- π Polynomials

4. Irreversible Squares

- Orbits

Kernel Geometry

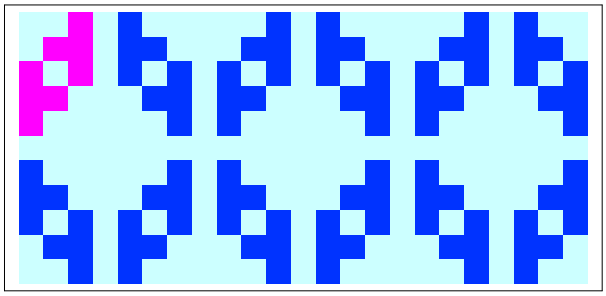

From 5×3 to 11×23.

A Multiplicative Sequence

Consider only irreversible squares:

$$\text{Irr} = \{\, n \mid \gcd(\pi_n(x), \pi_n(x+1)) \neq 1 \,\}$$

The first few terms of Irr are

$$5, 6, 10, 12, 15, 17, 18, 20, 24, 25, 30, 31, 33, 34, 35, 36, 40, 42, 45, 48, 50, 51, 54, 55, 60, 62, 63, 65, 66, 68, 70, 72, 75, 78, 80, 84, 85, 90, 93, 95, 96, 99, 100, \ldots$$

By kernel geometry, the sequence is multiplicatively closed.

Generators

So the question arises: what are the generators of Irr?

$$5, 6, 17, 31, 33, 63, 127, 129, 171, 257, 511, 683, 2047, 2731, 2979, 3277, \ldots$$

A Innocent Knuth Question:
Is Irr finitely generated?

No

Theorem (2006)

The sequence Irr *is not finitely generated.*

Sketch of proof:

$n \in$ Irr iff for some irreducible τ_1, τ_2 factors of π_n: $\tau_1(x) = \tau_2(x+1)$.

Focus on $\tau(x) = \tau(x+1)$: translation invariant irreducible polynomial (TIP).

So we only need to find lots of TIPs.

TIPs

We need to be able to construct TIPs at will. Sledgehammer approach: Let

$$\widehat{f}(x) = f(x(x+1))$$

Clearly invariant under the involution.

Lemma

Let f be irreducible of degree d.

Then either $\widehat{f}$ is TIP or $\widehat{f}(x) = f_1(x)\, f_2(x)$ where the f_i's are TIP.

Moreover, $\widehat{f}$ is TIP iff $[x^{d-1}]\widehat{f} = 1$.

Counting Irreducibles

Building on work by Niederreiter one can count irreducible polynomials of degree k with fixed coefficients c_1 and c_{k-1}.

The short answer: they exist.

Since you asked ...

The Long Answer

Let $\omega_{1/2} = (-1 \pm i\sqrt{7})/2$ be the two complex roots of $x^2 - x + 2 = 0$.

Lemma

Let $k \geq 4$ and write $k = k_0 k_1$ where k_0 is a power of 2 and k_1 is odd.

$$N_k^{ab} = \frac{1}{4k} \sum_{d \mid k} \mu(k_1/d) \left(2^{k_0 d} + (-1)^{a+b} (1 - \omega_1^{n_0 d} - \omega_2^{n_0 d}) - [a = b = 0, k_1 = 1]\, 4 \cdot 2^{n_0/2d} \right)$$

Infinitely Many Generators

Suppose there are finitely many generators $d_1, \ldots, d_s$ for Irr.

For any odd prime p there is a TIP f_p of degree $2p$ whose rank of apparition r divides $2^{2p} - 1$.

Hence $r \in$ Irr and we must have $d_i \mid r$.

But then there cannot be finitely many generators: for some $d = d_i$ we would have

$$d \mid \gcd(\mathsf{rap} f_p, \mathsf{rap} f_{p'}) \mid \gcd(2^{2p} - 1, 2^{2p'} - 1) = 3.$$

Contradiction. □

Analyzing Forward Orbits

To analyze forward orbits a convenient tool is the Fitting-decomposition of pattern space $\langle \mathcal{C}, \sigma \rangle$:

$$\mathcal{C} = K \oplus E$$

where K is the nilpotent part and E the regular part (with respect to linear operator σ).

- K corresponds to the transient part,
- E corresponds to the limit cycles.

Phase Space

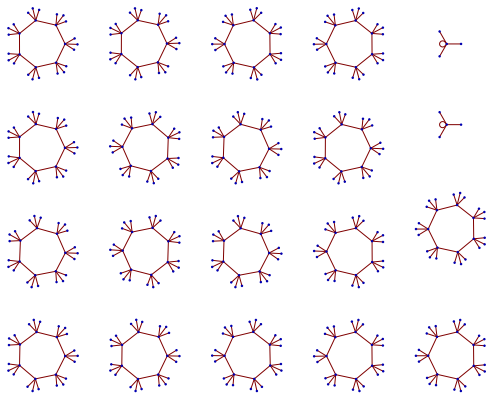

Trees and Cycles

The trees grafted on the limit cycles are isomorphic copies of K, the branching factor is $\mathrm{cork}(\sigma)$, their height is the nilpotency index:

$$\min(k \mid \mathrm{cork}(\sigma^k) = \mathrm{cork}(\sigma^{k+1}))$$

The limit cycles themselves correspond to the regular part: σ restricted to E is an automorphism.

We need to analyze E more closely to determine the structure of phasespace: elementary divisor decomposition.

The Key Ingredient

We need the minimal polynomial: the least degree polynomial $\mu \in \mathbb{F}_2[x]$ that annihilates σ:

$$\mu(\sigma) = 0$$

Given a factorization $\mu = \tau_1 \tau_2 \dots \tau_k$ we have a natural decomposition (elementary divisor decomposition)

$$\mathcal{C} = V_{\tau_1} \oplus V_{\tau_2} \oplus \dots \oplus V_{\tau_k}$$

into σ-invariant subspaces.

Minimal Polynomials

Theorem (2000)

The minimal polynomial of σ^- on a path of length n is $\pi_{n+1}(x)$.

The minimal polynomials for σ are then

$$\pi^+_{n+1}(x) = \pi_{n+1}(x+1)$$

and so forth.

Example: σ^- on P_{50}

$$\pi_{51}(x) = (x+1)^2 \left(x^4+x+1\right)^2 \left(x^4+x^3+x^2+x+1\right)^2 \left(x^8+x^4+x^3+x^2+1\right)^2 \left(x^8+x^7+x^3+x^2+1\right)^2$$

5 elementary divisor subspaces $E_1, \dots, E_5$, lengths and counts of pure cycles:

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

A little post-processing produces a count for all cycles.

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

Since the automaton is reversible this describes the whole phasespace.

Computational Walls

Theorem (1989)

Existence of a predecessor of bounded cardinality over $\mathbb{F}_2$ is $\mathbb{NP}$-complete.

Let $M = \langle\, a \mid a^2 = a^3 \,\rangle$ (three element Abelian monoid $\{0, 1, a\}$).

Theorem (1989)

Existence of a predecessor over M is $\mathbb{NP}$-complete.

Open Problems

- What exactly are the generators of Irr?
- Can reversibility of σ on an n by m grid be determined in time polynomial in $\log n$?
- How hard is it to compute the rank of apparition of $\tau(x + 1)$?
- Apply Fitting decomposition techniques to higher dimensional grids.
- Pin down the complexity of analyzing reversibility for algebraic dynamical systems.
- Pin down the complexity of analyzing the transition diagram for algebraic dynamical systems.