

**SOME RESULTS REGARDING FINITE LINEAR CELLULAR AUTOMATA
MODULO PRIME POWERS**

by

Zach Strong

Thesis submitted in partial fulfillment of the
requirements for the Degree of
Bachelor of Science with
Honours in Mathematics and Statistics

Acadia University

April 2025

© Copyright by Zach Strong, 2025

This thesis by Zach Strong
is accepted in its present form by the
Department of Mathematics and Statistics
as satisfying the thesis requirements for the degree of
Bachelor of Science with Honours

Approved by the Thesis Supervisor

Dr. Franklin Mendivil

Date

Approved by the Head of the Department

Dr. Richard Karsten

Date

Approved by the Honours Committee

Dr. Andrew Davis

Date

The author retains copyright in this thesis. Any substantial copying or any other actions that exceed fair dealing or other exceptions in the Copyright Act require the permission of the author.

Acknowledgements

I'd like to thank all my friends for putting up with my research-related ramblings throughout the past few years, even if they didn't understand a lick of what I was saying. Of course, I'd also like to thank Dr. Franklin Mendivil for supporting me through this project as well as through the last few years of my degree. Some credit should also go to Dr. Jeff Hooper and Dr. Richard Karsten for keeping Acadia's math department together. Finally, I'd like to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) for funding a large portion of this research through the Undergraduate Summer Research Award (USRA).

Contents

Abstract	xv
1 Motivating Examples	1
2 Background	7
2.1 Formal Representation of Finite Linear Cellular Automata	7
2.2 Vector Spaces and Modules	9
2.3 Iteration, Cycle Length, and Transient Length	12
2.4 Annihilating Polynomials & Minimal Polynomials	19
2.5 Theorems Relating to Minimal Polynomials	23
2.6 Lifting and Embedding	29
3 Understanding the Behaviour of Linear Cellular Automata	37
3.1 Multiplicative Orders	37
3.2 Cycle Converting Matrices	46
3.3 Relations Between Prime & Prime-Power Moduli	57
3.4 Cycle Spaces	64
4 Ideals of Annihilating Polynomials	73
4.1 When Does a “Minimal Polynomial” Exist?	74
4.2 Annihilating Ideal Generators	78

5	The Cores of Linear Cellular Automata	85
5.1	Dimensional Independence	87
5.2	Creating Dimensionally-Independent Sets	91
5.3	Prime-Power Cores	99
6	The Existence of Maximal Vectors	105
7	Conclusion	113
A	The Chinese Remainder Theorem	117
	Bibliography	121

List of Tables

List of Figures

1.1	The eight rules for the Wolfram 90 elementary cellular automata. Image taken from page 25 of <i>A New Kind of Science</i> (Wolfram [8]).	2
1.2	The first fifty time steps of the Wolfram 90 elementary cellular automata, visualised in the way described above. Image taken from page 25 of <i>A New Kind of Science</i> (Wolfram [8]).	3
1.3	Five time steps, starting from the left and working right, of a “glider” configuration in Conway’s Game of Life. Notice that the initial configuration is translated one unit diagonally down and to the right after these time steps, thus “gliding” the cells along the grid. This pattern of cell configurations will repeat indefinitely, “gliding” the cells across the board.	4
1.4	The configuration known as “Gosper’s Glider Gun”. After a certain number of time steps elapse, this configuration will “produce” a glider of the type shown in Figure 1.3 and return to the initial configuration, allowing for an infinite number of gliders to be produced over time.	4
2.1	On the left, a visual representation of a starting configuration’s sequence of time steps under the Wolfram 90 rule with four cells and cyclic boundary conditions. On the right, the same time steps are represented as vectors in the LCA $(\mathbb{Z}_2, \mathbb{Z}_2^4, [0\ 1\ 0\ 11\ 0\ 1\ 00\ 1\ 0\ 11\ 0\ 1\ 0])$	10
2.2	On the left, an example of a vector’s sequence of iterates when $\tau = 0$. On the right, an example of a vector’s sequence of iterates when $\tau = 2$. Notice that, when $\tau > 0$, some iterates will only ever be iterated to once.	14

Abstract

Our focus with this thesis will be on extending previous results obtained for finite linear cellular automata. Specifically, we will show that many properties of finite linear cellular automata with prime moduli also extend to the case of prime-power moduli. When the modulus is prime, the configuration space for our automata of interest forms a vector space, and many nice results regarding vector spaces can be directly utilised (such as the Primary Decomposition Theorem). However, when the modulus is a prime-power, our configuration space instead forms a module. While structurally very similar, modules are in general harder to work with. To overcome this difficulty, we will make use of the numerous connections between linear cellular automata with prime moduli and their corresponding systems with prime-power moduli (e.g. using a system modulo 5 to make conclusions about systems mod 25, 125, etc.). It turns out that, by exploring the similarities between systems with prime and prime-power moduli, we can show a lot to be true about the prime-power case which would otherwise be inaccessible if we focused solely on the prime-power case.

First, a few motivating examples will be discussed. Next, any relevant background information needed to understand the terminology, notation, or techniques used will be covered. The rest of the thesis will be dedicated to proving results about the following aspects of finite linear cellular automata: the multiplicative orders of matrices and vectors, the generators of annihilating polynomial ideals for vectors in the case of a prime-power moduli, the structure of a particular subset of vectors in our modules known as the “core”, and the existence of vectors with particular multiplicative orders known as “maximal vectors”. A brief section at the end will be set aside for mentioning possible avenues for future work on linear cellular automata.

Chapter 1

Motivating Examples

One of the chief concerns of mathematics is to describe and model complex behaviour. For instance, differential equations are one of math’s primary ways to model fluid flow and population dynamics, two very complex processes in the natural world. While differential equations give us a way to answer questions regarding complex phenomena, they also provide us with questions and complexities of their own. After all, the better we understand the models we use to describe the world, the better we understand the world itself. There’s a clear incentive to study not just the phenomenon of interest (fluid flow, population dynamics, etc.), but the models we create for them. The study of differential equations, consequently, is a wide, expansive topic, important enough to warrant a Millennium Prize Problem regarding them (solving the Navier-Stokes equations)!

We will not be studying differential equations in this thesis, but they illustrate an important point: in math, studying models on their own, separate from what they describe, is a worthwhile task. This will be the primary focus of this thesis: understanding a particular kind of model known as a *finite linear cellular automaton*.

A *cellular automata*, described intuitively by Stephen Wolfram on page 24 of *A New Kind of Science* (Wolfram [8]), is a set of coloured cells where “[a]t every [time] step there is a definite rule that determines the color of a given cell from the color of that and its immediate . . . neighbors on the step before”. In essence, a cellular automata is a set of cells that change state from time step to time step according to some simple rules.¹ Such a system

¹A more rigorous handling of cellular automata is given in Chapter 2.

Chapter 1. Motivating Examples

can be used to model phenomena where the behaviour is determined “locally”—that is, the behaviour at a particular location is determined solely by the state of the system around that location. Swarms of insects and the bulk behaviour of fluid are such examples.

Of course, the modelling capabilities of cellular automata can be quite useful, but it turns out that cellular automata have interesting-enough properties on their own, separate from any specific application. Despite cellular automata operating on *local* rules, many systems end up exhibiting *global* behaviour—that is, large-scale patterns that exceed the “reach” of the local rules imposed on the system.

The Wolfram 90 elementary cellular automata is an example of one such system. This automata makes use of an infinite strip of cells—cells arranged in a line. At each time step, cells change state according to their own state, as well as the states of their two neighbours (the cells to the immediate left and right of them on the strip). Figure 1.1 pictorially describes the local rules imposed on the Wolfram 90 cellular automata. The top three cells in each square represent a potential arrangement on our strip of cells, while the bottom cell represents what the middle cell in the line of three cells will become on the next time step.

As an example, the leftmost square says that, if at any time step there are three black cells in a row, then on the next time step, the middle of those three cells will become a white cell.² Using all eight rules provided, we have enough information to “update” the Wolfram 90 cellular automata, no matter the configuration the automata is in.




Figure 1.1: The eight rules for the Wolfram 90 elementary cellular automata. Image taken from page 25 of *A New Kind of Science* (Wolfram [8]).

To visualise the “evolution” of such an automata (that is, how the states of its cells change over time), we can “stack” the time steps on top of each other. The topmost row will represent the initial state of the automata, the row immediately below it will represent the time step after the initial configuration, the row below that will represent the time step after that, and so on. The farther down the image we look, the farther along in time

²Oftentimes, black cells are denoted as “alive” or “on” cells, while white cells are “dead” or “off”. This convention extends to the example of Conway’s Game of Life below.

the corresponding row represents. If our initial configuration for the Wolfram 90 cellular automata is a single black cell and an infinite number of white cells on either side of it, then our visualisation will look something like Figure 1.2.




Figure 1.2: The first fifty time steps of the Wolfram 90 elementary cellular automata, visualised in the way described above. Image taken from page 25 of *A New Kind of Science* (Wolfram [8]).

We notice that, despite the automata's rules only specifying how cells should evolve based on the states of itself and its immediate neighbours, we get a pattern that's consistent across the entire image. In fact, visualising the time steps in this fashion creates a Sierpinski triangle pattern. The local rules of the Wolfram 90 cellular automata create a global fractal pattern across time steps.

Emergent global behaviour is not unique to the Wolfram 90 cellular automata. Perhaps the most well-known example comes from Conway's Game of Life, a cellular automata which takes place on an infinite two-dimensional grid of cells rather than the infinite one-dimensional strip of the Wolfram 90 automata. In Conway's Game of Life, a white cell will turn black on the next time step if exactly three of its neighbours (the cells either orthogonally or diagonally touching it) are black, while a black cell will remain black on the next time step if two or three of its neighbours are black. Otherwise, a cell will turn white on the next time step.

Chapter 1. Motivating Examples




Figure 1.3: Five time steps, starting from the left and working right, of a “glider” configuration in Conway’s Game of Life. Notice that the initial configuration is translated one unit diagonally down and to the right after these time steps, thus “gliding” the cells along the grid. This pattern of cell configurations will repeat indefinitely, “gliding” the cells across the board.




Figure 1.4: The configuration known as “Gosper’s Glider Gun”. After a certain number of time steps elapse, this configuration will “produce” a glider of the type shown in Figure 1.3 and return to the initial configuration, allowing for an infinite number of gliders to be produced over time.

Via these simple local rules, extremely complex global behaviour arises. As an example, it’s possible to create configurations of cells that “glide” across the grid indefinitely like some strange cellular car. These configurations are commonly referred to as “gliders” (Patterson [5]). Figure 1.3 shows one such configuration. In fact, it’s actually possible to create configurations which *produce* gliders indefinitely, like a factory. The “Gosper’s Glider Gun” is one such configuration, shown in Figure 1.4 (Patterson [5]).

There are many, many interesting configurations one can discover within Conway’s Game of Life (Patterson [5]). The important part is that such configurations are possible at all, that the local rules governing Conway’s Game of Life give rise to global patterns and structures that extend beyond the scope of those initial local rules. Despite the definition of the cellular automata being very simple, very complex behaviour can emerge.

Through the examples of the Wolfram 90 and Game of Life cellular automata, we see

that cellular automata are interesting systems in their own right, even without coupling them with any particular application. The purpose of this thesis will be to explore some of the rich structure contained in a particular type of cellular automata: a *finite linear cellular automata*. In many ways, the automata we'll study are more complicated than the examples discussed here. Finite linear cellular automata allow cells to be in an *arbitrary*, finite number of states, not just “black” and “white”. As well, our automata will have periodic boundary conditions, meaning the cells won't exist on an infinitely-sized board, but rather one where the edges are connected. So, for instance, the leftmost cells on our boards will have the rightmost cells as their neighbours, and the topmost cells will have the bottom-most cells as neighbours. This will cause more interesting interactions to occur between cells, as a pattern that extends in one direction (like a glider) will inevitably appear on the opposite side of the board, allowing for “collisions” to occur that wouldn't be possible on an infinitely-sized board.

However, even with these added complications, the added property of linearity to our cellular automata will prove to be crucial to our findings. It's this added linearity that allows us to understand the behaviour of finite linear cellular automata much more completely than other cellular automata.

Chapter 1. Motivating Examples

Chapter 2

Background

2.1 Formal Representation of Finite Linear Cellular Automata

There are many different ways to conceptualise cellular automata. In the case of Conway's Game of Life, it's most helpful to visualise an infinite, two-dimensional grid composed of cells that can either be on or off. This is the standard way Conway's Game of Life is presented, and rightfully so. *Visually* showing how the grid of cells evolves over time is an effective way to demonstrate the emergent complex behaviour of the automata.

For formal analyses of cellular automata, a more abstract conceptualisation is typically used. A state space \mathcal{A} is defined, containing all the possible states a cell in the automata can have. In the case of Conway's Game of Life, this set would contain only two states: on and off. As well, some set Λ is used to index the collection of all cells in the automata. For Conway's Game of Life, this set would be a list of two-dimensional coordinate points, each coordinate representing a different cell. To assign a cell a state, we define a mapping that maps each cell indexed by Λ to a state in \mathcal{A} . These mappings are collected in a set denoted by

$$\mathcal{A}^\Lambda = \{f : \Lambda \rightarrow \mathcal{A}\}.$$

The update rule of the automata can then be represented as some function Φ that switches between particular given states to the automata's cells. In other words, the update rule is a

Chapter 2. Background

mapping between assigned mappings in \mathcal{A}^Λ :

$$\Phi : \mathcal{A}^\Lambda \rightarrow \mathcal{A}^\Lambda.$$

So, for Conway’s Game of Life, Φ can be viewed as taking the grid of cells at a particular time step as input and outputting the grid of cells for the next time step.

Note that the exact notation and treatment of cellular automata can change quite significantly between sources. The notation used here is adopted from “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]). For a slightly different set of conventions, see chapter 1.1 of *Algebraic Methods for Finite Linear Cellular Automata* (Dow [1]).

The cellular automata we’ll consider are one-dimensional, finite, *linear* cellular automata. Like the Wolfram 90 automata discussed in Chapter 1, our cells will be arranged in a strip, each cell taking on a particular state/value. Unlike the Wolfram 90 automata, our configurations will be finitely sized, meaning we’ll only have a certain number of cells to utilise. These types of automata lend themselves well to another type of notation: each configuration of our cellular automata can be treated as a vector of finite size.

The possible states for our cells will come from the integers modulo some positive integer (i.e. some set of numbers $\{0, 1, 2, \dots, N-1\}$). Let \mathbb{Z}_N denote the set of integers modulo N , and let \mathbb{Z}_N^L denote the set of $L \times 1$ vectors with components in \mathbb{Z}_N . Then, assuming L and N are set to be positive integers beforehand, each configuration for our cellular automata will be an element of \mathbb{Z}_N^L .¹

As well, since we’ll be restricting our attention to *linear* cellular automata, our update rules for these automata can be handled in a very particular way. Let \mathbf{A} denote our linear cellular automata’s update rule. For two given configurations of our automata $\vec{v}, \vec{w} \in \mathbb{Z}_N^L$, the update rule must abide by the following property (due to our automata being linear):

$$\mathbf{A}(c\vec{v} + d\vec{w}) \equiv c\mathbf{A}(\vec{v}) + d\mathbf{A}(\vec{w}) \pmod{N}$$

for constants $c, d \in \mathbb{Z}_N$.

¹Throughout this thesis, we’ll use the terms “configuration”, “initial configuration”, and “vector” interchangeably.

2.2. Vector Spaces and Modules

We see that our update rule \mathbf{A} is a linear transformation from \mathbb{Z}_N^L to \mathbb{Z}_N^L , meaning \mathbf{A} can be interpreted as an $L \times L$ matrix with components in \mathbb{Z}_N . Let $\mathbb{Z}_N^{L \times L}$ denote the set of $L \times L$ matrices with components in \mathbb{Z}_N . If N is omitted (e.g. $\mathbb{Z}^{L \times L}$), assume the components of the matrix come from the integers (i.e., \mathbb{Z}). Then, assuming L and N are set to be positive integers beforehand, the update rule for our finite linear cellular automata will be an element of $\mathbb{Z}_N^{L \times L}$. Applying the automata's update rule, then, is equivalent to multiplying our current vector by the matrix \mathbf{A} .

Now, we have everything we need to formally define finite linear cellular automata for our purposes.

Definition 2.1. A *finite linear cellular automata* (abbreviated as LCA) is a triplet (N, C, A) where N is the finite set of all possible states a cell can take (the *state space*), C is the finite set of all possible configurations (the *configuration space*), and A is the linear *update rule* or *update matrix* used. If an LCA takes the form $(\mathbb{Z}_n, \mathbb{Z}_n^L, \mathbf{A})$, then n is referred to as the *modulus* of the LCA.

As an example, for a finite version of the Wolfram 90 automata with cyclic boundary conditions and only four cells, its formal representation using Definition 2.1 would be $(\mathbb{Z}_2, \mathbb{Z}_2^4, \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix})$. Figure 2.1 illustrates both the “visual” way of representing the time steps of a particular Wolfram 90 starting configuration, as well as our more formal representation of the same time steps.

2.2 Vector Spaces and Modules

For LCAs of the form $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$, the ones we'll be considering, it's tempting to treat the configuration spaces as vector spaces. After all, vector spaces have plenty of nice properties, and applying those properties to an LCA would likely allow us to understand their behaviour better. However, in order for a configuration space \mathbb{Z}_N^L to be a vector space, the components of its vectors must come from a field. If N is taken to be a prime number, then \mathbb{Z}_N is indeed a field, and so \mathbb{Z}_N^L will be a vector space. If N is composite, then \mathbb{Z}_N is no longer a field, but a ring; the set \mathbb{Z}_N^L is not a vector space in this case, but a module, a slightly more general

Chapter 2. Background




Figure 2.1: On the left, a visual representation of a starting configuration's sequence of time steps under the Wolfram 90 rule with four cells and cyclic boundary conditions. On the right, the same time steps are represented as vectors in the LCA $\left(\mathbb{Z}_2, \mathbb{Z}_2^4, \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}\right)$.

object. Thus, unless N is taken to be a prime number, we cannot necessarily treat \mathbb{Z}_N^L as a vector space.

Consequently, the case where N is composite is much more difficult to analyse than the case where N is prime. To tackle the composite case, one of the most common approaches is to make use of the Chinese Remainder Theorem² to split up composite N into coprime, prime-power factors and analyse the corresponding LCAs with each coprime factor as their modulus. Using the behaviour of these resulting LCAs, we can deduce properties of the original LCA. In this way, we need only focus our attention on cases where N is prime, or N is a power of a prime.

When an LCA's configuration space is of the form $\mathbb{Z}_{p^k}^L$ for some prime-power p^k , it behaves as a module, not a vector space. However, it has many connections to the vector space \mathbb{Z}_p^L through various mappings, algebraic relations, and similar dynamics. These connections are the primary way in which the prime-power case is explored; they offer a means by which to bypass some of the shortcomings of having a module instead of a vector space. In later chapters, we'll use these sorts of connections to derive results for the prime-power case which would otherwise be unobtainable.

While we won't make extensive use of the formal definitions of rings, fields, vector spaces, and modules, they are included below for completeness.

²See Appendix A for an explanation as to how the Chinese Remainder Theorem is used to analyse LCAs.

2.2. Vector Spaces and Modules

Definition 2.2. A *ring* is a set S along with two operations, $+$: $S \times S \rightarrow S$ (addition) and \cdot : $S \times S \rightarrow S$ (multiplication) that satisfy the following properties:

- *Additive Identity:* There exists a $0 \in S$ such that, for all $a \in S$, $a + 0 = 0 + a = a$.
- *Additive Inverses:* For all $a \in S$, there exists a $-a \in S$ where $a + (-a) = (-a) + a = 0$.
- *Additive Associativity:* For all $a, b, c \in S$, $a + (b + c) = (a + b) + c$.
- *Additive Commutativity:* For all $a, b \in S$, $a + b = b + a$.
- *Multiplicative Distributivity:* For all $a, b, c \in S$, $a \cdot (b + c) = (a \cdot b) + (a \cdot c)$ and $(b + c) \cdot a = (b \cdot a) + (c \cdot a)$.

As well, to have an *associative ring* (which most rings are implicitly assumed to be), the following extra condition must hold:

- *Multiplicative Associativity:* For all $a, b, c \in S$, $a \cdot (b \cdot c) = (a \cdot b) \cdot c$.

Definition 2.3. A *field* is an associative ring $(S, +, \cdot)$ that satisfies the following properties:

- *Multiplicative Commutativity:* For all $a, b \in S$, $a \cdot b = b \cdot a$.
- *Multiplicative Identity:* There exists a $1 \in S$ such that, for all $a \in S \setminus \{0\}$, $1 \cdot a = a \cdot 1 = a$.
- *Multiplicative Inverses:* For all $a \in S \setminus \{0\}$, there exists a $a^{-1} \in S$ such that $a \cdot a^{-1} = a^{-1} \cdot a = 1$.

Definition 2.4. For an associative ring $(S, +, \cdot)$, a *module* over S is a set V with two operations, $+$: $V \times V \rightarrow V$ (vector addition) and \cdot : $S \times V \rightarrow V$ (scalar multiplication), that satisfies the following properties:

- *Additive Identity:* There exists a $\vec{0} \in V$ such that, for all $\vec{a} \in V$, $\vec{a} + \vec{0} = \vec{0} + \vec{a} = \vec{a}$.
- *Additive Inverses:* For all $\vec{a} \in V$ there exists a $-\vec{a} \in V$ where $\vec{a} + (-\vec{a}) = (-\vec{a}) + \vec{a} = \vec{0}$.
- *Additive Associativity:* For all $\vec{a}, \vec{b}, \vec{c} \in V$, $\vec{a} + (\vec{b} + \vec{c}) = (\vec{a} + \vec{b}) + \vec{c}$.
- *Additive Commutativity:* For all $\vec{a}, \vec{b} \in V$, $\vec{a} + \vec{b} = \vec{b} + \vec{a}$.
- *Scalar Sum Distributivity:* For all $x, y \in S$ and $\vec{a} \in V$, $(x + y) \cdot \vec{a} = x \cdot \vec{a} + y \cdot \vec{a}$.
- *Vector Sum Distributivity:* For all $\vec{a}, \vec{b} \in V$ and $x \in S$, $x \cdot (\vec{a} + \vec{b}) = x \cdot \vec{a} + x \cdot \vec{b}$.
- *Scalar Multiplication Associativity:* For all $x, y \in S$ and $\vec{a} \in V$, $x \cdot (y \cdot \vec{a}) = (x \cdot y) \cdot \vec{a}$.
- *Scalar Multiplication Identity:* There exists a $1 \in S$ such that, for all $\vec{a} \in V$, $1 \cdot \vec{a} = \vec{a}$.

Chapter 2. Background

Definition 2.5. A *vector space* is a module over S where $(S, +, \cdot)$ is a field.

2.3 Iteration, Cycle Length, and Transient Length

It is helpful to establish some terminology to describe some of the long-term behaviours of LCA configurations. This section will define three important ideas which underpin much of the work in describing LCAs and their properties.

Definition 2.6. An *iteration* is an application of an LCA’s update rule (i.e. a multiplication by the update matrix followed by a modular reduction by the modulus). To *iterate* is to apply an LCA’s update rule. An *iterate* is the result obtained by iteration via an LCA’s update rule.

By convention, we’ll consider a configuration to have itself as an iterate. That is, if \vec{v} is our configuration and \mathbf{A} is our update matrix, then $\mathbf{A}^0 \vec{v} \equiv \vec{v}$ is an iterate of \vec{v} .

Oftentimes, we’re not interested in the specific result of applying an update rule to a starting configuration, but rather the long-term behaviour of that starting configuration under repeated iteration of the update matrix. What happens to a vector if we keep repeatedly applying the update rule? Will it iterate back to its initial value? How many iterations will it take? Using the term “iteration” to mean applying an LCA’s update rule is to emphasise the iterative nature of repeatedly applying the update rule to an initial configuration to answer these sorts of questions.

A note on Definition 2.6: the language defined by this definition can apply to both configurations of an LCA and to an LCA’s update matrix itself. Since applying the LCA’s update rule is equivalent to multiplying by the update matrix, it makes perfect sense to iterate the update matrix—we take the product of the matrix and itself.

Briefly, let’s consider one of the questions posed above: given an initial configuration for some LCA, if we iterate enough times, will it eventually iterate back to itself? We know the configuration space for any LCA we consider will be a set of the form \mathbb{Z}_N^L , and this set has exactly N^L elements, meaning it is finite. Thus, we know that the set of unique iterates for a given initial configuration must also be finite since the set of iterates is a

2.3. Iteration, Cycle Length, and Transient Length

subset of the configuration space. So, it's within the realm of possibility that one of the initial configuration's iterates—excluding the “trivial” iterate, where the update rule is not applied—will be itself.

Let $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$ be the generic LCA we're considering, and let $\vec{v} \in \mathbb{Z}_N^L$ be our initial configuration. Then, the set of unique iterates³ for \vec{v} can be represented as

$$\mathcal{I} = \{\vec{v}, \mathbf{A}\vec{v}, \mathbf{A}^2\vec{v}, \dots, \mathbf{A}^c\vec{v}\} \pmod{N}$$

for some nonnegative integer c . Consider the “last” iterate $\mathbf{A}^c\vec{v}$. What happens if we apply the LCA's update rule to this configuration? We know that \mathcal{I} contains all unique iterates for the vector \vec{v} , so iterating $\mathbf{A}^c\vec{v}$ must give one of the vectors in \mathcal{I} . Therefore,

$$\mathbf{A}(\mathbf{A}^c\vec{v}) \equiv \mathbf{A}^\tau\vec{v} \pmod{N}$$

for some integer $0 \leq \tau \leq c$.

If $\tau = 0$, then the iterates of \vec{v} form a loop: repeatedly iterating \vec{v} will give $\mathbf{A}\vec{v}$, then $\mathbf{A}^2\vec{v}$, then $\mathbf{A}^3\vec{v}$, etc., until we get $\mathbf{A}^c\vec{v}$, at which point iterating again will give \vec{v} . Further iteration will simply cycle back through the same vectors since \vec{v} is our starting configuration. Thus, when $\tau = 0$, \vec{v} will eventually iterate back to itself.

Otherwise, if $\tau > 0$, then the iterates of \vec{v} will still form a sort of loop, but that loop won't include every possible iterate. For instance, consider the case where $\tau = 2$. Repeatedly iterating \vec{v} gives us $\mathbf{A}\vec{v}$, then $\mathbf{A}^2\vec{v}$, then $\mathbf{A}^3\vec{v}$, etc., the same as before, until we reach $\mathbf{A}^c\vec{v}$. If we iterate $\mathbf{A}^c\vec{v}$, we'll get $\mathbf{A}^2\vec{v}$, which is *not* the same as our initial configuration \vec{v} . Thus, if we continue iterating, we'll never get \vec{v} nor $\mathbf{A}\vec{v}$ as configurations again; only the vectors $\mathbf{A}^2\vec{v}$ through to $\mathbf{A}^c\vec{v}$ are obtainable via further applications of the update rule.

The value of τ , then, specifies the power of the update matrix \mathbf{A} that separates recurring iterates in the iteration of \vec{v} from one-time iterates that only occur a single time, no matter how many times the update rule is applied to \vec{v} . We call the value τ the *transient length* of \vec{v} under \mathbf{A} , as it counts the number of configurations that are “transient”, or non-permanent, in the long-term iteration of \vec{v} .

³This set is sometimes referred to as the *orbit* of \vec{v} under \mathbf{A} .

Chapter 2. Background




Figure 2.2: On the left, an example of a vector's sequence of iterates when $\tau = 0$. On the right, an example of a vector's sequence of iterates when $\tau = 2$. Notice that, when $\tau > 0$, some iterates will only ever be iterated to once.

An LCA's update matrix can also have a transient length. Below are formal definitions for both types of transient lengths.

Definition 2.7. The *transient length* of a matrix \mathbf{A} modulo N is the smallest nonnegative integer τ such that the congruence

$$\mathbf{A}^c \mathbf{A}^\tau \equiv \mathbf{A}^\tau \pmod{N}$$

has positive integer solutions for c .

Definition 2.8. The *transient length* of a vector $\vec{v} \in \mathbb{Z}_N^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_N^{L \times L}$ modulo N is the smallest nonnegative integer τ such that the congruence

$$\mathbf{A}^c \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v} \pmod{N}$$

has positive integer solutions for c .

Sometimes, it is useful to give a name to the set of configurations to which a vector will only ever iterate once, as the number of vectors in this set is, by construction, exactly the same as the vector's transient length.

Definition 2.9. Given a vector $\vec{v} \in \mathbb{Z}_N^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_N^{L \times L}$ modulo N , the *transient*

2.3. Iteration, Cycle Length, and Transient Length

region of \vec{v} is the set

$$\bigcup_{i=0}^{\tau-1} \{\mathbf{A}^i \vec{v}\} \pmod{N},$$

where τ is the transient length of \vec{v} .

From these definitions, and from the fact that the update rules of our LCAs are linear, we can already begin to draw conclusions about the general behaviour of LCA configurations. Proposition 2.1 provides an example of the types of reasoning we can employ.

Proposition 2.1. Let $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$ be an LCA, and let $\vec{v} \in \mathbb{Z}_N^L$ be a vector with transient length $\tau > 0$ such that $\mathbf{A}^\tau \vec{v} \not\equiv \vec{0} \pmod{N}$. Then there exists another vector $\vec{w} \in \mathbb{Z}_N^L$ with transient length τ such that $\mathbf{A}^\tau \vec{w} \equiv \vec{0} \pmod{N}$.

Proof. By definition of a vector's transient length, there exists a natural number c such that

$$\mathbf{A}^c \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v} \pmod{N}.$$

Consider the vector \vec{w} defined by

$$\vec{w} \equiv \mathbf{A}^{c-(\tau \bmod c)} \mathbf{A}^\tau \vec{v} - \vec{v} \pmod{N}.$$

The vector \vec{w} cannot be the zero vector since, if it were, it would imply that

$$\mathbf{A}^{c-(\tau \bmod c)} \mathbf{A}^\tau \vec{v} \equiv \vec{v} \pmod{N},$$

and this would contradict the fact that the transient length of \vec{v} is greater than zero.

As well, for $k \geq \tau$, we notice that $\mathbf{A}^k \vec{w} \equiv \vec{0} \pmod{N}$. To see this, let $k = \tau + t$ for some

Chapter 2. Background

nonnegative integer t . Then,

$$\begin{aligned}
& \mathbf{A}^k \vec{w} \\
& \equiv \mathbf{A}^{\tau+t} (\mathbf{A}^{c-(\tau \bmod c)} \mathbf{A}^\tau \vec{v} - \vec{v}) \\
& \equiv \mathbf{A}^t \mathbf{A}^{nc} \mathbf{A}^\tau \vec{v} - \mathbf{A}^t \mathbf{A}^\tau \vec{v} \\
& \equiv \mathbf{A}^t \mathbf{A}^\tau \vec{v} - \mathbf{A}^t \mathbf{A}^\tau \vec{v} \\
& \equiv \vec{0} \pmod{N}
\end{aligned}$$

using the definition of vector transient lengths, where n is some integer. Furthermore, $\mathbf{A}^k \vec{w} \not\equiv \vec{0} \pmod{N}$ for integers $0 \leq k < \tau$ since, if $\mathbf{A}^k \vec{w} \equiv \vec{0} \pmod{N}$, then we'd have that

$$\mathbf{A}^k \mathbf{A}^{c-(\tau \bmod c)} \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^k \vec{v} \pmod{N},$$

and this would contradict the fact that τ is the transient length of \vec{v} .

Now, consider the congruence

$$\mathbf{A}^x \mathbf{A}^T \vec{w} \equiv \mathbf{A}^T \vec{w} \pmod{N}. \tag{2.1}$$

We claim that there are no positive integer solutions for x when $T < \tau$. To see this, assume $T < \tau$, and assume we have some positive integer solution for x . Then, repeatedly applying Congruence (2.1) gives us that

$$\mathbf{A}^{x\tau} \mathbf{A}^T \vec{w} \equiv \mathbf{A}^T \vec{w} \not\equiv \vec{0} \pmod{N}$$

since $T < \tau$. However, $x\tau \geq \tau$, and so by what we showed above,

$$\mathbf{A}^{x\tau} \mathbf{A}^T \vec{w} \equiv \mathbf{A}^T (\mathbf{A}^{x\tau} \vec{w}) \equiv \mathbf{A}^T (\vec{0}) \equiv \vec{0} \pmod{N},$$

which is a contradiction.

Now, simply plugging in $T = \tau$ into Congruence (2.1) shows that τ is the first value of T where Congruence (2.1) has solutions for x (since $\mathbf{A}^\tau \vec{w} \equiv \vec{0}$), and so, by definition, τ is the transient length of \vec{w} . **QED**

2.3. Iteration, Cycle Length, and Transient Length

Proposition 2.1 demonstrates the power of defining quantities/properties of interest; with good definitions, certain behaviours of LCAs will make themselves known. This becomes even more apparent in Chapter 5, where we define a generalisation of linear independence to prove some properties of an important set of vectors within an LCA’s configuration space (known as the “core”). The remainder of this chapter will focus on introducing other definitions and concepts that prove to be useful in the analysis of LCAs.

Closely related to the transient length is the *cycle length* or *multiplicative order* of a vector/matrix. The transient length of a matrix \mathbf{A} is defined to be the smallest nonnegative integer τ such that

$$\mathbf{A}^c \mathbf{A}^\tau \equiv \mathbf{A}^\tau$$

has positive integer solutions for c . The smallest such c where this congruence is true is called the *cycle length* or *multiplicative order* of \mathbf{A} . The cycle length of a vector under a matrix is defined similarly.

Definition 2.10. The *cycle length*, or *multiplicative order*, of a matrix \mathbf{A} modulo N is the smallest positive integer ω such that

$$\mathbf{A}^\omega \mathbf{A}^t \equiv \mathbf{A}^t \pmod{N},$$

where t is the transient length of \mathbf{A} modulo N .

Definition 2.11. The *cycle length*, or *multiplicative order*, of a vector $\vec{v} \in \mathbb{Z}_N^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_N^{L \times L}$ modulo N is the smallest positive integer ω such that

$$\mathbf{A}^\omega \mathbf{A}^t \vec{v} \equiv \mathbf{A}^t \vec{v} \pmod{N},$$

where t is the transient length of \vec{v} under the matrix \mathbf{A} modulo N .

The term “cycle length” comes from a particular interpretation of this quantity. If we return to Figure 2.2, the cycle length of the vector \vec{v} is given by the number of vectors within the “loop” (for the diagram on the right, this does *not* include the “tail” that feeds into the

Chapter 2. Background

loop). Then, the cycle length of a vector is the number of iterations needed for the vector to “cycle back” to a previous iterate, hence the name. Likewise for the cycle length of a matrix.

The term “multiplicative order” comes from the similarity in this quantity’s definition to the multiplicative order of elements in a ring.

In this thesis, the terms “cycle length” and “multiplicative order” will be used interchangeably for vectors and matrices.

Notice that, for the definitions of cycle length, t is able to be greater than the transient length of the vector/matrix. The reason for this is straightforward; we’ll use the matrix case to illustrate, though the vector case is nearly identical. If we let τ be the transient length of our matrix \mathbf{A} , and ω its cycle length, then after iterating the matrix τ times, the iterates that follow can be listed as

$$\mathbf{A}\mathbf{A}^\tau, \mathbf{A}^2\mathbf{A}^\tau, \dots, \mathbf{A}^{\omega-1}\mathbf{A}^\tau, \mathbf{A}^\omega\mathbf{A}^\tau \equiv \mathbf{A}^\tau, \mathbf{A}\mathbf{A}^\tau, \mathbf{A}^2\mathbf{A}^\tau, \dots$$

Say we start at the iterate $\mathbf{A}^i\mathbf{A}^\tau$ in this list. It’ll take $\omega - i$ iterations to iterate to \mathbf{A}^τ , and then another i iterations to iterate back to $\mathbf{A}^i\mathbf{A}^\tau$. Thus, for any starting iterate in this list, it takes $\omega - i + i = \omega$ iterations to iterate back where we started. Looking back at our definition for cycle length, we see that this implies t can be any value greater than or equal to the transient length, and the obtained value of the cycle length will remain the same.

We point out this fact for ease of explanation later in this thesis. In most cases, the value of t in the definitions of cycle length will be taken to be the transient length of the respective vector/matrix, but it is sometimes useful to use a value greater than the transient length.

One other detail worth mentioning. If the transient length of a matrix \mathbf{A} is zero, then $\mathbf{A}^\omega \equiv \mathbf{I}$, where ω is the cycle length of \mathbf{A} and \mathbf{I} is the identity matrix of the appropriate size. This follows directly from the definition of a matrix’s cycle length. Thus, this means that \mathbf{A} must be invertible, since $\mathbf{A}^{\omega-1}\mathbf{A} \equiv \mathbf{A}^\omega \equiv \mathbf{I}$.

In fact, the converse is also true. If \mathbf{A} is invertible, then there exists an \mathbf{A}^{-1} such that

2.4. Annihilating Polynomials & Minimal Polynomials

$\mathbf{A}^{-1}\mathbf{A} \equiv \mathbf{I}$. Then, if ω is the cycle length of \mathbf{A} and τ is the transient length, we have that

$$\begin{aligned}\mathbf{A}^\omega \mathbf{A}^\tau &\equiv \mathbf{A}^\tau \\ \implies \mathbf{A}^{-\tau}(\mathbf{A}^\omega \mathbf{A}^\tau) &\equiv \mathbf{A}^{-\tau} \mathbf{A}^\tau \\ \implies \mathbf{A}^\omega &\equiv \mathbf{I},\end{aligned}$$

which means that $\tau = 0$ by definition.

Thus, we've proven the following proposition:

Proposition 2.2. A matrix $\mathbf{A} \in \mathbb{Z}_N^{L \times L}$ is invertible modulo N if and only if its transient length is zero.

This is another example of a definition (this time being the definition for the cycle length) proving to be useful in deducing properties regarding LCAs.

2.4 Annihilating Polynomials & Minimal Polynomials

Perhaps one of the more interesting ways we can analyse LCAs is by considering some special polynomials related to an LCA's update matrix.

Consider some arbitrary LCA $(\mathbb{Z}_p, \mathbb{Z}_p^L, \mathbf{A})$, where p is an odd prime. The matrix \mathbf{A} has a characteristic polynomial given by

$$\mu(\lambda) \equiv \det(\lambda \mathbf{I} - \mathbf{A}) \pmod{p}.$$

By the Cayley-Hamilton Theorem, $\mu(\lambda)$ is an *annihilating polynomial* for \mathbf{A} .

Definition 2.12. An *annihilating polynomial* of a matrix \mathbf{A} is any polynomial $r(x)$ such that $r(\mathbf{A}) = \mathbf{0}$, where $\mathbf{0}$ is the zero matrix of appropriate size.

Any polynomial multiple of $\mu(\lambda)$ is also an annihilating polynomial for \mathbf{A} . To see this,

Chapter 2. Background

simply let $\nu(\lambda)$ be an arbitrary polynomial and compute $\mu(\mathbf{A})\nu(\mathbf{A})$:

$$\begin{aligned}\mu(\mathbf{A})\nu(\mathbf{A}) &\equiv (\mathbf{0})\nu(\mathbf{A}) \pmod{p} \\ &\equiv \mathbf{0} \pmod{p}.\end{aligned}$$

In fact, any polynomial multiple of any annihilating polynomial for \mathbf{A} will also be an annihilating polynomial for \mathbf{A} . Thus, the set of all annihilating polynomials for \mathbf{A} forms an ideal⁴ of the ring of polynomials with coefficients in \mathbb{Z}_p . Let $\mathbb{Z}_N[X]$ represent the set of all polynomials with coefficients in \mathbb{Z}_N , and let $\text{Ann}_{\mathbb{Z}_p[X]}(\mathbf{A})$ represent the ideal of annihilating polynomials in $\mathbb{Z}_p[X]$ for \mathbf{A} .

Since $\mathbb{Z}_p[X]$ is a principal ideal domain (by virtue of the Euclidean Algorithm and the fact that only one indeterminate is used), our ideal $\text{Ann}_{\mathbb{Z}_p[X]}(\mathbf{A})$ is generated by a single element, meaning every annihilating polynomial for \mathbf{A} is a polynomial multiple of some “smallest” polynomial. This “smallest” polynomial is called the *minimal polynomial* of \mathbf{A} .

Definition 2.13. The *minimal polynomial* of a matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, for primes p , is its lowest degree, monic, annihilating polynomial.

The minimal polynomial is chosen to be monic (having a coefficient of 1 on its leading term) so that the minimal polynomial is unique. Otherwise, if the monic condition was dropped, there would be $p - 1$ valid generators for $\text{Ann}_{\mathbb{Z}_p[X]}(\mathbf{A})$.

The minimal polynomial of a matrix specifies, in a sense, the smallest algebraic property the matrix satisfies, which in turn allows us to make sense of its behaviour. As an example, say we’re given a matrix \mathbf{A} whose minimal polynomial is given by

$$m(x) = x^2 + x + 2.$$

Now, say we’re given the polynomial expression $\mathbf{A}^4 + 2\mathbf{A}^3 + 2\mathbf{A} + 4\mathbf{I}$. This polynomial is fairly intimidating, but using $m(x)$, we can reduce this to a nicer expression. By definition,

⁴An *ideal* \mathfrak{I} is an additive subgroup of a ring R such that, for any $x \in R$ and any $y \in \mathfrak{I}$, $xy \in \mathfrak{I}$ and $yx \in \mathfrak{I}$.

2.4. Annihilating Polynomials & Minimal Polynomials

we know $m(x)$ is an annihilating polynomial for \mathbf{A} , so

$$\begin{aligned}\mathbf{A}^2 + \mathbf{A} + 2\mathbf{I} &= \mathbf{0} \\ \implies \mathbf{A}^2 &= -(\mathbf{A} + 2\mathbf{I}).\end{aligned}$$

Multiplying both sides of this equation by \mathbf{A} :

$$\begin{aligned}\mathbf{A}^3 &= -\mathbf{A}(\mathbf{A} + 2\mathbf{I}) \\ &= -(\mathbf{A}^2 + 2\mathbf{A}).\end{aligned}$$

From above, we have an expression for \mathbf{A}^2 . Substituting this into our expression for \mathbf{A}^3 , we get that

$$\begin{aligned}\mathbf{A}^3 &= -(-(\mathbf{A} + 2\mathbf{I}) + 2\mathbf{A}) \\ &= 2\mathbf{I} - \mathbf{A}.\end{aligned}$$

Multiplying both sides of this equation by \mathbf{A} and substituting our found expression for \mathbf{A}^2 , we can derive an expression for \mathbf{A}^4 :

$$\begin{aligned}\mathbf{A}^4 &= \mathbf{A}(2\mathbf{I} - \mathbf{A}) \\ &= 2\mathbf{A} - \mathbf{A}^2 \\ &= 2\mathbf{A} - (-(\mathbf{A} + 2\mathbf{I})) \\ &= 3\mathbf{A} + 2\mathbf{I}.\end{aligned}$$

Using $m(x)$, we now have expressions for \mathbf{A}^2 , \mathbf{A}^3 , and \mathbf{A}^4 which are linear. If we plug these into our given polynomial expression:

$$\begin{aligned}&\mathbf{A}^4 + 2\mathbf{A}^3 + 2\mathbf{A} + 4\mathbf{I} \\ &= (3\mathbf{A} + 2\mathbf{I}) + 2(2\mathbf{I} - \mathbf{A}) + 2\mathbf{A} + 4\mathbf{I} \\ &= 3\mathbf{A} + 10\mathbf{I}.\end{aligned}$$

Chapter 2. Background

We see that our complicated quartic expression is equivalent to $3\mathbf{A} + 10\mathbf{I}$, which is evidently a much simpler expression. Using the minimal polynomial of a matrix, we can perform these kinds of simplifications without having to know the specific value of the matrix. This, in turn, leads to more efficient computations involving an LCA's update matrix. As well, in Section 2.5, we'll see that the minimal polynomial of a matrix can also give us insight into an LCA's behaviour (e.g. how vectors in its configuration space iterate, what their cycle lengths are, etc.).

If minimal polynomials help us make sense of complicated expressions involving an LCA's update matrix, is it possible they could also help us make sense of complicated expressions involving LCA configurations? Unlike with matrices, we can't simply substitute a vector into a univariate polynomial and evaluate it. To do so, we'd need to define what it means to take powers of a vector (like \vec{v}^2 , \vec{v}^3 , etc.), and this is an operation that doesn't have a clear interpretation. Instead, we define an annihilating polynomial for a vector to be a polynomial expression of a *matrix* that, when the vector is multiplied by the expression, gives a product of zero.

For example, given a matrix \mathbf{A} , any vector \vec{v} that satisfies the equation

$$(\lambda\mathbf{I} - \mathbf{A})\vec{v} = \vec{0}$$

for some value λ is an eigenvector for \mathbf{A} with eigenvalue λ . Thus, $\lambda\mathbf{I} - \mathbf{A}$ is an annihilating polynomial for any eigenvector with eigenvalue λ .

So long as the state space of our LCA is a field (\mathbb{Z}_p , in our case), the ring of polynomials we're dealing with will remain a principal ideal domain, meaning the set of annihilating polynomials for a vector \vec{v} (denoted as $\text{Ann}_{\mathbb{Z}_p[X]}(\vec{v})$) will be a principal ideal, and thus a minimal polynomial will exist just like the matrix case.

Definition 2.14. An *annihilating polynomial* of a vector $\vec{v} \in \mathbb{Z}_p^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$ (for odd primes p) is any polynomial $r(x) \in \mathbb{Z}_p[X]$ such that $r(\mathbf{A})\vec{v} \equiv \vec{0} \pmod{p}$.

Definition 2.15. The *minimal annihilating polynomial* of a vector $\vec{v} \in \mathbb{Z}_p^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$ (for odd primes p) is the vector's lowest degree, monic, annihilating polynomial.

To ease with explanation later in this thesis, we'll say that an annihilating polynomial for a *set* of vectors/matrices is a polynomial that's annihilating for all the vectors/matrices

2.5. Theorems Relating to Minimal Polynomials

in the set. The minimal annihilating polynomial for a set is defined similarly: it will be the lowest degree, monic polynomial that's annihilating for all the vectors/matrices in the set.

An important observation about vector minimal annihilating polynomials is that they will always be a factor of the corresponding matrix's minimal polynomial. Why? Consider a vector $\vec{v} \in \mathbb{Z}_p^L$ with minimal annihilating polynomial $m_{\vec{v}}(x) \in \mathbb{Z}_p[X]$ under the matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$. As well, let $m_{\mathbf{A}}(x) \in \mathbb{Z}_p[X]$ be the matrix's minimal polynomial. By definition, we have that

$$m_{\mathbf{A}}(\mathbf{A})\vec{v} \equiv (\mathbf{0})\vec{v} \equiv \vec{0} \pmod{p},$$

and so $m_{\mathbf{A}}(x)$ is an annihilating polynomial for \vec{v} . However, we know $\text{Ann}_{\mathbb{Z}_p[X]}(\vec{v})$ is a principal ideal generated by $m_{\vec{v}}(x)$, so $m_{\mathbf{A}}(x)$ must be a polynomial multiple of $m_{\vec{v}}(x)$. Thus, a vector's minimal polynomial will always be a factor of the corresponding matrix's minimal polynomial.

2.5 Theorems Relating to Minimal Polynomials

Using the idea of minimal polynomials, there are two important theorems we can employ to better understand LCAs of the form $(\mathbb{Z}_p, \mathbb{Z}_p^L, \mathbf{A})$ for odd primes p : mainly, the Primary Decomposition Theorem and the Minimal Polynomial Theorem. First, we'll look at the Primary Decomposition Theorem.

Theorem 1 (Primary Decomposition Theorem⁵). *Let \mathbb{K}^L be a vector space, $\mathbf{A} \in \mathbb{K}^{L \times L}$ a matrix, and $m(\lambda)$ the minimal polynomial for \mathbf{A} . If*

$$m(\lambda) = \prod_{i=1}^k (f_i(\lambda))^{n_i}$$

for monic and relatively prime polynomials $f_1(\lambda)$ to $f_k(\lambda)$ and positive integers n_1 to n_k , then

$$\mathbb{K}^L = \bigoplus_{i=1}^k \ker((f_i(\mathbf{A}))^{n_i}) \quad (2.2)$$

⁵See Theorem 12 of “Linear algebra” (Hoffman and Kunze [2]).

Chapter 2. Background

where each $\ker((f_i(\mathbf{A}))^{n_i}) \neq \{\vec{0}\}$.

For our purposes, Theorem 1 gives us a way to break up an LCA's configuration space into simpler pieces (subspaces) that are easier to analyse. From Equation (2.2), we see that each relatively prime factor of an update matrix's minimal polynomial corresponds to a subspace of the configuration space. What exactly is simpler about these subspaces? Well, each subspace is defined to be the kernel to some polynomial expression (a factor of the minimal polynomial), and so all the vectors in each respective subspace will, by definition, be annihilated by the polynomial expression. What this means is that, if we restrict our attention to any of the subspaces defined by Equation (2.2), the update matrix will behave, with respect to the vectors, as though its minimal polynomial is the polynomial expression that defines the subspace.

A specific example helps to make sense of Theorem 1. Consider the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, [\begin{smallmatrix} 3 & 4 \\ 2 & 0 \end{smallmatrix}])$. The minimal polynomial of the update matrix is given by

$$m(\lambda) = 2 + 2\lambda + \lambda^2 \equiv (4 + \lambda)(3 + \lambda) \pmod{5}.$$

The factors $4 + \lambda$ and $3 + \lambda$ are monic and relatively prime, and so the Primary Decomposition Theorem says we can break up our configuration space as the direct sum

$$\mathbb{Z}_5^2 = \ker(4 + \mathbf{A}) \oplus \ker(3 + \mathbf{A}).$$

How do the vectors in each subspace behave? Let's focus on $\ker(4 + \mathbf{A})$ as an example. By construction, all vectors within this subspace have $4 + \lambda$ as an annihilating polynomial, and so for any $\vec{v} \in \ker(4 + \mathbf{A})$:

$$\begin{aligned} (4\mathbf{I} + \mathbf{A})\vec{v} &\equiv \vec{0} \\ \implies \mathbf{A}\vec{v} &\equiv -4\vec{v} \equiv \vec{v} \pmod{5}. \end{aligned}$$

We see that, within this subspace, iterating by the update matrix does nothing to the starting vector; multiplying a vector by \mathbf{A} simply returns the same vector. Thus, within the subspace $\ker(4 + \mathbf{A})$ (which can be computed explicitly using a variety of methods, such as Gauss-Jordan Elimination), we have a complete description of the dynamics of all possible vectors

2.5. Theorems Relating to Minimal Polynomials

under iteration by the update matrix: iterating by \mathbf{A} leaves the vector alone.

Using the same reasoning, we can determine that all vectors within $\ker(3 + \mathbf{A})$ get scaled by 2 under iteration by \mathbf{A} . While the dynamics of these vectors are a tad more complex, we've still successfully reduced the complexity of the LCA's update rule down to one-dimensional dynamics. Rather than iterate vectors by a 2×2 update matrix, we can equivalently multiply by a scalar to obtain the same iterates within this particular subspace.

Knowing the simpler behaviour of vectors within these subspaces now allows us to understand the behaviour of *all* vectors in our configuration space. The Primary Decomposition Theorem provides a *direct* sum decomposition of the entire configuration space into these simpler subspaces, so any vector in the configuration space can be represented as exactly one sum of vectors where exactly one vector from each separate subspace is used.⁶ For our example LCA, $(\mathbb{Z}_5, \mathbb{Z}_5^2, [\begin{smallmatrix} 3 & 4 \\ 2 & 0 \end{smallmatrix}])$, this means every vector in \mathbb{Z}_5^2 can be expressed uniquely in the form $\vec{u} + \vec{v}$, where $\vec{u} \in \ker(4 + \mathbf{A})$ and $\vec{v} \in \ker(3 + \mathbf{A})$.

As well, each of the Primary Decomposition Theorem's kernels are invariant under multiplication by \mathbf{A} , meaning the iterates of vectors within a kernel will stay within the same kernel. This is straightforward to show. Consider an arbitrary kernel $\ker(f(\mathbf{A}))$ and some vector \vec{v} within it. Then,

$$\begin{aligned} \vec{v} &\in \ker(f(\mathbf{A})) \\ \implies f(\mathbf{A})\vec{v} &\equiv \vec{0} \\ \implies f(\mathbf{A})(\mathbf{A}^i\vec{v}) &\equiv \mathbf{A}^i(f(\mathbf{A})\vec{v}) \equiv \mathbf{A}^i(\vec{0}) \equiv \vec{0} \\ \implies \mathbf{A}^i\vec{v} &\in \ker(f(\mathbf{A})). \end{aligned}$$

So, if \vec{v} is in a kernel, then all its iterates are, too.

For us, this means the iterates of *any* vector in our configuration space are governed by the simpler actions of the update matrix within each of the Primary Decomposition Theorem's subspaces—if we write a vector as a sum of vectors from these subspaces, the iterates will remain in the same form, and no interaction between the vectors in the sum will occur. For

⁶This proves to be highly useful for constructing vectors in an LCA with particular properties. For instance, proposition 3 in “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]) uses the disjointedness of the subspaces described by the Primary Decomposition Theorem to construct a vector with maximal multiplicative order by summing together a vector from each subspace.

Chapter 2. Background

our example LCA, this gives the following relation: for any vector $\vec{w} \in \mathbb{Z}_5^2$ written as $\vec{u} + \vec{v}$, where $\vec{u} \in \ker(4 + \mathbf{A})$ and $\vec{v} \in \ker(3 + \mathbf{A})$, we have that

$$\begin{aligned}\vec{w} &\equiv \vec{u} + \vec{v}, \\ \mathbf{A}\vec{w} &\equiv \vec{u} + 2\vec{v}, \\ \mathbf{A}^2\vec{w} &\equiv \vec{u} + 4\vec{v}, \\ \mathbf{A}^3\vec{w} &\equiv \vec{u} + 8\vec{v} \equiv \vec{u} + 3\vec{v} \pmod{5}, \\ \mathbf{A}^4\vec{w} &\equiv \vec{u} + 6\vec{v} \equiv \vec{u} + \vec{v} \equiv \vec{w} \pmod{5}.\end{aligned}$$

After four iterations, \vec{w} iterates back to itself. Since \vec{w} is an arbitrary vector in our configuration space, we can conclude that *all* vectors in our LCA have a cycle length of at most 4. Furthermore, each iterate of \vec{w} can be computed by using only scalar multiplication on \vec{v} ; matrix multiplication, which is far more computationally expensive, isn't needed!

Thus, for the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, \begin{bmatrix} 3 & 4 \\ 2 & 0 \end{bmatrix})$, the Primary Decomposition Theorem allows us to completely understand how *every* vector in the configuration space iterates by restricting our attention to subspaces defined by the update matrix's minimal polynomial.

Now, we shift our attention to the Minimal Polynomial Theorem. Before that, however, we need to briefly introduce the concept of the “order” of a polynomial.

Definition 2.16. Given a polynomial of the form $x^k q(x)$ over some ring $\mathbb{Z}_N[X]$, where $q(x)$ is a polynomial such that $q(0) \neq 0$, the *order* of the polynomial over the ring is defined to be the smallest natural number c such that $q(x) \mid x^c - 1$. The specific ring over which the order of the polynomial is being found will be evident from context.

At first glance, Definition 2.16 may seem completely esoteric, but there's a clear reason why we care about the order of a polynomial, particularly for minimal annihilating polynomials. Recall the definition of a vector's multiplicative order: for a vector $\vec{v} \in \mathbb{Z}_p^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, it's the smallest positive integer ω such that

$$\mathbf{A}^\omega \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v} \pmod{p}$$

for the vector's transient length τ . Rearranging this expression, we get that

2.5. Theorems Relating to Minimal Polynomials

$$\begin{aligned} \mathbf{A}^\omega \mathbf{A}^\tau \vec{v} - \mathbf{A}^\tau \vec{v} &\equiv \vec{0} \pmod{p} \\ \implies \mathbf{A}^\tau (\mathbf{A}^\omega - \mathbf{I}) \vec{v} &\equiv \vec{0} \pmod{p}. \end{aligned}$$

Thus, $x^\tau(x^\omega - 1)$ is an annihilating polynomial for \vec{v} under \mathbf{A} . Using Definition 2.16, we see that the order of this annihilating polynomial is ω , the multiplicative order of \vec{v} . However, we know that the minimal annihilating polynomial of \vec{v} divides any annihilating polynomial for \vec{v} . So, if $m(x)$ is the minimal annihilating polynomial of \vec{v} , then $m(x) \mid x^\tau(x^\omega - 1)$, and so the order of $m(x)$ is at most ω (by divisibility arguments). In this way, the order of a vector's minimal annihilating polynomial is directly related to the vector's multiplicative order.

In fact, the Minimal Polynomial Theorem establishes an even stronger connection: the order of a vector's minimal annihilating polynomial is *exactly* its multiplicative order.

Theorem 2 (Minimal Polynomial Theorem). *Let $(\mathbb{Z}_p, \mathbb{Z}_p^L, \mathbf{A})$ be an LCA with p prime. If a vector $\vec{v} \in \mathbb{Z}_p^L$ or the matrix \mathbf{A} has the minimal polynomial $\lambda^k m(\lambda)$ for some monic polynomial $m(\lambda)$ where $m(0) \not\equiv 0 \pmod{p}$, then the multiplicative order of the vector/matrix is the order of $m(\lambda)$ and the transient length is k .⁷*

A proof of Theorem 2 is given in “Linear cellular automata” (Patterson [5]). Below, we'll detail a different way of approaching the theorem (for vectors, specifically) which hopefully illuminates *why* such a theorem should be true in the first place.

Consider an arbitrary LCA $(\mathbb{Z}_p, \mathbb{Z}_p^L, \mathbf{A})$ for prime p . Take some arbitrary vector $\vec{v} \in \mathbb{Z}_p^L$ with minimal annihilating polynomial $\lambda^k m(\lambda)$ where k is a nonnegative integer and $m(0) \not\equiv 0 \pmod{p}$. By the Primary Decomposition Theorem, we can break up this vector as $\vec{v} \equiv \vec{s} + \vec{t}$ where $\vec{s} \in \ker(m(\mathbf{A}))$ and $\vec{t} \in \ker(\mathbf{A}^k)$.

Now, every vector in $\ker(m(\mathbf{A}))$ has a minimal annihilating polynomial that divides $m(\lambda)$ (since $\mathbb{Z}_p[X]$ is a principal ideal domain). Because of this, no vector in $\ker(m(\mathbf{A}))$ has a positive transient length (i.e. every vector has a transient length of zero). Why? Well, Proposition 2.1 ensures that, if a vector with a positive transient length exists, then there must be some nonzero vector \vec{w} with the same transient length τ where $\mathbf{A}^\tau \vec{w} \equiv \vec{0} \pmod{p}$. However, such a vector \vec{w} can't possibly exist in $\ker(m(\mathbf{A}))$ since the minimal annihilating polynomial of \vec{w} would have to divide λ^τ while also dividing $m(\lambda)$, which is impossible.

⁷The original statement of Theorem 2 in “Linear cellular automata” (Patterson [5]) was concerned only with the multiplicative order and transient lengths of vectors. However, as the proof relies only on properties of the minimal polynomial itself, the theorem also applies to matrices.

Chapter 2. Background

While vectors in $\ker(m(\mathbf{A}))$ can't have positive transient lengths, vectors in $\ker(\mathbf{A}^k)$ certainly can. In fact, *every* vector in $\ker(\mathbf{A}^k)$ (minus the zero vector) has a transient length between 1 and k , inclusive. Why? By construction, every nonzero vector in $\ker(\mathbf{A}^k)$ is annihilated under multiplication by \mathbf{A}^k , so after at most k iterations by \mathbf{A} , every nonzero vector gets sent to zero. Once a vector iterates to zero, it isn't possible for it to iterate to anything other than the zero vector, so any nonzero iterates of a vector in $\ker(\mathbf{A}^k)$ must be part of its transient region. There can be between 1 and k nonzero vectors (including the vector's initial value) in a vector's transient region under these conditions, meaning the transient length can be between 1 and k , inclusive.

Since all the vectors in $\ker(\mathbf{A}^k)$ eventually iterate to zero, if we're given a vector $\vec{t} \in \ker(\mathbf{A}^k)$ with transient length τ , then the minimal annihilating polynomial of \vec{t} must be λ^τ . The converse is also true, since if λ^τ is the first power of λ to annihilate \vec{t} , then the transient region of \vec{t} has τ elements in it, meaning its transient length is τ . Note that this statement also works in the case where $k = 0$. In this case, $\ker(\mathbf{A}^k) = \{\vec{0}\}$, and so $\vec{t} = \vec{0}$ with transient length $\tau = 0$. Since its transient length is zero, its minimal annihilating polynomial is $\lambda^0 = 1$, and there are exactly zero vectors in its transient region.

Returning to our vector $\vec{v} \equiv \vec{s} + \vec{t}$, we can now say that the transient length of \vec{v} is the transient length of \vec{t} , say τ . To see why, consider the first τ iterates of \vec{v} . They'll look like $\mathbf{A}^i \vec{s} + \mathbf{A}^i \vec{t} \bmod p$, where $\mathbf{A}^i \vec{t}$ will be nonzero. The sum $\mathbf{A}^i \vec{s} + \mathbf{A}^i \vec{t}$ for $0 \leq i < \tau$ will never iterate back to itself since $\mathbf{A}^i \vec{t}$ never iterates back to itself, and by the Primary Decomposition Theorem, vectors have a *unique* representation as a sum of a vector from $\ker(\mathbf{A}^k)$ and from $\ker(m(\mathbf{A}))$ (and since kernels are invariant under multiplication by \mathbf{A} , \vec{s} can never iterate to a vector in $\ker(\mathbf{A}^k)$ to “replace” \vec{t}). Thus, for $0 \leq i < \tau$, $\mathbf{A}^i \vec{s} + \mathbf{A}^i \vec{t}$ must be in the transient region of \vec{v} . Any iterates of \vec{v} after the first τ , however, will *not* be in the transient region of \vec{v} since, for $i \geq \tau$, $\mathbf{A}^i \vec{s} + \mathbf{A}^i \vec{t} \equiv \mathbf{A}^i \vec{s} \bmod p$, and we know \vec{s} has a transient length of 0, so it'll eventually iterate back to itself. Thus, the transient region of \vec{v} will have τ elements in it, and so its transient length is τ .

Now, because λ^k is the *minimal* annihilating polynomial of \vec{t} (since $\lambda^k m(\lambda)$ is the minimal annihilating polynomial for \vec{v} , \vec{t} must take on λ^k as its minimal annihilating polynomial), we can deduce that the transient length of \vec{t} , and equivalently the transient length of \vec{v} , must be k . And, because after k iterations, \vec{v} iterates to an iterate of \vec{s} , the cycle length of \vec{v} is

2.6. Lifting and Embedding

the cycle length of \vec{s} , say ω .

How do we compute the cycle length of \vec{s} ? Well, the cycle length will be given by the smallest positive value of c such that

$$\begin{aligned} \mathbf{A}^c \vec{s} &\equiv \vec{s} \pmod{p} \\ \implies (\mathbf{A}^c - \mathbf{I}) \vec{s} &\equiv \vec{0} \pmod{p}. \end{aligned}$$

The matrix $\mathbf{A}^c - \mathbf{I}$ can only annihilate \vec{s} when it's a multiple of $m(\mathbf{A})$. Thus, ω is given by the smallest positive integer c such that $m(\mathbf{A}) \mid \mathbf{A}^c - \mathbf{I}$. In polynomial terms, we want to find the smallest positive integer c where $m(\lambda) \mid \lambda^c - 1$. This is exactly the definition of the order of $m(\lambda)$. So, the cycle length of \vec{s} , and thus the cycle length of \vec{v} , is the order of $m(\lambda)$.

So, we see that the Minimal Polynomial Theorem makes sense based on how the minimal polynomial of a vector behaves, along with the Primary Decomposition Theorem.

Theorem 2 allows the multiplicative orders and transient lengths of vectors/matrices to be calculated using only properties of the minimal polynomial. Using a computer, this is a much faster way to determine these values as compared to using something like Floyd's Cycle Detection Algorithm.⁸ This theorem also further establishes the connection between the algebraic properties (minimal polynomials, vector spaces, etc.) and the dynamical properties (multiplicative orders, transient lengths, etc.) of LCAs.

It's important to remember that Theorems 1 and 2 require that the modulus of our LCA be prime. If, instead, we're dealing with a more general LCA, say $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$, where N isn't necessarily prime, then the set of polynomials $\mathbb{Z}_N[X]$ isn't necessarily a principal ideal domain, and so the matrix \mathbf{A} doesn't necessarily have a minimal polynomial, nor do vectors necessarily have minimal annihilating polynomials. In Chapter 4, we'll develop some tools for working with annihilating polynomials when the modulus is a prime power.

2.6 Lifting and Embedding

One last concept worth familiarising ourselves with is the concept of lifting/embedding objects between LCAs with prime and prime-power moduli. To do this, let's first take a quick detour to discuss properties of modular reduction.

⁸See "Linear cellular automata" (Patterson [5]) for a description of Floyd's Cycle Detection Algorithm.

Chapter 2. Background

Consider a non-finite linear cellular automata, say $C = (\mathbb{Z}, \mathbb{Z}^L, \mathbf{A})$. Vectors within this automata will iterate in a nearly identical way to vectors under our previous LCAs, except now, we don't reduce by a modulus after multiplying by the update matrix \mathbf{A} . Because of this, C will encapsulate the behaviour of all LCAs of the form $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$ since we can simply take a vector $\vec{v} \in \mathbb{Z}_N^L$ (which can be thought of as a subset of \mathbb{Z}^L , the configuration space of C), iterate it by \mathbf{A} within C , then reduce the resulting iterates modulo N to obtain the respective iterates within the configuration space⁹ \mathbb{Z}_N^L .

Now, say we have some vector $\vec{v} \in \mathbb{Z}^L$ with corresponding iterates \vec{v} , $\mathbf{A}\vec{v}$, $\mathbf{A}^2\vec{v}$, etc., within the LCA C . If we want to find the iterates of \vec{v} within the LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ for some prime power p^k , we can start by expanding out the iterates from C into “base- p representation”. Just as how we can expand numbers like 123 into a base-10 representation

$$123 = (10^0 \times 3) + (10^1 \times 2) + (10^2 \times 1),$$

we can expand vectors into this same form, just replacing the powers of 10 with powers of p . In base- p representation, our iterates may look something like

$$\begin{aligned} \vec{v} &\equiv \vec{u}_{0,0} + p\vec{u}_{0,1} + p^2\vec{u}_{0,2} + \cdots + p^k\vec{u}_{0,k} + \cdots \\ \mathbf{A}\vec{v} &\equiv \vec{u}_{1,0} + p\vec{u}_{1,1} + p^2\vec{u}_{1,2} + \cdots + p^k\vec{u}_{1,k} + \cdots \\ \mathbf{A}^2\vec{v} &\equiv \vec{u}_{2,0} + p\vec{u}_{2,1} + p^2\vec{u}_{2,2} + \cdots + p^k\vec{u}_{2,k} + \cdots \\ \vdots &\equiv \vdots \end{aligned}$$

for vectors $\vec{u}_{i,j} \in \mathbb{Z}_p^L$.

Using base- p representation, reducing modulo p^k becomes simple: we just remove any vectors with coefficients greater than or equal to p^k from each sum. This works because we know the sum of $\vec{u}_{i,0}$ to $p^{k-1}\vec{u}_{i,k-1}$ for each $\mathbf{A}^i\vec{v}$ will remain the same under modular reduction modulo p^k —the resulting vector sum can never contain any components greater than or equal to p^k . The terms $p^k\vec{u}_{i,k}$ and above, however, will *always* contain components which are multiples of p^k , and thus they'll vanish under the modular reduction. This results

⁹More formally, we have a commutative diagram between \mathbb{Z}^L and \mathbb{Z}_N^L where the mappings $\pi : \mathbb{Z}^L \rightarrow \mathbb{Z}_N^L$, $\pi(x) = x \bmod N$ and $i(\vec{x}) = \mathbf{A}\vec{x}$ can be applied in either order to obtain the result of iteration within \mathbb{Z}_N^L .

2.6. Lifting and Embedding

in the following iterates:

$$\begin{aligned}\vec{v} &\equiv \vec{u}_{0,0} + p\vec{u}_{0,1} + p^2\vec{u}_{0,2} + \cdots + p^{k-1}\vec{u}_{0,k-1} \pmod{p^k} \\ \mathbf{A}\vec{v} &\equiv \vec{u}_{1,0} + p\vec{u}_{1,1} + p^2\vec{u}_{1,2} + \cdots + p^{k-1}\vec{u}_{1,k-1} \pmod{p^k} \\ \mathbf{A}^2\vec{v} &\equiv \vec{u}_{2,0} + p\vec{u}_{2,1} + p^2\vec{u}_{2,2} + \cdots + p^{k-1}\vec{u}_{2,k-1} \pmod{p^k} \\ \vdots &\equiv \vdots\end{aligned}$$

Notice that the two sets of iterates are very similar: the only difference with the iterates modulo p^k is that less terms are included in the sum for each iterate. Otherwise, they're identical. This will apply to *any* LCA whose modulus is a power of p and whose update matrix is \mathbf{A} . For instance, if we wanted to calculate the iterates of \vec{v} for the LCA $(\mathbb{Z}_{p^2}, \mathbb{Z}_{p^2}^L, \mathbf{A})$, we simply take our sum for each iterate $\mathbf{A}^i\vec{v}$ and include only the vectors $\vec{u}_{i,0}$ and $p\vec{u}_{i,1}$:

$$\begin{aligned}\vec{v} &\equiv \vec{u}_{0,0} + p\vec{u}_{0,1} \pmod{p^2} \\ \mathbf{A}\vec{v} &\equiv \vec{u}_{1,0} + p\vec{u}_{1,1} \pmod{p^2} \\ \mathbf{A}^2\vec{v} &\equiv \vec{u}_{2,0} + p\vec{u}_{2,1} \pmod{p^2} \\ \vdots &\equiv \vdots\end{aligned}$$

These iterates can be obtained by reducing the iterates of \vec{v} within the non-finite automata C , but they can also be obtained by reducing the iterates of \vec{v} within the LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ (so long as $k \geq 2$). This illustrates an intimate connection between LCAs whose update matrices are the same¹⁰: the iterates of vectors in LCAs with high prime-power moduli will always map down to iterates of vectors in LCAs with a lower prime-power moduli.

This property is useful enough to warrant its own terminology.

Definition 2.17. Given a vector $\vec{v} \in \mathbb{Z}_{p^k}^L$ for a prime power p^k (with $k \geq 1$), a *lift* of \vec{v} modulo $p^{k+\ell}$ is any vector of the form

$$\vec{v} + p^k\vec{w} \pmod{p^{k+\ell}}, \quad \vec{w} \in \mathbb{Z}_{p^\ell}^L.$$

¹⁰In fact, the update matrices only need to be congruent modulo the lesser of the two moduli, as that will ensure the relevant terms in each iterate's sum are identical between the two different LCAs.

Chapter 2. Background

In essence, a lift of a vector modulo a prime or prime power is any vector modulo a higher prime power that maps down to it via modular reduction. For our purposes, lifts give us a way to deduce properties of vectors within LCAs of different prime power moduli without having to explicitly calculate any of their iterates.

Example 2.1. Consider the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, [\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}])$ and the vector $[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}] \in \mathbb{Z}_5^2$. This particular vector has a cycle length of 4. Now, consider the similar LCA $(\mathbb{Z}_{25}, \mathbb{Z}_{25}^2, [\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}])$ and any vector of the form $([\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}] + 5\vec{w}) \in \mathbb{Z}_{25}^2$. These vectors are lifts of the vector $[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}]$ modulo 25, and so we know their iterates must map down to the iterates of $[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}]$ modulo 5. Immediately, then, we know that the cycle length of $[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}] + 5\vec{w}$ under iteration modulo 25 must be a multiple of 4, the cycle length under iteration modulo 5. Otherwise, the lift's iterates wouldn't map to the iterates modulo 5. So, just by determining the cycle length of a single vector modulo 5, lifts allow us to deduce something about the cycle lengths of an entire family of vectors modulo 25. \diamond

It turns out that considering lifts for matrices and polynomials is also useful under certain circumstances. They are defined similarly.

Definition 2.18. Given a matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$ for a prime power p^k (with $k \geq 1$), a *lift* of \mathbf{A} modulo $p^{k+\ell}$ is any matrix of the form

$$\mathbf{A} + p^k \mathbf{B} \pmod{p^{k+\ell}}, \quad \mathbf{B} \in \mathbb{Z}_{p^\ell}^{L \times L}.$$

Definition 2.19. Given some polynomial $r(x) \in \mathbb{Z}_{p^k}[X]$ for some prime power p^k (with $k \geq 1$), a *lift* of $r(x)$ modulo $p^{k+\ell}$ is any polynomial of the form

$$r(x) + p^k q(x) \pmod{p^{k+\ell}}, \quad q(x) \in \mathbb{Z}_{p^\ell}[X].$$

In Example 2.1, we demonstrated how calculating information about a vector under some prime or prime power modulus can allow us to deduce information about its lifts modulo higher powers of the prime. Using a concept very similar to lifts, we actually don't need to consider the lower power modulus; all of our computation can be done under a single modulus, within a single LCA!

2.6. Lifting and Embedding

Consider again the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix})$ and the vector $\begin{bmatrix} 1 \\ 2 \end{bmatrix} \in \mathbb{Z}_5^2$. The iterates for this vector within this LCA are

$$\begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \end{bmatrix}, \begin{bmatrix} 4 \\ 3 \end{bmatrix}, \begin{bmatrix} 2 \\ 4 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \dots$$

and so on. If we were then to consider the LCA $(\mathbb{Z}_{25}, \mathbb{Z}_{25}^2, \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix})$ and the vector $\begin{bmatrix} 1 \\ 2 \end{bmatrix} \in \mathbb{Z}_{25}^2$, we'd get a similar list of iterates within the higher-power LCA. Splitting up the iterates into their base- p representations, we get

$$\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix} + 5 \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right), \left(\begin{bmatrix} 3 \\ 1 \end{bmatrix} + 5 \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right), \left(\begin{bmatrix} 4 \\ 3 \end{bmatrix} + 5 \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right), \left(\begin{bmatrix} 2 \\ 4 \end{bmatrix} + 5 \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right), \left(\begin{bmatrix} 1 \\ 2 \end{bmatrix} + 5 \begin{bmatrix} 2 \\ 1 \end{bmatrix} \right), \dots$$

and so on. The “ p^0 ” terms of the iterates modulo 25 match exactly with the iterates modulo 5. This shouldn't be too surprising since $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ is a lift of $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ modulo 25 (we may call it the *trivial lift* of the vector modulo 25), and so we know that its iterates must map down. The only way for its iterates to map down is if this “ p^0 ” term matches the iterates modulo 5. Perhaps more surprising is the fact that we can use the linearity of our LCA to force the iterates modulo 25 to be even more similar to the iterates modulo 5.

Recall that, because matrix multiplication is linear, we have that

$$\mathbf{A}(p\vec{v}) \equiv p(\mathbf{A}\vec{v}) \pmod{p^k}.$$

Therefore, if we were to iterate the vector $5 \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ modulo 25, the iterates we obtain would be the same as those for $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$, just scaled by a factor of 5. Looking at the iterates for $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ above, we see that scaling by 5 would cause the “ p^1 ” terms to vanish (since they'd have a factor of 25 after the scaling) and the “ p^0 ” terms to remain, just scaled by 5. The resulting iterates look like

$$5 \begin{bmatrix} 1 \\ 2 \end{bmatrix}, 5 \begin{bmatrix} 3 \\ 1 \end{bmatrix}, 5 \begin{bmatrix} 4 \\ 3 \end{bmatrix}, 5 \begin{bmatrix} 2 \\ 4 \end{bmatrix}, 5 \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \dots$$

and so on. Thus, if we ignore the factor of 5, the iterates of $5 \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ under iteration modulo 25 match exactly with the iterates of $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ under iteration modulo 5.

In fact, we can see that this sort of behaviour applies generally by considering a particular

Chapter 2. Background

mapping between vectors under different moduli.

Define $\phi : \mathbb{Z}_{p^k}^L \rightarrow p\mathbb{Z}_{p^{k+1}}^L$ for prime p to be the mapping $\phi(\vec{v}) \equiv p\vec{v}$. This creates a bijective mapping that preserves matrix multiplication and vector addition between $\mathbb{Z}_{p^k}^L$ and $p\mathbb{Z}_{p^{k+1}}^L$. To see this, we first note that ϕ does, in fact, preserve vector addition. For arbitrary vectors $\vec{v}, \vec{w} \in \mathbb{Z}_{p^k}^L$,

$$\begin{aligned} & \phi(\vec{v} + \vec{w}) \\ & \equiv p(\vec{v} + \vec{w}) \\ & \equiv p\vec{v} + p\vec{w} \\ & \equiv \phi(\vec{v}) + \phi(\vec{w}). \end{aligned}$$

We can then show that ϕ is injective. Assuming $\phi(\vec{v}) \equiv \phi(\vec{w})$, we see that

$$\begin{aligned} & \phi(\vec{v}) \equiv \phi(\vec{w}) \quad \text{mod } p^{k+1} \\ \implies & p\vec{v} \equiv p\vec{w} \quad \text{mod } p^{k+1} \\ \implies & p(\vec{v} - \vec{w}) \equiv \vec{0} \quad \text{mod } p^{k+1} \\ \implies & \vec{v} - \vec{w} \equiv \vec{0} \quad \text{mod } p^k \\ \implies & \vec{v} \equiv \vec{w} \quad \text{mod } p^k, \end{aligned}$$

We also note that ϕ is surjective since, for any $p\vec{w} \in p\mathbb{Z}_{p^{k+1}}^L$, there exists a vector $\vec{w} \in \mathbb{Z}_{p^k}^L$ such that $\phi(\vec{w}) \equiv p\vec{w}$. Thus, ϕ is bijective.

Finally, we see that ϕ preserves matrix multiplication. For some vector $\vec{v} \in \mathbb{Z}_{p^k}^L$ and some matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$,

$$\phi(\mathbf{A}\vec{v}) \equiv p(\mathbf{A}\vec{v}) \equiv \mathbf{A}(p\vec{v}) \equiv \mathbf{A}\phi(\vec{v}) \quad \text{mod } p^{k+1},$$

and so matrix multiplication is preserved.

What does ϕ let us say about LCAs of differing prime-power moduli? Because ϕ is a bijection, it shows that the behaviour of vectors within LCAs of prime or prime-power moduli can be related to a subset of the vectors within LCAs of higher prime-power moduli. In particular, given two LCAs $A = (\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ and $B = (\mathbb{Z}_{p^{k+\ell}}, \mathbb{Z}_{p^{k+\ell}}^L, \mathbf{A})$ for a prime power p^k (with $k \geq 1$), the iterates of any vector $\vec{v} \in \mathbb{Z}_{p^k}^L$ within A can be matched with the iterates of $p^\ell \vec{v} \in \mathbb{Z}_{p^{k+\ell}}^L$ within B , notwithstanding the additional factor of p^ℓ present. So, if we wanted to compute the iterates of \vec{v} modulo p^k , we can instead compute the iterates of

2.6. Lifting and Embedding

$p^\ell \vec{v}$ modulo $p^{k+\ell}$ as they'll give us the same information. This means we can work under a single prime-power modulus to determine properties of all the LCAs with lower prime-power moduli!

As giving names to quantities has proven useful up to now, we'll denote the vector $p^\ell \vec{v}$ as the *embed* vector of \vec{v} .

Definition 2.20. Given a vector $\vec{v} \in \mathbb{Z}_{p^k}^L$ for some prime power p^k (with $k \geq 1$), an *embed* vector of \vec{v} is any vector of the form $p^\ell \vec{v} \in \mathbb{Z}_{p^{k+\ell}}^L$. Via the bijection ϕ , a vector and its embed can have their iterates matched up within their respective LCAs, so long as the update matrices are equivalent.

Lifts and embeds are conceptually similar, but they behave very differently. The following example demonstrates some of their differences.

Example 2.2. Take the LCAs $A = (\mathbb{Z}_5, \mathbb{Z}_5^2, \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix})$ and $B = (\mathbb{Z}_{25}, \mathbb{Z}_{25}^2, \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix})$ from earlier. Given the vector $\begin{bmatrix} 1 \\ 2 \end{bmatrix} \in \mathbb{Z}_5^2$, its one and only embed vector in B is $5\begin{bmatrix} 1 \\ 2 \end{bmatrix} \equiv \begin{bmatrix} 5 \\ 10 \end{bmatrix} \in \mathbb{Z}_{25}^2$. However, $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ has multiple possible lift vectors in B —25, to be exact. Let's consider the specific lift $\begin{bmatrix} 1 \\ 2 \end{bmatrix} + 5\begin{bmatrix} 2 \\ 3 \end{bmatrix} \equiv \begin{bmatrix} 11 \\ 17 \end{bmatrix} \in \mathbb{Z}_{25}^2$.

Knowing that the cycle length of $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ within A is 4, what do we know about our embed and lift vectors? The iterates of embed vectors will match the iterates of their corresponding “unembedded” vector via ϕ , so we know that the cycle length of $\begin{bmatrix} 5 \\ 10 \end{bmatrix}$ within B will also be 4. For lifts, we know that their iterates must map down to their corresponding “non-lifted” vector, so the cycle length of $\begin{bmatrix} 11 \\ 17 \end{bmatrix}$ within B must be a multiple of 4.

What happens if we reduce our lift and embed vectors modulo 5? Lift vectors map down to their “non-lifted” vectors, so the iterates of $\begin{bmatrix} 11 \\ 17 \end{bmatrix}$ will map down to the iterates of $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ within A . Embed vectors, on the other hand, have no such requirement. In fact, because the embed vector has a factor of 5 attached to it, $\begin{bmatrix} 5 \\ 10 \end{bmatrix}$ reduces to the zero vector modulo 5, giving us no useful information about the iterates of its “unembedded” vector. \diamond

As with lifts, it sometimes proves useful to use embeds of matrices and polynomials too.

Definition 2.21. Given a matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$ for a prime power p^k (with $k \geq 1$), an *embed* matrix of \mathbf{A} is any matrix of the form $p^\ell \mathbf{A} \in \mathbb{Z}_{p^{k+\ell}}^{L \times L}$. Via ϕ , a matrix and its embed will iterate equivalently under their respective moduli.

Chapter 2. Background

Definition 2.22. Given a polynomial $r(x) \in \mathbb{Z}_{p^k}[X]$ for a prime power p^k (with $k \geq 1$), an *embed* of $r(x)$ is any polynomial of the form $p^\ell r(x) \in \mathbb{Z}_{p^{k+\ell}}[X]$. Via ϕ , a polynomial and its embed will behave similarly under their respective moduli.

Throughout this thesis, lifts and embeds may be used without explicitly labelling them as lifts and embeds. This is simply due to how frequently they appear within proofs and algorithms involving LCAs of differing prime and prime power moduli.

With all the relevant terms now defined, we can begin to explore some more interesting properties of LCAs in the following chapters.

Chapter 3

Understanding the Behaviour of Linear Cellular Automata

3.1 Multiplicative Orders

One of the things we're interested in determining about an arbitrary LCA $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$ is the multiplicative order and transient length of its update matrix, as these values provide upper bounds for the corresponding values of all the vectors within the LCA. To see this, let τ be the transient length of \mathbf{A} and let ω be the multiplicative order of \mathbf{A} . We have that

$$\mathbf{A}^\omega \mathbf{A}^\tau \equiv \mathbf{A}^\tau \pmod{N},$$

and so for any vector $\vec{v} \in \mathbb{Z}_N^L$,

$$\mathbf{A}^\omega \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v} \pmod{N}.$$

Since a vector's transient length is defined to be the *smallest* value x where $\mathbf{A}^c \mathbf{A}^x \vec{v} \equiv \mathbf{A}^x \vec{v} \pmod{N}$ has positive solutions for c , the transient length of \vec{v} can be no higher than τ as τ satisfies this relation. Similarly, because a vector's multiplicative order is defined to be the *smallest* value c where $\mathbf{A}^c \mathbf{A}^t \vec{v} \equiv \mathbf{A}^t \vec{v} \pmod{N}$ (for t being greater than or equal to the transient length of \vec{v}), the cycle length of \vec{v} can be no higher than ω since ω satisfies this relation.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

In fact, something stronger is true for multiplicative orders. On top of being bounded above by an update matrix's multiplicative order, a vector's multiplicative order under said matrix must *divide* the matrix's multiplicative order. To see this, let \vec{v} be an arbitrary vector in an LCA with update matrix \mathbf{A} . Let c be the vector's multiplicative order. Further, let ω be the multiplicative order of \mathbf{A} and let τ be the transient length of \mathbf{A} . Now, we have that

$$\mathbf{A}^c \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v}$$

since the transient length of \vec{v} is bounded above by τ . However, we also have that

$$\mathbf{A}^\omega \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v}$$

by definition of the multiplicative order of \mathbf{A} . Then, repeatedly using the definition of a vector's multiplicative order, we have that

$$\mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^c \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^{2c} \mathbf{A}^\tau \vec{v} \equiv \dots \equiv \mathbf{A}^\omega \mathbf{A}^\tau \vec{v}.$$

Now, assume that ω is not a multiple of c . Let nc be the largest multiple of c which is less than ω . Then, this gives us that

$$\mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^{nc} \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\omega \mathbf{A}^\tau \vec{v}.$$

Notice that $\omega - nc$ must be a positive integer less than c by construction. This means we can say

$$\mathbf{A}^{\omega - nc} (\mathbf{A}^\tau \vec{v}) \equiv \mathbf{A}^{\omega - nc} (\mathbf{A}^{nc} \mathbf{A}^\tau \vec{v}) \equiv \mathbf{A}^\omega \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v}.$$

Condensing this congruence, we see

$$\mathbf{A}^{\omega - nc} \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v},$$

and so the multiplicative order of \vec{v} must be bounded above by $\omega - nc$. However, this is a contradiction since $\omega - nc < c$. Thus, our assumption that ω isn't a multiple of c must be false. Hence, $c \mid \omega$.

3.1. Multiplicative Orders

Thus, knowing the multiplicative order of an LCA's update matrix immediately restricts the possible multiplicative orders for all vectors within the configuration space, and so being able to determine the multiplicative order is something we'd like to do. We'll spend this section of the chapter establishing some tools we can use to work with and understand multiplicative orders.

The multiplicative order of a matrix is typically hard to determine without either explicitly calculating it via repeated matrix multiplication or by comparing it to another, similar LCA (such as an LCA which uses a lift of our matrix). However, there are other values we can find for the matrix that make this computation easier.

Proposition 3.1. Let p^k be an odd prime power with $k > 1$, and $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$ an invertible matrix where $\mathbf{A}^w \equiv \mathbf{I} + p^{k-1}\mathbf{B}$ for some positive integer w and some matrix $\mathbf{B} \in \mathbb{Z}_p^{L \times L}$. Then

$$\mathbf{A}^{nw} \equiv \mathbf{I} + np^{k-1}\mathbf{B} \pmod{p^k}$$

for any positive integer n .

Proof. We'll prove this statement using induction.

Base case: $n = 1$. If $n = 1$, we have that

$$\mathbf{A}^{nw} \equiv \mathbf{A}^w \equiv \mathbf{I} + p^{k-1}\mathbf{B} \pmod{p^k}.$$

This matches the form $\mathbf{A}^{nw} \equiv \mathbf{I} + np^{k-1}\mathbf{B}$, which is what we want. Thus, our statement is true when $n = 1$.

Induction step. Assume we have some positive integer r where

$$\mathbf{A}^{rw} \equiv \mathbf{I} + rp^{k-1}\mathbf{B} \pmod{p^k}.$$

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Expanding $\mathbf{A}^{(r+1)w}$, we see that

$$\begin{aligned}\mathbf{A}^{(r+1)w} &\equiv (\mathbf{I} + p^{k-1}\mathbf{B})(\mathbf{I} + rp^{k-1}\mathbf{B}) \\ &\equiv \mathbf{I} + p^{k-1}\mathbf{B} + rp^{k-1}\mathbf{B} + rp^{2k-2}\mathbf{B} \\ &\equiv \mathbf{I} + (r+1)p^{k-1}\mathbf{B} \pmod{p^k}.\end{aligned}$$

So

$$\mathbf{A}^{rw} \equiv \mathbf{I} + rp^{k-1}\mathbf{B} \implies \mathbf{A}^{(r+1)w} \equiv \mathbf{I} + (r+1)p^{k-1}\mathbf{B} \pmod{p^k}.$$

By induction, we can say that

$$\mathbf{A}^{nw} \equiv \mathbf{I} + np^{k-1}\mathbf{B} \pmod{p^k},$$

which is what we wanted to show.

QED

Proposition 3.1 gives us a feel for what certain powers of an invertible matrix will look like. In fact, once we find a value w that satisfies the assumptions of the proposition, we see that for pw , p being the prime base of the modulus used, $\mathbf{A}^{pw} \equiv \mathbf{I} \pmod{p^k}$, and so the multiplicative order of \mathbf{A} must divide pw (since the multiplicative order of \mathbf{A} is the smallest value of c where \mathbf{A}^c becomes the identity matrix for invertible matrices \mathbf{A}).

A more interesting result regarding the multiplicative orders of update matrices (in the prime modulus case) relates to their minimal polynomials. As the Minimal Polynomial Theorem (Theorem 2) reveals, there's a connection between the minimal polynomial (specifically, its order) and an update matrix's multiplicative order. It turns out that the multiplicity of the minimal polynomial's roots also relates to the matrix's multiplicative order. To show this relation, we must first establish a few facts about how minimal polynomials behave under *extension fields*.

For our purposes, an extension field is simply a field with an added element that solves a particular polynomial equation that didn't have a solution over the original field. An example of such an extension field is the complex numbers, which is an extension field of the real numbers with $\pm i$ added as a solution to the polynomial $x^2 + 1 = 0$. We may also make

3.1. Multiplicative Orders

use of the *splitting field* of a particular polynomial, which is the smallest extension field of a field where the given polynomial can be factored into linear terms.

An extension field of \mathbb{Z}_p is certainly more complicated than \mathbb{Z}_p itself, but many of the things we care about with regards to LCAs stay the same over extension fields, as the following propositions will show.

Proposition 3.2. Let p be an odd prime and \mathbb{K} be a finite algebraic extension field of \mathbb{Z}_p . Then, the multiplicative order of any element in \mathbb{K}^* —the set of invertible elements in \mathbb{K} —divides $p^d - 1$, where d is the degree of the extension.

Proof. The finite abelian group¹ \mathbb{K}^* under the operation of multiplication modulo p has $p^d - 1$ elements, and thus its order is $p^d - 1$. By Lagrange’s Theorem², any element in \mathbb{K}^* must have a multiplicative order dividing the group’s order. QED

Proposition 3.3. Let p be an odd prime. If $V \subseteq \mathbb{Z}_p^L$ is a set of linearly-independent vectors modulo p , then V is also linearly independent over an algebraic extension field \mathbb{K} of \mathbb{Z}_p .

Proof. Let $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$ where $n \leq L$. Consider the matrix

$$\mathbf{A} = \begin{bmatrix} \vec{v}_1 & \vec{v}_2 & \cdots & \vec{v}_n \end{bmatrix}.$$

Since V is linearly independent over \mathbb{Z}_p , we know that the equation $\mathbf{A}\vec{b} \equiv \vec{0} \pmod{p}$ has only the trivial solution $\vec{b} \equiv \vec{0}$. The algorithm for determining this—Gauss-Jordan Elimination—behaves the same over both \mathbb{Z}_p and \mathbb{K} since the entries in \mathbf{A} are all in \mathbb{Z}_p , and thus no operation will ever require elements from $\mathbb{K} \setminus \mathbb{Z}_p$. This means, over \mathbb{K} , the equation $\mathbf{A}\vec{b} \equiv \vec{0}$ also only has the trivial solution $\vec{b} \equiv \vec{0}$, which implies V is linearly independent over \mathbb{K} . QED

Proposition 3.4. Let p be an odd prime, $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, and \mathbb{K} an algebraic extension field of \mathbb{Z}_p . Then the minimal polynomial of \mathbf{A} over \mathbb{K} is the same as over \mathbb{Z}_p .

¹For us, a finite abelian group can be thought of as a finite ring with only the “addition” operation defined.

²For our purposes, Lagrange’s Theorem says that, given a finite abelian group, the order of any element in the group divides the order of the group itself.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Proof. The polynomial ring $\mathbb{K}[X]$ is a principal ideal domain, and so every ideal is generated by a single element. The set of annihilating polynomials for \mathbf{A} is indeed an ideal since any polynomial linear combination of annihilating polynomials is also an annihilating polynomial. Thus, the minimal polynomial of \mathbf{A} over \mathbb{K} , denoted as $m_{\mathbb{K}}(x)$, is the sole generator of the ideal of annihilating polynomials for \mathbf{A} over \mathbb{K} , meaning it must divide all other annihilating polynomials for \mathbf{A} .

We see that $\mathbb{Z}_p[X] \subseteq \mathbb{K}[X]$, and so the minimal polynomial of \mathbf{A} over \mathbb{Z}_p , denoted $m_p(x)$, must be in the ideal generated by $m_{\mathbb{K}}(x)$. Therefore, $m_{\mathbb{K}}(x) \mid m_p(x)$.

Now, let d denote the degree of $m_{\mathbb{K}}(x)$. Since $m_{\mathbb{K}}(\mathbf{A}) \equiv \mathbf{0}$, we know that the set of matrices $\{\mathbf{I}, \mathbf{A}, \mathbf{A}^2, \dots, \mathbf{A}^d\}$ is linearly dependent over \mathbb{K} . Using the contrapositive of Proposition 3.3 (which we can do since $L \times L$ matrices can be treated as $L^2 \times 1$ vectors), this means the set $\{\mathbf{I}, \mathbf{A}, \mathbf{A}^2, \dots, \mathbf{A}^d\}$ is also linearly dependent over \mathbb{Z}_p , which means a polynomial of degree d exists over \mathbb{Z}_p which annihilates \mathbf{A} . This implies that $m_p(x)$ has a degree which is no higher than d , since $m_p(x)$ is defined to be the monic polynomial of *least degree* which annihilates \mathbf{A} .

As well, we can see that the degree of $m_p(x)$ must be no lower than d since, if it was, then $m_{\mathbb{K}}(x)$ wouldn't be the monic polynomial of least degree over \mathbb{K} which annihilates \mathbf{A} since $m_p(x)$ is also an annihilating polynomial for \mathbf{A} over \mathbb{K} . Therefore, the degree of $m_p(x)$ is exactly d . Since we know $m_{\mathbb{K}}(x) \mid m_p(x)$, and since both polynomials are monic, we conclude that $m_{\mathbb{K}}(x) = m_p(x)$. **QED**

With these few ideas about extension fields, we can now proceed with proving the connection between an update matrix's multiplicative order and the multiplicity of the roots of its minimal polynomial.

Proposition 3.5. Let p be an odd prime. Let $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$ be an invertible matrix modulo p . Then the multiplicative order of \mathbf{A} is a multiple of p if and only if the minimal polynomial of \mathbf{A} has a multiple root.

Proof. Let the multiplicative order of \mathbf{A} be denoted as ω . Also, let the minimal polynomial of \mathbf{A} be represented as $m(x)$, and let \mathbb{K} be the splitting field of $m(x)$ over \mathbb{Z}_p . By Proposition 3.4, the minimal polynomial of \mathbf{A} remains the same over any algebraic extension field of \mathbb{Z}_p . Thus, we can work over the field \mathbb{K} to deduce properties of $m(x)$.

3.1. Multiplicative Orders

First, we'll show that $m(x)$ having a multiple root implies that ω is a multiple of p .

Assume that $m(x)$ has a multiple root, meaning we can write the minimal polynomial over \mathbb{K} as $m(x) = (x - c)^2 \mu(x)$, where $c \in \mathbb{K}$ and $\mu(x) \in \mathbb{K}[X]$.

Now, note that, since ω is the multiplicative order of \mathbf{A} , and since \mathbf{A} is invertible, we have that

$$\begin{aligned} \mathbf{A}^\omega &\equiv \mathbf{I} \pmod{p} \\ \implies \mathbf{A}^\omega - \mathbf{I} &\equiv \mathbf{0} \pmod{p}. \end{aligned}$$

This shows $x^\omega - 1$ is an annihilating polynomial for \mathbf{A} over \mathbb{Z}_p . All annihilating polynomials for \mathbf{A} are polynomial multiples of $m(x)$ (since $\mathbb{K}[X]$ is a principal ideal domain), so we have that

$$\begin{aligned} m(x) &\mid x^\omega - 1 \\ \implies (x - c)^2 &\mid x^\omega - 1 \\ \implies (x - c) &\mid x^\omega - 1. \end{aligned}$$

Dividing $(x^\omega - 1)$ by $(x - c)$ (via polynomial division), we see that

$$x^\omega - 1 = (x - c) \left(\sum_{i=0}^{\omega-1} c^i x^{\omega-1-i} \right) + c^\omega - 1. \quad (3.1)$$

We know the remainder of this division (i.e. $c^\omega - 1$) should be zero since $(x - c) \mid x^\omega - 1$. Therefore,

$$c^\omega \equiv 1 \pmod{p}.$$

For this congruence to hold, ω must be a multiple of the multiplicative order of c . Since c is an eigenvalue for \mathbf{A} over \mathbb{K} , this will always hold true.

We also know that $(x - c)^2 \mid (x^\omega - 1)$, so $(x - c)$ should divide the quotient obtained in Equation (3.1). Using polynomial division, we get the following:

$$\sum_{i=0}^{\omega-1} c^i x^{\omega-1-i} = (x - c) \left(\sum_{i=0}^{\omega-2} (i+1) c^i x^{\omega-2-i} \right) + c^{\omega-1} + (\omega - 1) c^{\omega-1}.$$

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Again, we know the remainder of the division (i.e. $c^{\omega-1} + (\omega - 1)c^{\omega-1}$) is zero since, from above, $(x - c)^2 \mid x^\omega - 1$. Therefore,

$$(1 - \omega)c^{\omega-1} \equiv c^{\omega-1} \pmod{p}.$$

The only way this congruence can be satisfied is if $(1 - \omega) \equiv 1 \pmod{p}$, which implies that $\omega \equiv 0 \pmod{p}$. So, ω is a multiple of p .

Now, we'll show that ω being a multiple of p implies $m(x)$ has a multiple root.

To show this, assume otherwise. That is, assume that $p \mid \omega$, but $m(x)$ does *not* have a multiple root. Thus, over the splitting field \mathbb{K} , $m(x)$ can be written as

$$m(x) = \prod_{i=0}^{\ell} (x - c_i), \quad c_i \in \mathbb{K},$$

for some positive integer ℓ , where for all $0 \leq i \leq \ell$ and $0 \leq j \leq \ell$, $i \neq j \implies c_i \not\equiv c_j \pmod{p}$. Then, by the Primary Decomposition Theorem (Theorem 1), this means the vector space \mathbb{K}^L can be decomposed as

$$\mathbb{K}^L = \bigoplus_{i=0}^{\ell} \ker(\mathbf{A} - c_i \mathbf{I})$$

where each kernel contains at least one nonzero vector. Due to the fact that each factor in $m(x)$ is linear, each kernel corresponds to an eigenspace with an eigenvalue given by one of c_i . Let $E_i = \ker(\mathbf{A} - c_i \mathbf{I})$.

Now, by proposition 3 in “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]), there exists at least one nonzero vector from each of E_i we can sum together to obtain a maximal vector, a vector whose multiplicative order equals ω . From this same proposition, we know $\omega = \text{lcm}(\omega_1, \omega_2, \dots, \omega_\ell)$ where ω_i is the maximum multiplicative order possible for a vector in E_i .

Since each E_i is an eigenspace, each ω_i is simply the multiplicative order of the eigenvalue of E_i . If the eigenvalue of E_i is in \mathbb{Z}_p , then by Fermat's Little Theorem³ we know the multiplicative order of the eigenvalue must divide $p - 1$, so ω_i must also divide $p - 1$. Otherwise, if the eigenvalue of E_i is in $\mathbb{K} \setminus \mathbb{Z}_p$, then Proposition 3.2 tells us that the multiplicative order

³Fermat's Little Theorem states that, for any integer n coprime to a prime p , we have that $n^{p-1} \equiv 1 \pmod{p}$.

3.1. Multiplicative Orders

of the eigenvalue must divide $p^d - 1$, where d is the degree of the particular extension where the eigenvalue is obtained from. So, ω_i must also divide $p^d - 1$.

Importantly, we see that p doesn't divide $p - 1$ nor $p^d - 1$. This means that $\omega = \text{lcm}(\omega_1, \omega_2, \dots, \omega_\ell)$ cannot be a multiple of p . This is a contradiction since $p \mid \omega$. Thus, our assumption that $m(x)$ doesn't contain a multiple root must be incorrect. **QED**

As an example of Proposition 3.5, consider the invertible matrix

$$\begin{bmatrix} 1 & 4 & 4 & 0 & 1 \\ 4 & 2 & 0 & 3 & 4 \\ 1 & 3 & 1 & 3 & 1 \\ 1 & 3 & 1 & 2 & 0 \\ 4 & 2 & 4 & 3 & 1 \end{bmatrix} \pmod{5}.$$

The minimal polynomial of this matrix is

$$1 + 2\lambda + 2\lambda^2 + \lambda^3 + 3\lambda^4 + \lambda^5 = (4 + \lambda)(3 + 2\lambda + \lambda^2)(3 + 2\lambda + \lambda^2).$$

We see that it has a repeated factor (and thus a multiple root), and so by Proposition 3.5, its multiplicative order should be a multiple of 5. Sure enough, its multiplicative order is $120 = 5(24)$.

As another example, consider the invertible matrix

$$\begin{bmatrix} 3 & 3 & 4 & 1 & 3 \\ 0 & 0 & 2 & 4 & 0 \\ 3 & 3 & 1 & 2 & 3 \\ 0 & 4 & 3 & 2 & 3 \\ 0 & 2 & 2 & 2 & 3 \end{bmatrix} \pmod{5}.$$

The minimal polynomial of this matrix is

$$3 + 3\lambda + 3\lambda^2 + 2\lambda^3 + \lambda^4 + \lambda^5 = (3 + \lambda)(1 + 4\lambda + 3\lambda^2 + 3\lambda^3 + \lambda^4).$$

This polynomial has no multiple roots, and so Proposition 3.5 tells us that its multiplicative

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

order shouldn't be a multiple of 5. Calculating the multiplicative order, we get 156, which is not a multiple of 5.

Proposition 3.5 is quite useful in deducing information about an LCA update matrix's multiplicative order without having to actually calculate it (which can take a lot of time and/or computational resources; computing the minimal polynomial is much faster).

3.2 Cycle Converting Matrices

Another object we can study to gain info on multiplicative orders is the *cycle converting matrix*, or CCM.

Definition 3.1. Let \mathbf{A} be an invertible matrix with multiplicative order ω . A *cycle converting matrix* for \mathbf{A} is any matrix of the form

$$\sum_{i=0}^{\frac{\beta-\alpha}{\alpha}} \mathbf{A}^{i\alpha},$$

where $\beta \mid \omega$ and $\alpha \mid \beta$. This matrix is denoted as $\mathbf{C}_{\beta \rightarrow \alpha}$. The number β is called the *intended* cycle length, while α is called the *target* cycle length.

As their name suggests, CCMs can be used to take a vector of a particular cycle length under the matrix \mathbf{A} and “convert” it to another vector with a different cycle length. While the structure of CCMs may seem odd initially, the reason for their structure isn't terribly difficult to see.

Consider a matrix \mathbf{A} with cycle length 20, and consider the form for the CCM $\mathbf{C}_{10 \rightarrow 5}$:

$$\mathbf{C}_{10 \rightarrow 5} = \mathbf{I} + \mathbf{A}^5.$$

What would happen if we were given a vector \vec{v} with cycle length 10 and multiplied it by $\mathbf{C}_{10 \rightarrow 5}$? We'd end up with the expression

$$\mathbf{C}_{10 \rightarrow 5} \vec{v} = \vec{v} + \mathbf{A}^5 \vec{v}.$$

3.2. Cycle Converting Matrices

Not too exciting. However, things get more interesting if we then multiply the resulting vector by \mathbf{A}^5 :

$$\mathbf{A}^5 \mathbf{C}_{10 \rightarrow 5} \vec{v} = \mathbf{A}^5 \vec{v} + \mathbf{A}^5 \mathbf{A}^5 \vec{v} = \mathbf{A}^5 \vec{v} + \mathbf{A}^{10} \vec{v}.$$

Recall that \vec{v} has a cycle length of 10 (and a transient length of 0 since \mathbf{A} is invertible). This means

$$\mathbf{A}^5 \vec{v} + \mathbf{A}^{10} \vec{v} = \mathbf{A}^5 \vec{v} + \vec{v} = \mathbf{C}_{10 \rightarrow 5} \vec{v}.$$

We see that $\mathbf{A}^5 \mathbf{C}_{10 \rightarrow 5} \vec{v} = \mathbf{C}_{10 \rightarrow 5} \vec{v}$, meaning the cycle length of $\mathbf{C}_{10 \rightarrow 5} \vec{v}$ must divide 5. This is far from a coincidence. The form of the CCM $\mathbf{C}_{10 \rightarrow 5}$ is specifically so that vectors with cycle lengths of 10 will get transformed to vectors of cycle lengths which divide 5 after multiplication by it.

In general, we have that a CCM of the form $\mathbf{C}_{\beta \rightarrow \alpha}$ takes vectors of cycle length β and transforms them to vectors of cycle lengths that divide α .⁴

CCMs are interesting objects in their own right. The following propositions establish some of their many properties.

Proposition 3.6. Let \mathbf{A} be an invertible matrix with multiplicative order ω , and let α, β , and γ be positive integers such that $\alpha \mid \omega$, $\beta \mid \alpha$, and $\gamma \mid \beta$. Then

$$(\mathbf{C}_{\alpha \rightarrow \beta})(\mathbf{C}_{\beta \rightarrow \gamma}) = (\mathbf{C}_{\beta \rightarrow \gamma})(\mathbf{C}_{\alpha \rightarrow \beta}) = \mathbf{C}_{\alpha \rightarrow \gamma}.$$

⁴In fact, any vector whose cycle length *divides* the intended cycle length will be transformed into a vector whose cycle length divides the target cycle length.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Proof. By definition of the CCM, we have that

$$\begin{aligned}
& (\mathbf{C}_{\alpha \rightarrow \beta})(\mathbf{C}_{\beta \rightarrow \gamma}) \\
&= \left(\sum_{i=0}^{\frac{\alpha-\beta}{\beta}} \mathbf{A}^{\beta i} \right) \left(\sum_{j=0}^{\frac{\beta-\gamma}{\gamma}} \mathbf{A}^{\gamma j} \right) \\
&= \mathbf{I}(\mathbf{I} + \mathbf{A}^{\gamma} + \mathbf{A}^{2\gamma} + \dots + \mathbf{A}^{\beta-2\gamma} + \mathbf{A}^{\beta-\gamma}) \\
&+ \mathbf{A}^{\beta}(\mathbf{I} + \mathbf{A}^{\gamma} + \mathbf{A}^{2\gamma} + \dots + \mathbf{A}^{\beta-2\gamma} + \mathbf{A}^{\beta-\gamma}) \\
&+ \mathbf{A}^{2\beta}(\mathbf{I} + \mathbf{A}^{\gamma} + \mathbf{A}^{2\gamma} + \dots + \mathbf{A}^{\beta-2\gamma} + \mathbf{A}^{\beta-\gamma}) \\
&+ \dots \\
&+ \mathbf{A}^{\alpha-2\beta}(\mathbf{I} + \mathbf{A}^{\gamma} + \mathbf{A}^{2\gamma} + \dots + \mathbf{A}^{\beta-2\gamma} + \mathbf{A}^{\beta-\gamma}) \\
&+ \mathbf{A}^{\alpha-\beta}(\mathbf{I} + \mathbf{A}^{\gamma} + \mathbf{A}^{2\gamma} + \dots + \mathbf{A}^{\beta-2\gamma} + \mathbf{A}^{\beta-\gamma}).
\end{aligned}$$

This sum creates a series of increasing powers of \mathbf{A} , from \mathbf{I} to $\mathbf{A}^{\alpha-\gamma}$, each different by a factor of \mathbf{A}^{γ} . This can be rewritten as

$$\sum_{i=0}^{\frac{\alpha-\gamma}{\gamma}} \mathbf{A}^{\gamma i} = \mathbf{C}_{\alpha \rightarrow \gamma}.$$

Note that, because the above expressions are composed entirely of powers of the matrix \mathbf{A} , the order in which $\mathbf{C}_{\alpha \rightarrow \beta}$ and $\mathbf{C}_{\beta \rightarrow \gamma}$ are multiplied makes no difference. **QED**

Proposition 3.6 shows that CCMs have a sort of transitivity, where combining multiple CCMs creates another CCM that converts from the first intended cycle length to the last target cycle length.

Another property of CCMs involves considering LCAs of differing prime-power moduli.

Proposition 3.7. Let p^k be an odd prime power. For an invertible matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$ with multiplicative order ω modulo p^k , if the multiplicative order of \mathbf{A} increases to $p\omega$ modulo p^{k+1} , then

$$\mathbf{C}_{p\omega \rightarrow \omega} \equiv p\mathbf{I} \pmod{p^{k+1}}.$$

3.2. Cycle Converting Matrices

Proof. Computing $\mathbf{C}_{p\omega \rightarrow \omega} \bmod p^{k+1}$, we get that

$$\begin{aligned} & \mathbf{C}_{p\omega \rightarrow \omega} \\ & \equiv \sum_{i=0}^{\frac{p\omega - \omega}{\omega}} \mathbf{A}^{\omega i} \\ & \equiv \mathbf{I} + \mathbf{A}^{\omega} + \mathbf{A}^{2\omega} + \dots + \mathbf{A}^{p\omega - 2\omega} + \mathbf{A}^{p\omega - \omega} \bmod p^{k+1}. \end{aligned}$$

Using Proposition 3.1 and the fact that the multiplicative order of \mathbf{A} modulo p^k is ω , we have that

$$\begin{aligned} & \mathbf{I} + \mathbf{A}^{\omega} + \mathbf{A}^{2\omega} + \dots + \mathbf{A}^{p\omega - 2\omega} + \mathbf{A}^{p\omega - \omega} \\ & \equiv \mathbf{I} + (\mathbf{I} + p^k \mathbf{B}) + \dots + (\mathbf{I} + (p-2)p^k \mathbf{B}) + (\mathbf{I} + (p-1)p^k \mathbf{B}) \\ & \equiv p\mathbf{I} + \frac{p(p-1)}{2} p^k \mathbf{B} \bmod p^{k+1} \end{aligned}$$

for some matrix $\mathbf{B} \in \mathbb{Z}_p^{L \times L}$. Because p is an odd prime, this simplifies to

$$p\mathbf{I} \bmod p^{k+1}.$$

QED

Proposition 3.7 demonstrates one of the many quirks of CCMs. For any matrix whose cycle length increases when going from $\bmod p^k$ to $\bmod p^{k+1}$ (for an odd prime power p^k), the CCM $\mathbf{C}_{p\omega \rightarrow \omega}$ (for ω being the cycle length of the matrix modulo p^k) will always scale the given vector by p . This behaviour makes some amount of sense, seeing as any vector multiplied by p will be an embedded vector, meaning its behaviour will mimic some vector from the $\bmod p^k$ case. We know all the vectors in the $\bmod p^k$ case have cycle lengths that divide ω (since the matrix's cycle length is the maximum cycle length for any vector), and so the resulting vector after being transformed by the CCM will necessarily have a cycle length that divides ω , as desired.

This behaviour, however, raises some interesting questions about how the behaviour of CCMs can be interpreted. As Proposition 3.7 shows, we can't always expect the range of a

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

CCM (i.e. the set of all vectors that can be obtained by multiplying appropriate vectors by the CCM) to give *all* possible vectors with cycle lengths that divide the target cycle length. As an example, consider the LCA $(\mathbb{Z}_{121}, \mathbb{Z}_{121}^3, \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 120 \end{bmatrix})$ and the vector $\vec{v} \equiv \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$. We see \vec{v} has cycle length 10, while the update matrix has cycle length 110. However, because we know

$$\mathbf{C}_{110 \rightarrow 10} \equiv 11\mathbf{I}$$

by Proposition 3.7 (since the cycle length of the update matrix modulo 11 is 10), we can be sure that no vector can be multiplied by the CCM and give \vec{v} (since \vec{v} doesn't have a multiple of 11 in its components), meaning there's at least one vector of the correct target cycle length that the CCM "misses".

What vectors, then, should we expect CCMs to be able to give us? Why is the vector \vec{v} from the example above excluded from the range of the CCM? How many vectors can a CCM "miss"? In fact, it's possible for a CCM to "miss" all vectors of a particular cycle length. For the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix})$, we notice that $\text{span}(\begin{bmatrix} 1 \\ 2 \end{bmatrix})$ is an eigenspace with every vector (excluding the zero vector) having a cycle length of 4. As well, we note that the cycle length of the update matrix is 20. However, it turns out that $\mathbf{C}_{20 \rightarrow 4} \equiv \mathbf{0}$, and so the only vector in this CCM's range is the zero vector. The CCM completely misses our span of vectors!

Thus, CCMs cannot necessarily be used as a reliable way to count vectors with a particular cycle length within an LCA. However, under certain circumstances, CCMs can give us some very useful information regarding cycle lengths that *no* vector has within an LCA.

Proposition 3.8. Let \mathbf{A} be an invertible matrix with cycle length ω , and let α be a positive integer greater than 1 that divides ω . If $\mathbf{C}_{\alpha \rightarrow 1}$ is invertible, then no vectors with cycle length α can exist under \mathbf{A} .

Proof. Let the set of all vectors with cycle lengths that divide α under \mathbf{A} be denoted as V , and let the set of all vectors with cycle length 1 under \mathbf{A} be denoted as N .

If $\mathbf{C}_{\alpha \rightarrow 1}^{-1}$ exists, then there exists a bijection between the vectors in N and the vectors in V . Vectors with cycle length 1 necessarily have cycle lengths that divide α , so $N \subseteq V$. In order for there to be a bijection between N and V , $|N| = |V|$. If $N \subseteq V$, then this implies

3.2. Cycle Converting Matrices

that $N = V$. Therefore, every vector in V has a cycle length of 1, and therefore no vectors with a cycle length of α can exist. **QED**

Proposition 3.8 gives us a way to check whether vectors of a particular cycle length exist within an LCA without having to iterate through every possible vector in the configuration space. As an example, consider again the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, [\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}])$. We compute that

$$\mathbf{C}_{20 \rightarrow 5} \equiv \mathbf{A}^0 + \mathbf{A}^5 + \mathbf{A}^{10} + \mathbf{A}^{15} \equiv \begin{bmatrix} 2 & 2 \\ 2 & 0 \end{bmatrix} \pmod{5},$$

and $[\begin{smallmatrix} 2 & 2 \\ 2 & 0 \end{smallmatrix}] \pmod{5}$ is an invertible matrix. Thus, we immediately know that no vectors of cycle length 5 exist within the given LCA.

In the case of a prime modulus, we can extend the argument of Proposition 3.8 to apply to multiple LCAs at once, though we first need to briefly introduce the concept of a *quotient ring*.

Definition 3.2. Given a ring R and an ideal \mathfrak{I} of R , define the congruence \sim as

$$a \sim b \iff a - b \in \mathfrak{I},$$

for elements $a, b \in R$. The *quotient ring* R/\mathfrak{I} is the ring of all unique equivalence classes

$$[a] = \{a + i : i \in \mathfrak{I}\}$$

under \sim for all $a \in R$ with corresponding operations $+$ and \cdot defined as

$$[a] + [b] = [a + b]$$

and

$$[a] \cdot [b] = [a \cdot b]$$

for $a, b \in R$.

For our purposes, Definition 3.2 is quite formal. At a high level, a quotient ring can be thought of as a way to generalise modular arithmetic. For example, the ring of integers

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

modulo N (which we've denoted as \mathbb{Z}_N) can also be represented as the quotient ring $\mathbb{Z}/N\mathbb{Z}$, where $N\mathbb{Z}$ is the set of all integers which are multiples of N . Each equivalence class in $\mathbb{Z}/N\mathbb{Z}$ represents the set of all integers with a particular remainder after division by N , and so the addition and multiplication specified by Definition 3.2 behave exactly as our usual modular arithmetic.

The quotient rings we're interested in are those of the form $\mathbb{Z}_p[X]/\langle f(X) \rangle$, where p is prime and $f(x)$ is some monic polynomial in $\mathbb{Z}_p[X]$. As with $\mathbb{Z}/N\mathbb{Z}$, the equivalence classes for $\mathbb{Z}_p[X]/\langle f(X) \rangle$ represent sets of polynomials with a particular remainder after polynomial division by $f(x)$. In this way, every polynomial in $\mathbb{Z}_p[X]$ can be reduced to some representative in $\mathbb{Z}_p[X]/\langle f(X) \rangle$ with a degree less than the degree of $f(x)$ (by polynomial division).

A specific example will show why such a construction may be useful to us. Consider a random LCA, say $(\mathbb{Z}_5, \mathbb{Z}_5^3, \begin{bmatrix} 4 & 2 & 4 \\ 0 & 0 & 1 \\ 1 & 4 & 4 \end{bmatrix})$. The update matrix has a minimal polynomial of $m(x) = 4 + 3x + 2x^2 + x^3$. Since our modulus is prime, we know that every annihilating polynomial for the update matrix will be a multiple of $m(x)$.

A prime modulus also means we can use polynomial long division to write every polynomial $f(x) \in \mathbb{Z}_5[X]$ in the form

$$f(x) = q(x)m(x) + r(x)$$

for polynomials $q(x), r(x) \in \mathbb{Z}_5[X]$ where the degree of $r(x)$ is strictly less than the degree of $m(x)$. If \mathbf{A} is our update matrix, then we have that

$$f(\mathbf{A}) \equiv q(\mathbf{A})m(\mathbf{A}) + r(\mathbf{A}) \equiv q(\mathbf{A})(\mathbf{0}) + r(\mathbf{A}) \equiv r(\mathbf{A}) \pmod{5},$$

since $m(x)$ is an annihilating polynomial for \mathbf{A} . We see that the value of $f(\mathbf{A})$ depends only on the value of the remainder $r(\mathbf{A})$. Coincidentally, within the quotient ring $\mathbb{Z}_5[X]/\langle m(X) \rangle$, $f(x) \equiv r(x)$. In essence, the quotient ring $\mathbb{Z}_5[X]/\langle m(X) \rangle$ associates each polynomial with a polynomial of degree strictly less than the degree of $m(x)$ which evaluates to the same matrix when the update matrix is plugged into it. Furthermore, instead of performing operations on polynomials over $\mathbb{Z}_5[X]$, we can take the associated polynomials over $\mathbb{Z}_5[X]/\langle m(X) \rangle$ and perform the operations on them without changing what the polynomial evaluates to when

3.2. Cycle Converting Matrices

plugging in \mathbf{A} (just by how the operations over the quotient ring are defined).

Using quotient rings, questions about a matrix suddenly become questions about how polynomials behave in quotient rings defined by the ideal generated by the matrix's minimal polynomial. While this may seem like a more complicated setting in which to analyse LCAs, it turns out that considering the behaviour of polynomials like this brings us to some pretty interesting conclusions.

Proposition 3.9. Given the LCA $(\mathbb{Z}_p, \mathbb{Z}_p^L, \mathbf{A})$ where \mathbf{A} is invertible, if the polynomial $\sum_{i=0}^{\alpha-1} x^i$ is invertible in the quotient ring $\mathbb{Z}_p[X]/\text{Ann}_{\mathbb{Z}_p[X]}(\mathbf{A})$, then vectors with cycle length α do not exist within the LCA.

Proof. Let $f(x) = \sum_{i=0}^{\alpha-1} x^i$. We see that $f(\mathbf{A}) = \mathbf{C}_{\alpha \rightarrow 1}$. Thus, if we can show that $f(\mathbf{A})$ is invertible, then Proposition 3.8 guarantees that no vectors of cycle length α exist within our given LCA.

Consider the case when $f(x)$ is invertible within $\mathbb{Z}_p[X]/\text{Ann}_{\mathbb{Z}_p[X]}(\mathbf{A})$. Then, there exists some polynomial $g(x) \in \mathbb{Z}_p[X]$ such that $f(x)g(x) \equiv 1 + n(x)$, where $n(x) \in \text{Ann}_{\mathbb{Z}_p[X]}(\mathbf{A})$. This means

$$f(\mathbf{A})g(\mathbf{A}) \equiv \mathbf{I} + n(\mathbf{A}) \equiv \mathbf{I}$$

since $n(x)$ is an annihilating polynomial for \mathbf{A} . Thus, $g(\mathbf{A})$ is the inverse of $f(\mathbf{A})$, meaning $f(\mathbf{A})$ is invertible, meaning vectors of cycle length α cannot exist within our given LCA.

QED

Put simply, Proposition 3.9 allows us to deduce whether vectors of a particular cycle length exist in an LCA simply by considering properties of a few particular polynomials, mainly the update matrix's annihilating polynomials and a sort of polynomial analogue to the special CCM form used in Proposition 3.8. What makes this proposition so powerful is the fact that it doesn't rely on the LCA's specific update matrix, only its ideal of annihilating polynomials. Thus, anything we prove for a particular set of annihilating polynomials and a modulus applies to any LCA with those same annihilating polynomials and that same modulus.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Example 3.1. Consider the quotient ring $\mathbb{Z}_5[X]/\langle X^4 + 2X^3 + 2X^2 + 3X + 3 \rangle$. We find that

$$\sum_{i=0}^9 x^i \equiv (x^4 + 2x^3 + 2x^2 + 3x + 3)(x^5 + 4x^4 + x^3 + 3x^2 + 3x + 4) + 4,$$

and so within the quotient ring, $\sum_{i=0}^9 x^i \equiv 4$. Modulo 5, the inverse of 4 is 4, so

$$4 \sum_{i=0}^9 x^i \equiv 1.$$

Therefore, by Proposition 3.9, any LCA with modulus 5 and an invertible update matrix with $x^4 + 2x^3 + 2x^2 + 3x + 3$ as an annihilating polynomial will never have a vector with a cycle length of 10.

Other than CCMs of the form $\mathbf{C}_{\alpha \rightarrow 1}$, there are many special CCM forms that have interesting properties. The following two propositions establish two such forms.

Proposition 3.10. Let $\mathbf{A} \in \mathbb{Z}_N^{L \times L}$ be an invertible matrix with multiplicative order ω , and α a positive integer that divides ω . If a vector $\vec{v} \in \mathbb{Z}_N^L$ has a multiplicative order under \mathbf{A} that divides α , then

$$\mathbf{C}_{\alpha\beta \rightarrow \alpha} \vec{v} \equiv \beta \vec{v} \pmod{N}$$

for any positive integer β .

Proof. Calculating $\mathbf{C}_{\alpha\beta \rightarrow \alpha} \vec{v}$, we get

$$\begin{aligned} \mathbf{C}_{\alpha\beta \rightarrow \alpha} \vec{v} &\equiv (\mathbf{I} + \mathbf{A}^\alpha + \mathbf{A}^{2\alpha} + \cdots + \mathbf{A}^{\alpha\beta - \alpha}) \vec{v} \\ &\equiv \vec{v} + \mathbf{A}^\alpha \vec{v} + \mathbf{A}^{2\alpha} \vec{v} + \cdots + \mathbf{A}^{\alpha\beta - \alpha} \vec{v} \\ &\equiv \vec{v} + \vec{v} + \vec{v} + \cdots + \vec{v} \\ &\equiv \beta \vec{v} \pmod{N}. \end{aligned}$$

QED

Proposition 3.10 tells us that CCMs of the form $\mathbf{C}_{\alpha\beta \rightarrow \alpha}$ create eigenvectors out of vectors whose cycle lengths divide α under the original update matrix. It might be tempting to try

3.2. Cycle Converting Matrices

and apply this proposition in the other direction; that is, are the eigenvectors of a CCM those vectors with cycle lengths that divide α under the original update matrix?

Unfortunately, Proposition 3.10 is only true in the one direction specified. As a counterexample, consider the LCA $\left(\mathbb{Z}_{11}, \mathbb{Z}_{11}^5, \begin{bmatrix} 6 & 0 & 6 & 7 & 0 \\ 0 & 8 & 10 & 2 & 10 \\ 0 & 10 & 3 & 0 & 4 \\ 0 & 0 & 0 & 2 & 4 \\ 0 & 0 & 0 & 10 & 4 \end{bmatrix}\right)$. If we let \mathbf{A} be the LCA's update matrix, we find that

$$\mathbf{C}_{12 \rightarrow 4} \equiv \mathbf{A}^0 + \mathbf{A}^4 + \mathbf{A}^8 \equiv \begin{bmatrix} 3 & 0 & 0 & 2 & 6 \\ 0 & 3 & 0 & 6 & 6 \\ 0 & 0 & 3 & 3 & 3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \pmod{11}.$$

The basis vector \vec{e}_1 is an eigenvector for $\mathbf{C}_{12 \rightarrow 4}$ since

$$\mathbf{C}_{12 \rightarrow 4} \vec{e}_1 \equiv \begin{bmatrix} 3 & 0 & 0 & 2 & 6 \\ 0 & 3 & 0 & 6 & 6 \\ 0 & 0 & 3 & 3 & 3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \equiv \begin{bmatrix} 3 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \equiv 3\vec{e}_1 \pmod{11}.$$

However, under the LCA's update matrix, \vec{e}_1 has a cycle length of 10, and 10 doesn't divide 4. Thus, we see that Proposition 3.10 does not apply in the other direction; that is, a vector being an eigenvalue for a particular CCM does not imply that its cycle length divides the CCM's target cycle length.

CCMs of the form $\mathbf{C}_{\gamma\alpha \rightarrow \gamma\beta}$ also have some noteworthy properties.

Proposition 3.11. Let \mathbf{A} be an invertible matrix, $\mathbf{C}_{\gamma\alpha \rightarrow \gamma\beta}$ a CCM for \mathbf{A} where $\beta \mid \alpha$, and \vec{v} a vector whose cycle length under \mathbf{A} divides α . If $\gcd(\alpha, \gamma) = 1$, then $\mathbf{C}_{\gamma\alpha \rightarrow \gamma\beta} \vec{v}$ will have a cycle length that divides β .

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Proof. Consider the sequence

$$\beta, 2\beta, 3\beta, \dots, \alpha - 2\beta, \alpha - \beta \pmod{\alpha}. \quad (3.2)$$

Because $\beta \mid \alpha$, this sequence contains all the elements that are multiples of β modulo α . Note that, if we multiply any element in this sequence by some constant modulo α , the element will remain a multiple of β .

Since $\gcd(\alpha, \gamma) = 1$, we know γ^{-1} exists modulo α . Thus, multiplication by γ modulo α is an injective transformation. Now, consider the sequence

$$\gamma\beta, 2\gamma\beta, \dots, \gamma(\alpha - 2\beta), \gamma(\alpha - \beta) \pmod{\alpha}. \quad (3.3)$$

We can be sure that Sequence (3.3) doesn't contain any repeat elements since Sequence (3.2) contains no repeat elements (and since multiplying by γ is an injective transformation). Furthermore, each element in Sequence (3.3) remains a multiple of β , and since there are the same number of elements in Sequence (3.3) as there is in (3.2), this means Sequence (3.3) also contains all multiples of β modulo α , just in a different order.

Now, consider two sums of powers of \mathbf{A} where the matrix powers are taken from the two sequences (and with the identity matrix included):

$$\mathbf{I} + \mathbf{A}^\beta + \mathbf{A}^{2\beta} + \dots + \mathbf{A}^{\alpha-\beta} \quad \text{and} \quad \mathbf{I} + \mathbf{A}^{\gamma\beta} + \mathbf{A}^{2\gamma\beta} + \dots + \mathbf{A}^{\gamma(\alpha-\beta)}.$$

If we multiplied \vec{v} by these two series (and reduced by the appropriate modulus), we'd get the same result for each series since the cycle length of \vec{v} divides α , and so the powers on the matrices can be reduced modulo α when evaluating the vector-matrix multiplication without affecting the result (by virtue of \mathbf{A} being invertible), giving the same sets of matrix powers.

Therefore, we have that

$$\begin{aligned} \mathbf{C}_{\gamma\alpha \rightarrow \gamma\beta} \vec{v} &\equiv \mathbf{I}\vec{v} + \mathbf{A}^{\gamma\beta} \vec{v} + \mathbf{A}^{2\gamma\beta} \vec{v} + \dots + \mathbf{A}^{\gamma(\alpha-\beta)} \vec{v} \\ &\equiv \mathbf{I}\vec{v} + \mathbf{A}^\beta \vec{v} + \mathbf{A}^{2\beta} \vec{v} + \dots + \mathbf{A}^{\alpha-\beta} \vec{v} \\ &\equiv \mathbf{C}_{\alpha \rightarrow \beta} \vec{v}, \end{aligned}$$

3.3. Relations Between Prime & Prime-Power Moduli

and so the cycle length of \vec{v} must divide β .

QED

Proposition 3.11 shows that the behaviour of certain CCMs can be “embedded” into the behaviour of other CCMs. Specifically, for CCMs of the form $\mathbf{C}_{\gamma\alpha\rightarrow\gamma\beta}$ where $\beta \mid \alpha$ and $\gcd(\alpha, \gamma) = 1$, the behaviour of $\mathbf{C}_{\alpha\rightarrow\beta}$ will be present, too, demonstrating a connection between CCMs of similar forms.

All these propositions show how CCMs can not only aid in determining properties of interest for LCAs, but how they are objects in their own right worth exploring.

3.3 Relations Between Prime & Prime-Power Moduli

Throughout this thesis, we’ll make heavy use of the connections between LCAs with prime moduli and LCAs with prime-power moduli. The following propositions will establish some tools for doing this.

Proposition 3.12. If a is a positive integer that’s invertible modulo some prime power p^k (where $k \geq 1$), then a is invertible modulo all prime powers p^ℓ (where p stays fixed).

Proof. Since a is invertible modulo p^k , there exists some number $\alpha \in \mathbb{Z}_{p^k}$ such that

$$a\alpha \equiv 1 \pmod{p^k}. \quad (3.4)$$

Then, for all prime powers p^i where $1 \leq i < k$, reducing Equation (3.4) modulo p^i gives

$$a\alpha \equiv 1 \pmod{p^i},$$

meaning a remains invertible modulo p^i .

Now, considering Equation (3.4) modulo p^j where $j > k$, we have that

$$a\alpha \equiv 1 + p^k n \pmod{p^j}$$

for some integer n . Then we have

$$a(\alpha - p^k n \alpha) \equiv 1 + p^k n - p^k n(1 + p^k n) \equiv 1 - p^{2k} n^2 \pmod{p^j}.$$

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

This means

$$a(1 - p^k n + p^{2k} n^2) \alpha \equiv 1 - p^{2k} n^2 + p^{2k} n^2 (1 + p^k n) \equiv 1 + p^{3k} n^3 \pmod{p^j}.$$

A general pattern emerges:

$$a \alpha \sum_{i=0}^N (-1)^i p^{ik} n^i \equiv 1 + (-1)^N p^{(N+1)k} n^{N+1} \pmod{p^j}$$

for some positive integer N . Then, if $(\mathcal{L} + 1)k \geq j$,

$$a \alpha \sum_{i=0}^{\mathcal{L}} (-1)^i p^{ik} n^i \equiv 1 + (-1)^{\mathcal{L}} p^{(\mathcal{L}+1)k} n^{\mathcal{L}+1} \equiv 1 \pmod{p^j}.$$

So, $\alpha \sum_{i=0}^{\mathcal{L}} (-1)^i p^{ik} n^i$ is the inverse of a modulo p^j , meaning a is invertible modulo p^j .

Putting these two results together, we see that a being invertible modulo p^k implies a is invertible modulo any prime power p^ℓ where p remains fixed. **QED**

Proposition 3.12 is a simple result, but it'll allow us to more easily work with invertible elements modulo prime powers, so it's worth stating explicitly.

Proposition 3.13. Given a prime power p^k , if a set of vectors V is linearly independent modulo p , then they'll be linearly independent modulo p^k .

Proof. Assume otherwise. Then, for the vectors \vec{v}_1 to \vec{v}_r in V , there exist constants a_1 to a_r in \mathbb{Z}_{p^k} , at least one of which is nonzero, such that

$$\sum_{i=1}^r a_i \vec{v}_i \equiv \vec{0} \pmod{p^k}. \tag{3.5}$$

There are two cases we need to consider.

Case 1: At least one of $a_i \not\equiv 0 \pmod{p}$. In this case, reducing Equation (3.5) modulo p gives us that

$$\sum_{i=1}^r a_i \vec{v}_i \equiv \vec{0} \pmod{p}$$

3.3. Relations Between Prime & Prime-Power Moduli

where at least one $a_i \not\equiv 0 \pmod{p}$. This means the vectors \vec{v}_1 to \vec{v}_r are linearly dependent modulo p , which is a contradiction. Therefore, such a situation cannot happen.

Case 2: All of $a_i \equiv 0 \pmod{p}$. In this case, let $\vec{v}_i \equiv a_i \vec{v}_i$, and let

$$P = \log_p(\gcd(p^k, a_1, a_2, \dots, a_r)).$$

By construction of P , every a_i has at least P factors of p in its factorisation, and at least one a_i has *exactly* P . Note that $P < k$ since, if it wasn't, then each of a_i would be zero modulo p^k , which is a contradiction.

Now, rewriting Equation (3.5) using our new notation, we get that

$$\sum_{i=1}^r \vec{v}_i \equiv \vec{0} \pmod{p^k}. \quad (3.6)$$

Because every a_i has at least P factors of p attached to it, every \vec{v}_i is a scalar multiple of an embed vector for \vec{v}_i . Thus, applying the inverse of the bijection ϕ to both sides of Equation (3.6), we have that

$$\sum_{i=1}^r \phi^{-P}(\vec{v}_i) \equiv \phi^{-P}(\vec{0}) \equiv \vec{0} \pmod{p^{k-P}}.$$

At least one of $\phi^{-P}(\vec{v}_i)$ is a scalar multiple of \vec{v}_i which doesn't contain a factor of p (by construction of P). So, if we reduce this congruence modulo p , we'd have that

$$\sum_{i=1}^r b_i \vec{v}_i \equiv \vec{0} \pmod{p}$$

for constants b_1 to b_r in \mathbb{Z}_p , one of which is nonzero. Such a sum cannot happen as the vectors \vec{v}_1 to \vec{v}_r are linearly independent modulo p .

Both cases lead to a contradiction, and so our assumption that V isn't a linearly-independent set modulo p^k must be false. **QED**

Proposition 3.13 tells us that a basis for \mathbb{Z}_p^L , with p a prime and L a positive integer, is also a basis for $\mathbb{Z}_{p^k}^L$, with k another positive integer. Note that a linearly-independent set of vectors modulo p^k is also linearly independent modulo p (which can be seen using the

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

“base- p representation” of vectors used earlier in this thesis), and so Proposition 3.13 works just as well when given a set of linearly-independent vectors modulo p^k for any $k \geq 1$.

Proposition 3.14. Let p^k be an odd prime power for $k > 1$. Given an invertible matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$ and some positive integer w such that $\mathbf{A}^w \equiv \mathbf{I} + p^{k-1}\mathbf{B} \pmod{p^k}$ for some $\mathbf{B} \in \mathbb{Z}_p^{L \times L}$, then

$$\mathbf{A}^{pw} \equiv \mathbf{I} + p^k\mathbf{B} \pmod{p^{k+1}}.$$

Proof. Firstly, we’ll prove that, for any positive integer n , we have that

$$\mathbf{A}^{nw} \equiv \mathbf{I} + np^{k-1}\mathbf{B} + np^k\mathbf{C} + \binom{n}{2}p^{2k-2}\mathbf{B}^2 \pmod{p^{k+1}}$$

for some matrix $\mathbf{C} \in \mathbb{Z}_p^{L \times L}$ and where $\binom{x}{y}$ is the binomial coefficient function. We’ll prove this statement using induction.

Base case: $n = 1$. If $n = 1$, we have that

$$\mathbf{A}^{nw} \equiv \mathbf{A}^w \equiv \mathbf{I} + p^{k-1}\mathbf{B} + p^k\mathbf{C} \pmod{p^{k+1}}$$

for some matrix $\mathbf{C} \in \mathbb{Z}_p^{L \times L}$. This matches the form we want (using the convention that $\binom{1}{2} = 0$), and so our statement is true for $n = 1$.

Induction step. Assume we’re given some positive integer r such that

$$\mathbf{A}^{rw} \equiv \mathbf{I} + rp^{k-1}\mathbf{B} + rp^k\mathbf{C} + \binom{r}{2}p^{2k-2}\mathbf{B}^2 \pmod{p^{k+1}}.$$

3.3. Relations Between Prime & Prime-Power Moduli

Then, expanding $\mathbf{A}^{(r+1)w}$, we see that

$$\begin{aligned}
\mathbf{A}^{(r+1)w} &\equiv (\mathbf{I} + p^{k-1}\mathbf{B} + p^k\mathbf{C})(\mathbf{I} + rp^{k-1}\mathbf{B} + rp^k\mathbf{C} + \binom{r}{2}p^{2k-2}\mathbf{B}^2) \\
&\equiv \mathbf{I} + p^{k-1}\mathbf{B} + p^k\mathbf{C} + rp^{k-1}\mathbf{B} + rp^{2k-2}\mathbf{B}^2 + rp^{2k-1}\mathbf{CB} + rp^k\mathbf{C} \\
&\quad + rp^{2k-1}\mathbf{BC} + rp^{2k}\mathbf{C}^2 + \binom{r}{2}p^{2k-2}\mathbf{B}^2 + \binom{r}{2}p^{3k-3}\mathbf{B}^3 + \binom{r}{2}p^{3k-2}\mathbf{CB}^2 \\
&\equiv \mathbf{I} + (r+1)p^{k-1}\mathbf{B} + (r+1)p^k\mathbf{C} + \left(r + \binom{r}{2}\right)p^{2k-2}\mathbf{B}^2 \pmod{p^{k+1}} \\
&\equiv \mathbf{I} + (r+1)p^{k-1}\mathbf{B} + (r+1)p^k\mathbf{C} + \binom{r+1}{2}p^{2k-2}\mathbf{B}^2.
\end{aligned}$$

So,

$$\begin{aligned}
\mathbf{A}^{rw} &\equiv \mathbf{I} + rp^{k-1}\mathbf{B} + rp^k\mathbf{C} + \binom{r}{2}p^{2k-2}\mathbf{B}^2 \pmod{p^{k+1}} \\
\implies \mathbf{A}^{(r+1)w} &\equiv \mathbf{I} + (r+1)p^{k-1}\mathbf{B} + (r+1)p^k\mathbf{C} + \binom{r+1}{2}p^{2k-2}\mathbf{B}^2 \pmod{p^{k+1}}.
\end{aligned}$$

Therefore, by induction, we can say

$$\mathbf{A}^{nw} \equiv \mathbf{I} + np^{k-1}\mathbf{B} + np^k\mathbf{C} + \binom{n}{2}p^{2k-2}\mathbf{B}^2 \pmod{p^{k+1}}.$$

Now, setting $n = p$ yields

$$\begin{aligned}
\mathbf{A}^{pw} &\equiv \mathbf{I} + pp^{k-1}\mathbf{B} + pp^k\mathbf{C} + \binom{p}{2}p^{2k-2}\mathbf{B}^2 \\
&\equiv \mathbf{I} + p^k\mathbf{B} + p^{k+1}\mathbf{C} + \frac{p-1}{2}p^{2k-1}\mathbf{B}^2 \\
&\equiv \mathbf{I} + p^k\mathbf{B} \pmod{p^{k+1}}
\end{aligned}$$

since p is an odd prime, and since $2k-1 \geq k+1$ for all $k > 1$. Thus, we get that

$$\mathbf{A}^{pw} \equiv \mathbf{I} + p^k\mathbf{B} \pmod{p^{k+1}},$$

which is what we wanted to show.

QED

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

The next proposition shows that properties like multiplicative orders also have relations between LCAs of differing prime-power moduli.

Proposition 3.15. Let p^k be an odd prime power. Suppose that \mathbf{A} is an invertible matrix with multiplicative order ω modulo p^k . Then, the multiplicative order of \mathbf{A} modulo p^{k+1} is either ω or $p\omega$. As well, if the multiplicative order of \mathbf{A} increased to $p\omega$ modulo p^{k+1} , then the multiplicative order of \mathbf{A} modulo $p^{k+\ell}$ will be $p^\ell\omega$ for any positive integer ℓ .

Proof. Let the multiplicative order of \mathbf{A} modulo p^{k+1} be Ω . Assume that $\omega \nmid \Omega$. Then $\Omega = m\omega + n$ for some nonnegative integers m and n where $1 \leq n < \omega$. This means

$$\begin{aligned} \mathbf{A}^\Omega &\equiv \mathbf{I} \pmod{p^{k+1}} \\ \implies \mathbf{A}^\Omega &\equiv \mathbf{I} \pmod{p^k} \\ \implies \mathbf{A}^{m\omega} \mathbf{A}^n &\equiv \mathbf{I} \pmod{p^k} \\ \implies \mathbf{A}^n &\equiv \mathbf{I} \pmod{p^k}. \end{aligned}$$

This is a contradiction since $n < \omega$, yet ω is the multiplicative order of \mathbf{A} modulo p^k . Therefore, $\omega \mid \Omega$.

Now, we know that $\mathbf{A}^\omega \equiv \mathbf{I} \pmod{p^k}$, so

$$\mathbf{A}^\omega \equiv \mathbf{I} + p^k \mathbf{B} \pmod{p^{k+1}}$$

for some $\mathbf{B} \in \mathbb{Z}_p^{L \times L}$.

If $\mathbf{B} \equiv \mathbf{0}$, then we have

$$\mathbf{A}^\omega \equiv \mathbf{I} \pmod{p^{k+1}}.$$

We can show that ω is the smallest such solution to $\mathbf{A}^x \equiv \mathbf{I} \pmod{p^{k+1}}$ for x . Assume a smaller solution y exists. Then

$$\begin{aligned} \mathbf{A}^y &\equiv \mathbf{I} \pmod{p^{k+1}} \\ \implies \mathbf{A}^y &\equiv \mathbf{I} \pmod{p^k}, \end{aligned}$$

3.3. Relations Between Prime & Prime-Power Moduli

and this is a contradiction since $y < \omega$, yet ω is the multiplicative order of \mathbf{A} modulo p^k . So, by definition of matrix multiplicative orders, it must be that $\Omega = \omega$.

Otherwise, if $\mathbf{B} \not\equiv \mathbf{0}$, we can check successive powers of \mathbf{A}^ω to find the first that is equivalent to the identity matrix. Since $\omega \mid \Omega$, we can be sure that the first value of $n\omega$ such that $\mathbf{A}^{n\omega} \equiv \mathbf{I} \pmod{p^{k+1}}$ is indeed the multiplicative order of \mathbf{A} modulo p^{k+1} . Using Proposition 3.1, we see that $n = p$ will be the first such power of \mathbf{A} we're looking for, as

$$\begin{aligned}\mathbf{A}^{p\omega} &\equiv \mathbf{I} + pp^k\mathbf{B} \\ &\equiv \mathbf{I} \pmod{p^{k+1}}.\end{aligned}$$

No positive value of n smaller than p will cause this cancellation. Therefore, in the case where $\mathbf{B} \not\equiv \mathbf{0}$, $\Omega = p\omega$.

These are all the possibilities we need to consider, so $\Omega = \omega$ or $\Omega = p\omega$.

Now, we'll prove that the multiplicative order of \mathbf{A} increasing by a factor of p for a lower power of the modulus means that the multiplicative order of \mathbf{A} must always increase by a factor of p for higher powers of the modulus. If ψ is the multiplicative order of \mathbf{A} modulo p^ℓ for some positive integer ℓ , and $p\psi$ is the multiplicative order of \mathbf{A} modulo $p^{\ell+1}$, then we know

$$\mathbf{A}^\psi \equiv \mathbf{I} + p^\ell \mathbf{C} \pmod{p^{\ell+1}}$$

for some nonzero $\mathbf{C} \in \mathbb{Z}_p^{L \times L}$. Repeatedly applying Proposition 3.14 guarantees that

$$\mathbf{A}^{p^w\psi} \equiv \mathbf{I} + p^{\ell+w} \mathbf{C} \pmod{p^{\ell+1+w}}$$

for positive integers w . There is no possible way $\mathbf{A}^{p^w\psi} \equiv \mathbf{I} \pmod{p^{\ell+1+w}}$ since $\mathbf{C} \not\equiv \mathbf{0}$, meaning $p^w\psi$ can never be the multiplicative order of \mathbf{A} modulo $p^{\ell+1+w}$. Thus, the multiplicative order will be $p^{w+1}\psi$. By inductively applying this reasoning, we see the multiplicative order of \mathbf{A} must increase by a factor of p for each increment of w . **QED**

A statement similar to Proposition 3.15 can be made for the multiplicative orders of vectors. The proof is nearly identical to the proof of Proposition 3.15, and so it will be omitted.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Proposition 3.16. Let p^k be an odd prime power (with $k \geq 1$), and let $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$ be an invertible matrix whose multiplicative order modulo p^k is ω . If $\vec{v} \in \mathbb{Z}_{p^k}^L$ is a maximal vector⁵ under \mathbf{A} modulo p^k , then the multiplicative order of \vec{v} modulo p^{k+1} is either ω or $p\omega$.

3.4 Cycle Spaces

An object of interest to us is the span of a vector's set of iterates under a given update matrix. We call this space the *cycle space* of a vector⁶.

Definition 3.3. The *cycle space* of a vector $\vec{v} \in \mathbb{Z}_p^L$ under a matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$ for some prime p , denoted as $\mathcal{S}_{\vec{v}}$, is defined to be the subspace spanned by the iterates of \vec{v} under iteration by \mathbf{A} . Symbolically,

$$\mathcal{S}_{\vec{v}} = \text{span} \left(\bigcup_{i=0}^{\infty} \{\mathbf{A}^i \vec{v}\} \right) \pmod{p}.$$

The vector \vec{v} is said to be the *generating vector* for $\mathcal{S}_{\vec{v}}$.

For primes p , because \mathbb{Z}_p^L is a finite vector space, cycle spaces, which are subspaces of \mathbb{Z}_p^L , are also finite, despite our definition including an infinite union. In fact, for the definition of the cycle space, it would be sufficient to take the union from $0 \leq i \leq \omega + \tau - 1$, where ω is the cycle length of the generating vector, and τ is the transient length. This is because all the iterates for a vector after the $(\omega + \tau - 1)$ -th iterate will have been iterated to before by definition of the cycle length and transient length.

Cycle spaces are of interest to us as they represent the finest decomposition of an LCA's configuration space into invariant⁷ subspaces under multiplication by the update matrix. It isn't too difficult to see why. Say we wanted the smallest possible invariant subspace under a matrix \mathbf{A} that contained the vector \vec{v} . For the space to be invariant under \mathbf{A} , it better be the case that $\mathbf{A}\vec{v}$, $\mathbf{A}^2\vec{v}$, etc., are in the subspace. In order for us to have a subspace, it must

⁵See Definition 6.1.

⁶In "Dynamics of finite linear cellular automata over \mathbb{Z}_N " (Mendivil and Patterson [4]), cycle spaces are instead called *orbit spaces*.

⁷As a reminder, a subspace is *invariant* under a matrix \mathbf{A} if, for all vectors \vec{v} in the subspace, $\mathbf{A}\vec{v}$ is also in the subspace.

3.4. Cycle Spaces

also be the case that any linear combination of \vec{v} , $\mathbf{A}\vec{v}$, $\mathbf{A}^2\vec{v}$, etc., is also in the subspace. The definition of cycle spaces satisfies these criteria without introducing any additional vectors which don't need to be in the subspace, and so cycle spaces are indeed the smallest possible invariant subspaces that contain a given vector.

There are many reasons to care about invariant subspaces. In a more general context, one may want to consider the eigenspaces of a linear transformation in order to characterise the transformation. Eigenspaces are, by definition, invariant subspaces. For our purposes, we wish to better understand the dynamics of the vectors in an LCA's configuration space, and one way to do that is by understanding the action of the LCA's update matrix on a vector's iterates. The cycle space of a vector provides the smallest subspace in which the iterates of a vector are contained, and so it is in some sense the simplest space to consider for this purpose. In both cases, having a better understanding of how cycle spaces behave could prove useful, and so we'll dedicate this section of the thesis to proving some results regarding these spaces.

A property we can deduce about a vector's cycle space right away is its dimension—the number of linearly-independent vectors needed to span it.

Proposition 3.17. Given a vector $\vec{v} \in \mathbb{Z}_p^L$, p prime, with minimal annihilating polynomial $m(x) \in \mathbb{Z}_p[X]$ under $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, the dimension of the cycle space of \vec{v} is equal to the degree of $m(x)$. Symbolically, $\dim(\mathcal{S}_{\vec{v}}) = \deg(m(x))$.⁸

Proof. The minimal annihilating polynomial is the smallest polynomial $m(x)$ such that $m(\mathbf{A})\vec{v} \equiv \vec{0}$. By rearranging this relation, we find that $\mathbf{A}^{\deg(m(x))}\vec{v}$ is the first iterate of \vec{v} that can be written as a linear combination of its previous iterates (i.e. \vec{v} to $\mathbf{A}^{\deg(m(x))-1}\vec{v}$). So, $\mathcal{S}_{\vec{v}}$ requires $\deg(m(x))$ basis vectors (\vec{v} to $\mathbf{A}^{\deg(m(x))-1}\vec{v}$), and so $\dim(\mathcal{S}_{\vec{v}}) = \deg(m(x))$. **QED**

We can also show that, given a subspace whose minimal annihilating polynomial is a particular degree, there must exist a cycle space within that subspace whose dimension equals the degree of the polynomial.

⁸Note that proofs of this fact have appeared previously. An example is theorem 2 on page 69 of “Lectures in abstract algebra: II. linear algebra” (Jacobson [3]).

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Proposition 3.18. Given an invariant subspace $V \subseteq \mathbb{Z}_p^L$ under $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, p prime, with minimal annihilating polynomial $m(x) \in \mathbb{Z}_p[X]$, there exists a vector $\vec{v} \in V$ such that $\dim(\mathcal{S}_{\vec{v}}) = \deg(m(x))$. Thus, the dimension of V is at least $\deg(m(x))$.

Proof. Let $m(x) = (f_1(x))^{n_1} \times (f_2(x))^{n_2} \times \cdots \times (f_k(x))^{n_k}$ for irreducible, monic, relatively-prime polynomials $f_i(x) \in \mathbb{Z}_p[X]$ and positive integers n_i . If we let $\mathbf{B} = \mathbf{A}|_V$, the restriction of \mathbf{A} to the subspace V , then the Primary Decomposition Theorem says that

$$V = \bigoplus_{i=1}^k \ker((f_i(\mathbf{B}))^{n_i})$$

where each $\ker((f_i(\mathbf{B}))^{n_i}) \neq \{\vec{0}\}$. Thus, for each $\ker((f_i(\mathbf{B}))^{n_i})$, there must exist at least one vector \vec{v}_i such that $(f_i(x))^{n_i}$ is the minimal annihilating polynomial of \vec{v}_i . If this wasn't the case, then there would exist some $\eta_i < n_i$ where $\ker((f_i(\mathbf{B}))^{\eta_i}) = \ker((f_i(\mathbf{B}))^{n_i})$ and so

$$(f_1(x))^{n_1} \times \cdots \times (f_{i-1}(x))^{n_{i-1}} \times (f_i(x))^{\eta_i} \times (f_{i+1}(x))^{n_{i+1}} \times \cdots \times (f_k(x))^{n_k}$$

would be an annihilating polynomial for V . However, this polynomial has a smaller degree than $m(x)$, which contradicts the fact that $m(x)$ is the minimal annihilating polynomial for V , and so this cannot happen.

Because each $\ker((f_i(\mathbf{B}))^{n_i})$ shares only the zero vector with all the other kernels, the sum $\vec{v} \equiv \vec{v}_1 + \vec{v}_2 + \cdots + \vec{v}_k$ has minimal annihilating polynomial

$$(f_1(x))^{n_1} \times (f_2(x))^{n_2} \times \cdots \times (f_k(x))^{n_k} \equiv m(x).$$

By Proposition 3.17, then, $\mathcal{S}_{\vec{v}}$ has dimension $\deg(m(x))$. As well, because V is invariant under \mathbf{A} , $\mathcal{S}_{\vec{v}} \subseteq V$, so $\dim(V) \geq \dim(\mathcal{S}_{\vec{v}}) = \deg(m(x))$. **QED**

Propositions 3.17 and 3.18 show that vectors' cycle spaces are directly related to minimal annihilating polynomials. It's natural, then, to wonder if cycle spaces could prove to be as useful in analysing LCAs as the subspaces given by the Primary Decomposition Theorem. After all, both types of subspaces are directly linked to minimal annihilating polynomials, and both share similar properties (such as being invariant under multiplication by the LCA's

3.4. Cycle Spaces

update matrix). However, as cycle spaces provide a more granular decomposition of an LCA's configuration space, it's possible they could reveal somehow “finer” details about an LCA's behaviour.⁹

Though cycle spaces share many traits with the subspaces given by the Primary Decomposition Theorem, we quickly notice that they don't behave the same. One immediate difference we notice is that the intersections of cycle spaces need not contain only the zero vector, as is the case with the Primary Decomposition Theorem's subspaces. As an example, take the LCA $(\mathbb{Z}_5, \mathbb{Z}_5^2, \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix})$ and consider the two vectors $\vec{v} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $\vec{w} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. We see that

$$\mathcal{S}_{\vec{v}} = \text{span} \left(\left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\} \right) \pmod{5}$$

since $\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} \equiv \begin{bmatrix} 3 \\ 1 \end{bmatrix} \equiv 3 \begin{bmatrix} 1 \\ 2 \end{bmatrix} \pmod{5}$, and

$$\mathcal{S}_{\vec{w}} = \text{span} \left(\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\} \right) \pmod{5},$$

and so

$$\mathcal{S}_{\vec{v}} \cap \mathcal{S}_{\vec{w}} = \mathcal{S}_{\vec{v}}.$$

This example shows that the intersections of cycle spaces are not as predictable as the intersections of the subspaces provided by the Primary Decomposition Theorem. However, our example also displays a potential pattern: is the intersection of two cycle spaces always another cycle space? In this case, the result is trivial since one of the cycle spaces is the entirety of our LCA's configuration space, and so the intersection comes out as the other cycle space. If we were to try other examples, though, we'd see the same behaviour: intersections of cycle spaces seem to yield other cycle spaces. Thus, it may be worthwhile to think about this question. What can we deduce about the intersection of cycle spaces?

Because cycle spaces are invariant subspaces under multiplication by an update matrix, anything we prove about general invariant subspaces will also apply to cycle spaces. The

⁹Our cycle spaces are a special case of a more general space: a *cyclic subspace*. Cyclic subspaces can be used to show lots of interesting results in areas of linear algebra (such as proving the Cayley-Hamilton Theorem; see pages 280 to 285 of *Linear algebra* (Stephen H. Friedberg [7])), so it makes sense to wonder whether they can be used for our purposes, too.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

following proposition is an example.

Proposition 3.19. Let V and W be invariant subspaces of \mathbb{Z}_p^L under $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, p prime. Then $V \cap W$ is a subspace of \mathbb{Z}_p^L which is invariant under \mathbf{A} .

Proof. Let $\vec{x}, \vec{y} \in V \cap W$. This means \vec{x} and \vec{y} are in both V and W . We can then say $\vec{x} + \vec{y}$ is in both V and W since subspaces are closed under addition. Thus, $\vec{x} + \vec{y} \in V \cap W$.

Furthermore, we can say $t\vec{x}$ is in both V and W for $t \in \mathbb{Z}_p$ since subspaces are closed under taking scalar multiples. Thus, $t\vec{x} \in V \cap W$ for $t \in \mathbb{Z}_p$.

Finally, because $V \cap W \subseteq \mathbb{Z}_p^L$ by virtue of both V and W being within \mathbb{Z}_p^L , we have enough to show that $V \cap W$ is a subspace of \mathbb{Z}_p^L .

To show $V \cap W$ is invariant under \mathbf{A} , note that $\mathbf{A}\vec{x}$ is in both V and W since both subspaces are invariant under \mathbf{A} by assumption. Thus, $\mathbf{A}\vec{x} \in V \cap W$, and so any arbitrary vector in the intersection stays within the intersection under multiplication by \mathbf{A} . In other words, $V \cap W$ is invariant under \mathbf{A} . **QED**

Another operation of interest is the *sum* of invariant subspaces. Note that, when we refer to the sum of two subspaces, say $V + W$ for subspaces V and W , we mean the set of vectors $\{\vec{v} + \vec{w} : \vec{v} \in V, \vec{w} \in W\}$. This is slightly different from the *direct* sum of subspaces V and W which requires that for every $\vec{x} \in V + W$ there exists a *unique* pair of vectors $\vec{v} \in V$ and $\vec{w} \in W$ such that $\vec{x} = \vec{v} + \vec{w}$. For this thesis, \oplus will be used to denote the direct sum, while $+$ will be used for this looser sum.

Proposition 3.20. Let V and W be invariant subspaces of \mathbb{Z}_p^L under $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, p prime. Then $V + W$ is a subspace of \mathbb{Z}_p^L which is invariant under \mathbf{A} .

Proof. Let $\vec{x}, \vec{y} \in V + W$. This means both \vec{x} and \vec{y} can be written as a sum of two vectors: one from V and one from W . Therefore, $\vec{x} + \vec{y}$ can also be written in this form simply by adding the respective vectors from each subspace (which can be done since subspaces are closed under addition). So, $\vec{x} + \vec{y} \in V + W$.

Furthermore, a scalar multiple of a sum of two vectors, one from each subspace, will result in another sum of two vectors, one from each subspace, since subspaces are closed under scalar multiplication. Thus, $t\vec{x} \in V + W$ for $t \in \mathbb{Z}_p$.

3.4. Cycle Spaces

By virtue of V and W being subsets of \mathbb{Z}_p^L , $V + W$ must also be a subset of \mathbb{Z}_p^L due to \mathbb{Z}_p^L being a subspace (which is closed under addition). This, combined with what was shown above, is enough to show that $V + W$ is a subspace of \mathbb{Z}_p^L .

As with scalar multiples, a sum of two vectors, one from each subspace, multiplied by \mathbf{A} , must also give a sum of two vectors, one from each subspace, since both W and V are invariant under \mathbf{A} . Therefore, $\mathbf{A}\vec{x} \in V + W$, meaning $V + W$ is invariant under \mathbf{A} . **QED**

Our current conjecture is that the intersection of cycle spaces is itself a cycle space. Proposition 3.20 may have us wondering whether something more general is true. Perhaps *any* combination of two cycle spaces that results in an invariant subspace will be a cycle space.

Unfortunately, examples of sums of cycle spaces exist where the resulting sum is not a cycle space. For instance, take the LCA $\left(\mathbb{Z}_3, \mathbb{Z}_3^3, \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}\right)$ and add the two cycle spaces generated by the vectors $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$. The resulting space, \mathbb{Z}_3^3 , cannot be spanned by a single vector under iteration by the given update matrix as the minimal polynomial for $\begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ modulo 3 is a degree 2 polynomial, meaning every vector in \mathbb{Z}_3^3 under this matrix will have a minimal annihilating polynomial of degree 2 or less, and by Proposition 3.17, this means every vector's cycle space will be of dimension 2 or less.

So, it isn't the case that all "combinations" of cycle spaces that result in invariant subspaces give cycle spaces. Thus, if there is some unique property that intersections of cycle spaces possess so that they themselves are cycle spaces, it goes beyond the fact that they're invariant subspaces. The following proposition will give us some insight into what that unique property may be.

Proposition 3.21. Let $\vec{v} \in \mathbb{Z}_p^L$ be a vector, p prime, with cycle space $\mathcal{S}_{\vec{v}}$ under $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$. If S is an invariant subspace such that $S \subseteq \mathcal{S}_{\vec{v}}$, then

$$\dim(S) = \deg(m(x)),$$

where $m(x) \in \mathbb{Z}_p[X]$ is the minimal annihilating polynomial of S .

Proof. Let $D = \dim(\mathcal{S}_{\vec{v}})$. To show what we want to show, there are two cases we need to consider.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Case 1: $S = \mathcal{S}_{\vec{v}}$. Since $S = \mathcal{S}_{\vec{v}}$, the minimal annihilating polynomial for S will be equal to the minimal annihilating polynomial for $\mathcal{S}_{\vec{v}}$. By Proposition 3.17, the degree of the minimal annihilating polynomial for $\mathcal{S}_{\vec{v}}$ is D , which is the dimension of $\mathcal{S}_{\vec{v}}$ and therefore the dimension of S .

Case 2: $S \subset \mathcal{S}_{\vec{v}}$. Assume the opposite of our proposition. Then, we have that $S \subset \mathcal{S}_{\vec{v}}$, but $\dim(S) > \deg(m(x))$. Note that it isn't possible for the dimension of an invariant subspace to be less than the degree of its minimal annihilating polynomial by Proposition 3.18, and S is certainly an invariant subspace by assumption. So, if $\dim(S) \neq \deg(m(x))$, then it must be the case that the dimension is greater, not less.

To ease with notation, let $n = \deg(m(x))$ and $d = \dim(S)$.

Now, let $B = \{\vec{b}_1, \vec{b}_2, \dots, \vec{b}_d, \vec{b}_{d+1}, \dots, \vec{b}_D\}$ be a basis for $\mathcal{S}_{\vec{v}}$ where $\{\vec{b}_1, \vec{b}_2, \dots, \vec{b}_d\}$ is a basis for S . As well, let $\tilde{B} = \{\vec{b}_{d+1}, \vec{b}_{d+2}, \dots, \vec{b}_D\}$. Then, we can write $\mathbf{A}^i \vec{v}$ as

$$\mathbf{A}^i \vec{v} \equiv \vec{t}_i + \vec{g}_i \pmod{p},$$

where $\vec{t}_i \in S$ and $\vec{g}_i \in \text{span}(\tilde{B})$.

Because $|\tilde{B}| = D - d$ (and because \tilde{B} is a basis for $\text{span}(\tilde{B})$), the maximum number of vectors in a subset of $\text{span}(\tilde{B})$ that can be linearly independent is $D - d$. Thus, it must be the case that the set $\{\vec{g}_0, \vec{g}_1, \dots, \vec{g}_{D-d}\}$ is linearly dependent, meaning there are constants a_0 to a_{D-d} in \mathbb{Z}_p such that

$$a_0 \vec{g}_0 + a_1 \vec{g}_1 + \dots + a_{D-d} \vec{g}_{D-d} \equiv \vec{0} \pmod{p}.$$

Let $r(x) = a_0 + a_1 x + a_2 x^2 + \dots + a_{D-d} x^{D-d}$. Note that $r(x)$ is at most a degree $D - d$

3.4. Cycle Spaces

polynomial. We have that

$$\begin{aligned}
& r(\mathbf{A})\vec{v} \\
& \equiv a_0\vec{v} + a_1\mathbf{A}\vec{v} + a_2\mathbf{A}^2\vec{v} + \cdots + a_{D-d}\mathbf{A}^{D-d}\vec{v} \\
& \equiv a_0(\vec{t}_0 + \vec{g}_0) + a_1(\vec{t}_1 + \vec{g}_1) + \cdots + a_{D-d}(\vec{t}_{D-d} + \vec{g}_{D-d}) \\
& \equiv a_0\vec{t}_0 + a_1\vec{t}_1 + \cdots + a_{D-d}\vec{t}_{D-d} + a_0\vec{g}_0 + a_1\vec{g}_1 + \cdots + a_{D-d}\vec{g}_{D-d} \\
& \equiv a_0\vec{t}_0 + a_1\vec{t}_1 + \cdots + a_{D-d}\vec{t}_{D-d} + \vec{0} \\
& \in S.
\end{aligned}$$

Because $m(x)$ is the minimal annihilating polynomial of S , this means

$$m(\mathbf{A})r(\mathbf{A})\vec{v} \equiv \vec{0} \pmod{p}.$$

This implies that an annihilating polynomial of at most degree $n + D - d$ exists for \vec{v} , and so by Proposition 3.17, $\dim(\mathcal{S}_{\vec{v}}) \leq n + D - d < D$ since $d > n$. This is a contradiction since $\dim(\mathcal{S}_{\vec{v}}) = D$. Thus, our assumption that $\dim(S) > \deg(m(x))$ must be false. **QED**

Proposition 3.21 massively restricts what sorts of cycle space behaviour are possible within another cycle space. For example, say we had a cycle space $\mathcal{S}_{\vec{v}}$ with minimal annihilating polynomial $(f(x))^3$ for some irreducible $f(x)$, and within $\mathcal{S}_{\vec{v}}$ we had two cycle spaces \mathcal{S}_1 and \mathcal{S}_2 with minimal annihilating polynomials $f(x)$ and $(f(x))^2$, respectively. By Proposition 3.21, it must be the case that $\mathcal{S}_1 \subset \mathcal{S}_2$. Otherwise, $\mathcal{S}_1 + \mathcal{S}_2$ (an invariant subspace by Proposition 3.20) would have a dimension greater than $2\deg(f(x))$, while the minimal annihilating polynomial would have a degree of exactly $2\deg(f(x))$, and Proposition 3.21 guarantees that this isn't possible.

Reading Proposition 3.21, we see what sets apart the intersection of two cycle spaces from other “combinations” of cycle spaces (such as a sum): the intersection is necessarily contained within another cycle space. Thus, the proposition applies, and we find that our initial conjecture is correct; the intersection of two cycle spaces must also be a cycle space.

Chapter 3. Understanding the Behaviour of Linear Cellular Automata

Corollary 3.1. Let $\vec{v}, \vec{w} \in \mathbb{Z}_p^L$, p prime. Then for some vector $\vec{x} \in \mathbb{Z}_p^L$,

$$\mathcal{S}_{\vec{v}} \cap \mathcal{S}_{\vec{w}} = \mathcal{S}_{\vec{x}}$$

under some matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$.

Proof. Let $S = \mathcal{S}_{\vec{v}} \cap \mathcal{S}_{\vec{w}}$. By Proposition 3.19, S is an invariant subspace under \mathbf{A} . As well, by definition of intersections, we have that $S \subseteq \mathcal{S}_{\vec{v}}$. Thus, Proposition 3.21 tells us that the dimension of S is equal to the degree of its minimal annihilating polynomial, which we'll label as d . By Proposition 3.18, there must exist some vector $\vec{x} \in S$ whose cycle space has dimension d . Because $\mathcal{S}_{\vec{x}} \subseteq S$, we have that $\mathcal{S}_{\vec{x}} = S$ since their dimensions are the same. **QED**

Corollary 3.1 provides one possible tool for making sense of cycle spaces within an LCA—it gives us a way to understand how they could potentially interact with each other. As cycle spaces are intimately connected to the dynamics of particular vectors within an LCA, this result may prove useful in understanding what sorts of behaviours are possible for the iterates of vectors within a particular LCA.

In Chapter 6, we'll return to cycle spaces by discussing another result relating to them, the Cyclic Decomposition Theorem.

Chapter 4

Ideals of Annihilating Polynomials

As briefly mentioned in Section 2.5, minimal polynomials (of both LCA update matrices and vectors in a configuration space) prove to be some of the most useful objects in analysing LCAs. Via the Minimal Polynomial Theorem, they allow us to more easily determine multiplicative orders and transient lengths, and via the Primary Decomposition Theorem, they allow us to break an LCA’s configuration space into simpler “chunks” that obey simpler algebraic properties. However, minimal polynomials are only guaranteed to exist when the modulus of our LCA is prime, as $\mathbb{Z}_N[X]$ is only a principal ideal domain when N is prime. Put another way, the ideal of annihilating polynomials for a vector/matrix may have more than one “minimal” element when N isn’t prime, and so a minimal polynomial might not exist.

What can we say in the case when the modulus is a prime power? While ideals aren’t guaranteed to be principal over $\mathbb{Z}_{p^k}[X]$ for prime powers p^k , they aren’t guaranteed to *not* be principal, either. Are there cases where an ideal of annihilating polynomials is principal when the modulus is a prime power? If there are, then the Minimal Polynomial Theorem would be able to be used, as there’s nothing in the proof of the theorem that requires a prime modulus—it only assumes the ideal of annihilating polynomials has a single, monic generator.

Chapter 4. Ideals of Annihilating Polynomials

4.1 When Does a “Minimal Polynomial” Exist?

We can immediately show for any prime power p^k and any matrix \mathbf{A} that, however many generators there are for $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$, no generator can have a degree lower than the minimal polynomial of \mathbf{A} .

Proposition 4.1. Let $m(x) \in \mathbb{Z}_p[X]$ be the minimal polynomial for $\mathbf{A} \in \mathbb{Z}^{L \times L}$, p prime. Then, no nonzero polynomial with lesser degree than $m(x)$ can annihilate \mathbf{A} modulo a prime power p^k .

Proof. We'll prove this statement using induction.

Base case: modulo p . By definition, no nonzero polynomial with degree less than the degree of $m(x)$ can annihilate \mathbf{A} modulo p .

Inductive case: modulo p^j . Our induction hypothesis will be that no nonzero polynomial with degree less than the degree of $m(x)$ can annihilate \mathbf{A} modulo p up to modulo p^{j-1} .

Assume that a nonzero annihilating polynomial $r(x) \in \mathbb{Z}_{p^j}[X]$ exists for \mathbf{A} modulo p^j with a lower degree than $m(x)$. Either $r(x)$ reduces to 0 modulo p , or it doesn't.

Since $r(x)$ annihilates \mathbf{A} modulo p^j , it would also annihilate \mathbf{A} when reduced modulo p . However, if $r(x)$ doesn't reduce to 0 modulo p , this creates a contradiction since $r(x)$ has degree less than $m(x)$, and any polynomial with degree less than the minimal polynomial cannot be an annihilating polynomial modulo p by definition.

So, $r(x)$ must reduce to 0 modulo p . Then $r(x)$ must be of the form $pt(x)$ for some polynomial $t(x) \in \mathbb{Z}_{p^{j-1}}[X]$. Note that $t(x)$ must have the same degree as $r(x)$. We find that $t(x)$ must be an annihilating polynomial modulo p^{j-1} :

$$\begin{aligned} r(\mathbf{A}) &\equiv \mathbf{0} \pmod{p^j} \\ \implies pt(\mathbf{A}) &\equiv \mathbf{0} \pmod{p^j} \\ \implies t(\mathbf{A}) &\equiv \mathbf{0} \pmod{p^{j-1}}. \end{aligned}$$

4.1. When Does a “Minimal Polynomial” Exist?

By the induction hypothesis, $t(x)$ is the zero polynomial. Thus, $r(x) \equiv pt(x) \equiv 0 \pmod{p^j}$. This is a contradiction since we assumed $r(x)$ was nonzero. Therefore, no nonzero annihilating polynomial exists for \mathbf{A} modulo p^j with degree less than the degree of $m(x)$.

By induction, our proposition is proved for any arbitrary prime-power modulus p^k .

QED

The proof of the vector case is nearly identical, so we'll list the result without rewriting what we showed above.

Proposition 4.2. Let $m(x) \in \mathbb{Z}_p[X]$ be the minimal annihilating polynomial for $\vec{v} \in \mathbb{Z}^L$ under $\mathbf{A} \in \mathbb{Z}^{L \times L}$, p prime. Then, no nonzero polynomial with lesser degree than $m(x)$ can annihilate \vec{v} under \mathbf{A} modulo a prime power p^k .

Furthermore, we can say something about annihilating polynomials with degrees higher than that of the minimal polynomial. If a monic annihilating polynomial $t(x)$ exists for a matrix modulo p^k with the same degree as the minimal polynomial, then no annihilating polynomial of higher degree can exist for the matrix without being a multiple of $t(x)$.

Proposition 4.3. Let $m(x) \in \mathbb{Z}_p[X]$ be the minimal polynomial for $\mathbf{A} \in \mathbb{Z}^{L \times L}$, p prime. If a monic annihilating polynomial $q(x) \in \mathbb{Z}_{p^k}[X]$ for \mathbf{A} modulo a prime power p^k has the same degree as $m(x)$, then $q(x)$ divides all annihilating polynomials for \mathbf{A} modulo p^k which have degree greater than or equal to the degree of $m(x)$.

Proof. Assume we have another annihilating polynomial $r(x) \in \mathbb{Z}_{p^k}[X]$ modulo p^k with degree greater than or equal to the degree of $q(x)$. Because $q(x)$ is monic, we can use polynomial long division to give us that

$$r(x) \equiv q(x)\tilde{q}(x) + \tilde{r}(x)$$

for polynomials $\tilde{q}(x), \tilde{r}(x) \in \mathbb{Z}_{p^k}[X]$ where the degree of $\tilde{r}(x)$ is strictly less than the degree

Chapter 4. Ideals of Annihilating Polynomials

of $q(x)$. We know $r(x)$ annihilates \mathbf{A} modulo p^k , so

$$\begin{aligned} r(\mathbf{A}) &\equiv \mathbf{0} \\ \implies q(\mathbf{A})\tilde{q}(\mathbf{A}) + \tilde{r}(\mathbf{A}) &\equiv \mathbf{0} \\ \implies (\mathbf{0})\tilde{q}(\mathbf{A}) + \tilde{r}(\mathbf{A}) &\equiv \mathbf{0} \\ \implies \tilde{r}(\mathbf{A}) &\equiv \mathbf{0}, \end{aligned}$$

and so $\tilde{r}(x)$ is an annihilating polynomial for \mathbf{A} modulo p^k . By Proposition 4.1, then, $\tilde{r}(x) \equiv 0 \pmod{p^k}$. Thus, $r(x) \equiv q(x)\tilde{q}(x)$, and so $q(x) \mid r(x)$. **QED**

Again, the proof of the vector case is nearly identical, so we'll list the result here without rewriting what was shown above.

Proposition 4.4. Let $m(x) \in \mathbb{Z}_p[X]$ be the minimal annihilating polynomial for $\vec{v} \in \mathbb{Z}^L$ under $\mathbf{A} \in \mathbb{Z}^{L \times L}$, p prime. If a monic annihilating polynomial $q(x) \in \mathbb{Z}_{p^k}[X]$ for \vec{v} modulo a prime power p^k has the same degree as $m(x)$, then $q(x)$ divides all annihilating polynomials for \vec{v} modulo p^k which have degree greater than or equal to the degree of $m(x)$.

Putting Propositions 4.1 and 4.3 together gives us a fairly nice result about when a matrix's ideal of annihilating polynomials is principal, and therefore has a “minimal polynomial”.

Theorem 3. Let $m(x) \in \mathbb{Z}_p[X]$ be the minimal polynomial for $\mathbf{A} \in \mathbb{Z}^{L \times L}$, p prime. If a monic annihilating polynomial $q(x) \in \mathbb{Z}_{p^k}[X]$ for \mathbf{A} modulo a prime power p^k has the same degree as $m(x)$, then $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$ is generated by a single element, that element being $q(x)$.

Proof. By Proposition 4.1, $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$ can have no elements with degree less than the degree of $q(x)$. By Proposition 4.3, any element in $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$ with degree greater than or equal to the degree of $q(x)$ must be a multiple of $q(x)$. Thus, $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$ is generated by $q(x)$. **QED**

The vector result can be proved using the analogous propositions for vector annihilating polynomials.

4.1. When Does a “Minimal Polynomial” Exist?

Theorem 4. *Let $m(x) \in \mathbb{Z}_p[X]$ be the minimal annihilating polynomial for $\vec{v} \in \mathbb{Z}^L$ under $\mathbf{A} \in \mathbb{Z}^{L \times L}$, p prime. If a monic annihilating polynomial $q(x) \in \mathbb{Z}_{p^k}[X]$ for \vec{v} modulo a prime power p^k has the same degree as $m(x)$, then $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\vec{v})$ is generated by a single element, that element being $q(x)$.*

Whenever Theorem 3 applies, the Minimal Polynomial Theorem also applies, even if the modulus isn’t prime. Though checking whether a monic annihilating polynomial of a certain degree exists for a matrix/vector isn’t a trivial task, there are a few special cases where we can show Theorem 3 always applies.

Consider an LCA where the update matrix’s characteristic polynomial modulo a prime p is equal to its minimal polynomial. By the Cayley-Hamilton Theorem (which works under any ring \mathbb{Z}_N), we know that a matrix’s characteristic polynomial will always be an annihilating polynomial. As well, the characteristic polynomial is always monic, and its degree is given by the size of the matrix; it doesn’t change with the LCA’s modulus. Thus, for any prime-power modulus p^k , the update matrix will always have a monic annihilating polynomial with the same degree as that of its minimal polynomial (namely the characteristic polynomial), and so Theorem 3 will apply.

As a sort of corollary to the above, any update matrix whose characteristic polynomial modulo a prime is irreducible will also satisfy Theorem 3, as in these cases, the characteristic polynomial is automatically equal to the minimal polynomial modulo a prime, and so the above argument will apply.

For another, slightly more contrived example, consider the matrix $n\mathbf{I}$, a scalar multiple of the identity matrix. Modulo a prime, this matrix’s minimal polynomial will be $m(\lambda) = \lambda - n$. No matter the modulus, it’s clear that $\lambda - n$ will always be a monic annihilating polynomial. Thus, Theorem 3 will be satisfied for any LCA with a prime-power modulus that uses $n\mathbf{I}$ as its update matrix.

In all these cases, a “minimal polynomial” of sorts exists for all LCAs with prime-power moduli which use these matrices as their update matrices. Thus, the Minimal Polynomial Theorem can be used to calculate cycle lengths and transient lengths without having to compute them through the dynamics of the system (e.g. using Floyd’s Cycle Detection Algorithm).

Chapter 4. Ideals of Annihilating Polynomials

One final note: if it's the case that an ideal of annihilating polynomials is principal modulo a prime power, it's necessarily the case that the ideal's single generator has a leading term coefficient that's invertible modulo the prime power. Why? Consider the Cayley-Hamilton Theorem, which guarantees that the characteristic polynomial of an LCA's update matrix is always an annihilating polynomial for the matrix (and therefore for the LCA's configuration space). This polynomial is monic. Thus, if an ideal of annihilating polynomials is principal, the characteristic polynomial must be a polynomial multiple of the sole generator, and so it must be possible for the generator's leading term's coefficient to become 1 after some scalar multiplication. Thus, the generator's leading term must be invertible.

Observation 4.1. Modulo a prime power, the generator for a principal ideal of annihilating polynomials must have an invertible coefficient on its leading term. Equivalently, every principal ideal of annihilating polynomials modulo a prime power can be generated by a monic polynomial.

4.2 Annihilating Ideal Generators

It won't always be the case that Theorem 3 will be satisfied for an LCA with a prime-power modulus. In these cases, ideals of annihilating polynomials will have more than one generator, and thus we won't be able to use results like the Minimal Polynomial Theorem to analyse the LCA's behaviour. Given an LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ for some prime power p^k , is there anything we can say about $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$ in the general case?

Although there won't necessarily be a single generator for $\text{Ann}_{\mathbb{Z}_{p^k}[X]}(\mathbf{A})$, the generators we get end up having some fairly nice properties. The following propositions will help us understand this structure.

Proposition 4.5. For some prime power p^k , say we're given a vector $\vec{v} \in \mathbb{Z}_{p^k}^L$, a matrix $\mathbf{A} \in \mathbb{Z}^{L \times L}$, and an annihilating polynomial $f(x) \in \mathbb{Z}_{p^k}[X]$ for \vec{v} modulo p^k with leading term $p^\alpha a_d x^d$, $\alpha < k$, and where a_d^{-1} exists. Any annihilating polynomial for \vec{v} modulo p^k with leading term $p^{\alpha+\beta} b_{d+q} x^{d+q}$, $\beta, q \geq 0$, $\alpha + \beta < k$, $b_{d+q} \not\equiv 0$, can be written as the sum of a polynomial multiple of $f(x)$ and another annihilating polynomial with degree less than $d+q$.

4.2. Annihilating Ideal Generators

Proof. Let

$$f(x) \equiv p^\alpha a_d x^d + \sum_{i=0}^{d-1} a_i x^i \pmod{p^k}$$

for constants $a_i \in \mathbb{Z}_{p^k}$ and let an arbitrary annihilating polynomial $g(x)$ matching our assumptions be written as

$$g(x) \equiv p^{\alpha+\beta} b_{d+q} x^{d+q} + \sum_{j=0}^{d+q-1} b_j x^j \pmod{p^k}$$

for constants $b_j \in \mathbb{Z}_{p^k}$. Multiplying $f(x)$ by $p^\beta b_{d+q} a_d^{-1} x^q$ makes its leading term equal to the leading term in $g(x)$:

$$p^\beta b_{d+q} a_d^{-1} x^q f(x) \equiv p^{\alpha+\beta} b_{d+q} x^{d+q} + p^\beta b_{d+q} a_d^{-1} x^q \left(\sum_{i=0}^{d-1} a_i x^i \right) \pmod{p^k}.$$

Therefore, there exists some polynomial $r(x)$ with degree less than $d+q$ such that

$$p^\beta b_{d+q} a_d^{-1} x^q f(x) + r(x) \equiv g(x) \pmod{p^k}.$$

Also, since $f(x)$ and $g(x)$ are annihilating polynomials for \vec{v} , $r(x)$ must also be an annihilating polynomial since

$$\begin{aligned} p^\beta b_{d+q} a_d^{-1} x^q f(\mathbf{A}) \vec{v} + r(\mathbf{A}) \vec{v} &\equiv g(\mathbf{A}) \vec{v} \pmod{p^k} \\ \Rightarrow \vec{0} + r(\mathbf{A}) \vec{v} &\equiv \vec{0} \pmod{p^k}. \end{aligned}$$

QED

Proposition 4.5 acts like a generalised Division Algorithm for polynomials where the divisor isn't monic. Given two annihilating polynomials $f(x), g(x) \in \mathbb{Z}_{p^k}[X]$ for some vector $\vec{v} \in \mathbb{Z}_{p^k}^L$ where p^k is a prime power, if we have that $\deg(f(x)) \leq \deg(g(x))$ and the leading term of $f(x)$ has less factors of p on it than $g(x)$, then the proposition guarantees there exists polynomials $r(x), u(x) \in \mathbb{Z}_{p^k}[X]$ such that

$$u(x)f(x) + r(x) \equiv g(x) \pmod{p^k}$$

Chapter 4. Ideals of Annihilating Polynomials

where the degree of $r(x)$ is strictly less than the degree of $g(x)$.

Clearly, Proposition 4.5 applies equally well to annihilating polynomials of a matrix.

Proposition 4.6. For some prime power p^k , say we're given a matrix $\mathbf{A} \in \mathbb{Z}^{L \times L}$ and an annihilating polynomial $f(x) \in \mathbb{Z}_{p^k}[X]$ for \mathbf{A} modulo p^k with leading term $p^\alpha a_d x^d$, $\alpha < k$, and where a_d^{-1} exists. Any annihilating polynomial for \mathbf{A} modulo p^k with leading term $p^{\alpha+\beta} b_{d+q} x^{d+q}$, $\beta, q \geq 0$, $\alpha + \beta < k$, $b_{d+q} \not\equiv 0$, can be written as the sum of a polynomial multiple of $f(x)$ and another annihilating polynomial with degree less than $d + q$.

Using Proposition 4.5, we gain a sense as to what the generators for ideals of annihilating polynomials modulo a prime power must look like in the general case.

Proposition 4.7. Modulo a prime power p^k , a maximum of one polynomial per degree is needed as a generator for a matrix's/vector's ideal of annihilating polynomials—namely, a polynomial with the fewest factors of p in its leading term's coefficient compared to all other annihilating polynomials of the same degree.

Proof. Assume we're given an annihilating polynomial $q(x)$ modulo p^k of degree n with the fewest number of factors of p in its leading term's coefficient when compared to all other degree n annihilating polynomials. Let $ap^\alpha x^n$ be the leading term of $q(x)$ where $a \in \mathbb{Z}_{p^k}$ and where a^{-1} exists.

Now, let $s(x)$ be another degree n annihilating polynomial modulo p^k with leading term $bp^{\alpha+\beta} x^n$ where $\beta \geq 0$, $\alpha + \beta < k$, $b \in \mathbb{Z}_{p^k}$, and b^{-1} exists. Proposition 4.5/4.6 says that $s(x)$ can be written as

$$s(x) \equiv a^{-1}bp^\beta q(x) + r(x) \pmod{p^k},$$

where $r(x)$ is an, at most, $n - 1$ degree annihilating polynomial. Therefore, any polynomial of degree n can be written as the sum of a scaled $q(x)$ and a smaller degree annihilating polynomial, meaning $q(x)$ is the only degree n polynomial needed as a generator for the ideal of annihilating polynomials. **QED**

Just from algebraic properties of annihilating polynomials, Proposition 4.7 gives some nice structure to what the set of generators for *any* ideal of annihilating polynomials must look like modulo a prime power. The next proposition will add onto this structure.

4.2. Annihilating Ideal Generators

Proposition 4.8. Let p^k be a prime power, and let $q(x)$ be an annihilating polynomial modulo p^k of degree n with a factor of p^α in its leading term's coefficient ($\alpha < k$). Any annihilating polynomial modulo p^k of degree $n + 1$ with at least a factor of p^α in its leading term's coefficient is not necessary as a generator for the ideal of annihilating polynomials modulo p^k .

Proof. Let

$$q(x) \equiv a_n p^\alpha x^n + \sum_{i=0}^{n-1} a_i x^i \pmod{p^k}$$

for coefficients $a_i \in \mathbb{Z}_{p^k}$ with a_n being invertible, and let

$$s(x) \equiv b_{n+1} p^\alpha x^{n+1} + \sum_{i=0}^n b_i x^i \pmod{p^k}$$

be an arbitrary annihilating polynomial with degree $n + 1$ that has at least a factor of p^α in its leading term's coefficient, with constants $b_i \in \mathbb{Z}_{p^k}$. By Proposition 4.5/4.6, we can write $s(x)$ as

$$s(x) \equiv b_{n+1} a_n^{-1} x q(x) + r(x) \pmod{p^k}$$

with $r(x)$ an annihilating polynomial of degree at most n . Since $q(x)$ and $r(x)$ are both at most degree n annihilating polynomials, they must be generated by polynomials in the ideal of annihilating polynomials with at most degree n . Therefore, $s(x)$ does not need to be a generator for the ideal since it can be represented as a sum of annihilating polynomials that are generated by generators with lower degree. **QED**

While the wording of Proposition 4.8 may be a bit difficult to wrap our heads around, it specifies another nice restriction that the generators for an ideal of annihilating polynomials modulo a prime power must satisfy: as the degrees of the polynomial generators increase, the number of factors of the prime on their leading terms must decrease (this can be shown by repeatedly applying Proposition 4.8 to higher and higher degree polynomials). Immediately, this gives us an upper bound for how many generators an ideal of annihilating polynomials modulo a prime power must have.

Chapter 4. Ideals of Annihilating Polynomials

Proposition 4.9. Let $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ be an LCA for some prime power p^k , and let $m(x)$ be a vector's/matrix's minimal annihilating polynomial modulo p . Then the number of generators for the vector's/matrix's ideal of annihilating polynomials modulo p^k is bounded above by $\min(k, L - \deg(m(x)))$.

Proof. Via Proposition 4.8, we know that the number of factors of p on our generators' leading terms must decrease as their degrees increase. Therefore, the maximum number of generators we can have is the maximum number of leading term coefficients we can list where the number of factors of p decreases with each term. Modulo p^k , the maximum number of terms in such a sequence is k .

As well, by the Cayley-Hamilton Theorem, we know that the characteristic polynomial of \mathbf{A} will always be a monic, degree L annihilating polynomial for both \mathbf{A} and all vectors in $\mathbb{Z}_{p^k}^L$. This means no annihilating polynomial with degree higher than L can be a generator for the ideal of annihilating polynomials since the characteristic polynomial is monic, meaning it isn't possible for higher degree polynomials to have less factors of p on their leading terms, which by Proposition 4.8 means they aren't necessary as generators.

Furthermore, $p^{k-1}m(x)$ will always be an annihilating polynomial modulo p^k since it's an embed of $m(x)$ modulo p^k . By Proposition 4.1/4.2, no polynomial with degree less than $m(x)$ can be an annihilating polynomial. Thus, all generators for our ideal of annihilating polynomials must have degrees of at least $\deg(m(x))$. From above, we know that no generator can have degree higher than L . Since Proposition 4.7 guarantees that only one polynomial per degree is needed as a generator, this implies that there can be no more than $L - \deg(m(x))$ generators.

Therefore, the number of generators for our ideal of annihilating polynomials is bounded above by $\min(k, L - \deg(m(x)))$. **QED**

Because Proposition 4.9 gives us an upper bound on the number of generators for an ideal of annihilating polynomials, and because Proposition 4.7 guarantees that only a maximum of one polynomial per degree is needed as a generator, it's actually possible to describe a general process for finding *all* generators for a vector's/matrix's ideal of annihilating polynomials modulo a prime power.

4.2. Annihilating Ideal Generators

At a high level, the process merely requires us to calculate the iterates of our vector/matrix and use matrix row reduction to determine whether it's possible to write each successive iterate (or a prime-power multiple of it) as a linear combination of the previous iterates. Once we find an iterate which is a linear combination of the previous iterates, *and* where the most recent previous iterate has no factors of the prime in its coefficient in the linear combination, then we know we've found all the ideal generators.

This process is far, far faster than manually checking all possible annihilating polynomials for a given vector/matrix and seeing which are needed to generate the ideal.

Chapter 4. Ideals of Annihilating Polynomials

Chapter 5

The Cores of Linear Cellular Automata

For any LCA, a particular submodule of the configuration space we're interested in understanding is the *core* of the LCA. Intuitively, the core of an LCA is the largest submodule within the configuration space where the update matrix acts as an invertible transformation. Such a submodule is of interest to us as having an update matrix whose transformation is invertible often makes analysing the behaviour of vectors within the LCA much easier. From Figure 2.2, we see that if a vector has a nonzero transient length, there will exist a vector in its set of iterates that has at least two vectors that map to it under the update matrix, meaning an inverse function to undo the transformation can't possibly exist. Thus, to be in the core, a vector must have a transient length of zero. In fact, this is the *only* requirement since it implies only one other vector will iterate to any of the iterates in the vector's set of iterates: the previous iterate.

As an example of why considering the core is useful, suppose that for vectors \vec{v} and \vec{u} in the core of an LCA with update matrix \mathbf{A} , we have the relation

$$\mathbf{A}\vec{v} \equiv \mathbf{A}\vec{u}.$$

Chapter 5. The Cores of Linear Cellular Automata

Because the update matrix's transformation is invertible within the core, this implies that

$$\vec{v} \equiv \vec{u}.$$

If \vec{v} and \vec{u} weren't in the core, we wouldn't be able to immediately conclude this.¹

When the update matrix for an LCA is invertible, an inverse function exists (that function being the inverse of the update matrix) that undoes the transformation of the update matrix on the entire configuration space, and so the core will be the entire configuration space. A consequence of this is that, when the update matrix is invertible, every vector in the configuration space has a transient length of zero. To see this, let τ be the transient length of some vector \vec{v} under the invertible matrix \mathbf{A} . This means there exists some positive integer c where

$$\mathbf{A}^c \mathbf{A}^\tau \vec{v} \equiv \mathbf{A}^\tau \vec{v}.$$

Because \mathbf{A}^{-1} exists, we have that

$$\begin{aligned} \mathbf{A}^{-\tau}(\mathbf{A}^c \mathbf{A}^\tau \vec{v}) &\equiv \mathbf{A}^{-\tau}(\mathbf{A}^\tau \vec{v}) \\ \implies \mathbf{A}^c \vec{v} &\equiv \vec{v}. \end{aligned}$$

By definition of vector transient lengths, this means the transient length of \vec{v} is zero.

For an LCA with a prime modulus, the core is always a subspace of the configuration space. This can be shown by noting that for any two vectors with transient lengths of zero (i.e. two vectors in the core of the LCA), a linear combination of those vectors will also have a transient length of zero, and so the core is closed under linear combinations, making it a subspace. For LCAs with prime-power moduli, the same reasoning can be used to show that the core is always a submodule of the configuration space.

Our goal in this chapter is to show that something stronger is true: the core of an LCA with a prime-power modulus is always a *free* module, a module spanned by a set of linearly-independent vectors. In other words, we want to show that the core of an LCA with a prime-power modulus always has a basis. If this is the case, then for any LCA

¹As an example where this reasoning doesn't apply, consider the LCA $(\mathbb{Z}_9, \mathbb{Z}_9^2, \begin{bmatrix} 3 & 0 \\ 0 & 3 \end{bmatrix})$ and the vectors $\begin{bmatrix} 3 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 3 \end{bmatrix}$. Both vectors iterate to the zero vector, yet they aren't the same vector.

5.1. Dimensional Independence

with a prime-power modulus, the restriction of the update matrix to the LCA's core can be represented by an invertible matrix. In this way, if we concern ourselves only with the cores of LCAs with prime-power moduli, then we can always assume that the update matrix is invertible, allowing for easier algebraic manipulations to be carried out on the vectors in our configuration space.

Before we delve into the cores of LCAs with prime-power moduli, let's specify a more formal definition of the core, along with some notation.

Definition 5.1. Given the LCA $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$, the *core* of the LCA, represented as $\mathcal{K}_{\mathbb{Z}_N^L}(\mathbf{A})$, is the largest submodule of \mathbb{Z}_N^L such that

$$\mathbf{A}\mathcal{K}_{\mathbb{Z}_N^L}(\mathbf{A}) \equiv \mathcal{K}_{\mathbb{Z}_N^L}(\mathbf{A}).$$

In other words, the core is the largest submodule within \mathbb{Z}_N^L where \mathbf{A} is an invertible linear transformation.

There are many equivalent definitions of the core. The definition used in “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]), for instance, takes an infinite intersection of the sequence of submodules $\mathbb{Z}_N^L, \mathbf{A}\mathbb{Z}_N^L, \mathbf{A}^2\mathbb{Z}_N^L$, etc. Another equivalent definition is to set the core as the set of all vectors within the configuration space that have a transient length of zero. The definition we use is chosen to reflect the core's important property of being the largest submodule where the update matrix's transformation is invertible.

5.1 Dimensional Independence

In the case where an LCA's configuration space is of the form $\mathbb{Z}_{p^k}^L$, we must work in the context of modules rather than vector spaces (as would be the case if our configuration space was of the form \mathbb{Z}_p^L). Within modules, the concept of linear independence behaves slightly differently from how we'd like (see below example). As such, we introduce here the idea of *dimensional independence*.

Chapter 5. The Cores of Linear Cellular Automata

Definition 5.2. A set of nonzero vectors V is said to be *dimensionally independent* if

$$\forall \vec{v} \in V, \quad \text{span}(\vec{v}) \cap \text{span}(V \setminus \{\vec{v}\}) = \{\vec{0}\}.$$

As an example for why dimensional independence may be useful, consider the vectors $\begin{bmatrix} 5 \\ 0 \\ 0 \end{bmatrix}$, $\begin{bmatrix} 0 \\ 5 \\ 0 \end{bmatrix}$, and $\begin{bmatrix} 0 \\ 0 \\ 5 \end{bmatrix}$ modulo 25. These vectors aren't linearly independent since

$$5 \begin{bmatrix} 5 \\ 0 \\ 0 \end{bmatrix} + 5 \begin{bmatrix} 0 \\ 5 \\ 0 \end{bmatrix} + 5 \begin{bmatrix} 0 \\ 0 \\ 5 \end{bmatrix} \equiv \vec{0} \pmod{25}. \quad (5.1)$$

However, in some sense, these vectors are still independent since one cannot be made from a linear combination of the others. Other than the zero vector, their spans are disjoint, similar to the spans of linearly-independent vectors. Although these vectors are *not* linearly independent, there may still be certain properties of linear independence that apply to these vectors. Dimensional independence aims to capture these properties.

The following propositions will help solidify the connection between linear independence and dimensional independence.

Proposition 5.1. If $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_g\}$ is a dimensionally-independent set, then the set $\{a_1\vec{v}_1, a_2\vec{v}_2, \dots, a_g\vec{v}_g\}$ is also a dimensionally-independent set for nonzero vectors $a_1\vec{v}_1$ to $a_g\vec{v}_g$.

Proof. For any a_i , $\text{span}(a_i\vec{v}_i) \subseteq \text{span}(\vec{v}_i)$, so

$$\text{span}(\vec{v}_i) \cap \text{span}(V \setminus \{\vec{v}_i\}) = \{\vec{0}\} \implies \text{span}(a_i\vec{v}_i) \cap \text{span}(V \setminus \{\vec{v}_i\}) = \{\vec{0}\}.$$

As well, for any $\vec{v}_j \in V \setminus \{\vec{v}_i\}$, we have that

$$\begin{aligned} \text{span}(v_j) \cap \text{span}(\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_i, \dots, \vec{v}_g\} \setminus \{\vec{v}_j\}) &= \{\vec{0}\} \\ \implies \text{span}(v_j) \cap \text{span}(\{\vec{v}_1, \vec{v}_2, \dots, a_i\vec{v}_i, \dots, \vec{v}_g\} \setminus \{\vec{v}_j\}) &= \{\vec{0}\}. \end{aligned}$$

This means swapping any \vec{v}_i in V with $a_i\vec{v}_i$ (so long as $a_i\vec{v}_i \neq \vec{0}$) will keep V dimensionally independent. **QED**

5.1. Dimensional Independence

Proposition 5.2. If $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_g\}$ is dimensionally independent, then the congruence

$$a_1\vec{v}_1 + a_2\vec{v}_2 + \dots + a_g\vec{v}_g \equiv \vec{0}$$

has only the solution where each $a_i\vec{v}_i \equiv \vec{0}$.

Proof. Assume that a solution to

$$a_1\vec{v}_1 + a_2\vec{v}_2 + \dots + a_g\vec{v}_g \equiv \vec{0}$$

exists where at least one term is nonzero. Rearranging, we see

$$-a_i\vec{v}_i \equiv a_1\vec{v}_1 + a_2\vec{v}_2 + \dots + a_{i-1}\vec{v}_{i-1} + a_{i+1}\vec{v}_{i+1} + \dots + a_g\vec{v}_g$$

for some nonzero term $-a_i\vec{v}_i$. This implies

$$-a_i\vec{v}_i \in \text{span}(V \setminus \{\vec{v}_i\}),$$

which implies

$$\text{span}(\vec{v}_i) \cap \text{span}(V \setminus \{\vec{v}_i\}) \neq \{\vec{0}\}.$$

This is a contradiction since V is dimensionally independent. Therefore, no such solution can exist. **QED**

Much like linearly-independent vectors, dimensionally-independent vectors can only sum to zero if each vector itself is zero. However, since nonzero multiples of dimensionally-independent vectors can be zero (as in Equation (5.1)), there may be multiple distinct linear combinations (that is, distinct sets of coefficients for each linear combination) that result in a sum of zero.

It turns out that, aside from dimensionally-independent vectors behaving similarly to linearly-independent vectors, there's a direct link between sets of linearly-independent vectors and dimensionally-independent vectors.

Proposition 5.3. For some prime-power p^k , if we have a set of dimensionally-independent vectors $V = \{p^{\alpha_1}\vec{v}_1, p^{\alpha_2}\vec{v}_2, \dots, p^{\alpha_g}\vec{v}_g\}$ for integers $0 \leq \alpha_i < k$ and where $\forall \vec{v}_i \in V, \vec{v}_i \not\equiv \vec{0}$

Chapter 5. The Cores of Linear Cellular Automata

$\vec{0} \bmod p$, then the set

$$W = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_g\}$$

is linearly independent modulo p^k .

Proof. Assume W isn't linearly independent modulo p^k . Then there exists a non-trivial solution to

$$\sum_i p^{\beta_i} a_i \vec{v}_i \equiv \vec{0} \bmod p^k$$

for integers a_i and β_i where $a_i \not\equiv 0 \bmod p$ and $\beta_i \geq 0$.

Now, define P as

$$P = \max(\{p^{\alpha_1 - \beta_1}, p^{\alpha_2 - \beta_2}, \dots, p^{\alpha_g - \beta_g}, 1\}).$$

Then, we have that

$$\sum_i P p^{\beta_i} a_i \vec{v}_i \equiv \vec{0} \bmod p^k. \quad (5.2)$$

By construction of P , each $P p^{\beta_i} a_i \vec{v}_i$ is now a multiple of some vector in V . At least one $P p^{\beta_i} a_i \vec{v}_i$ is nonzero, since if $\alpha_j - \beta_j$ is the maximum difference between any $\alpha_i - \beta_i$, then

$$P p^{\beta_j} a_j \vec{v}_j \equiv p^{\alpha_j - \beta_j} p^{\beta_j} a_j \vec{v}_j \equiv p^{\alpha_j} a_j \vec{v} \bmod p^k,$$

which cannot be zero since a_j has no factors of p . By Proposition 5.2, Equation (5.2) implies the set V isn't dimensionally independent. This is a contradiction, so W must be linearly independent. **QED**

Proposition 5.3 shows that, given a set of dimensionally-independent vectors, there's a corresponding set of linearly-independent vectors. This fact will prove useful in Section 5.3, where we use it to strengthen the result of Theorem 5.

We can also say something about dimensionally-independent vectors if they sum to a vector containing factors of p , the base of a prime-power modulus.

Proposition 5.4. Given a prime-power p^k , let $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_g\}$ be dimensionally independent. If

$$\vec{v}_1 + \vec{v}_2 + \dots + \vec{v}_g \equiv p^c \vec{w} \bmod p^k$$

5.2. Creating Dimensionally-Independent Sets

where $\vec{w} \not\equiv \vec{0} \pmod{p}$ and $0 \leq c < k$, then each $\vec{v}_i \equiv \vec{0} \pmod{p^c}$. Furthermore, at least one $\vec{v}_i \not\equiv \vec{0} \pmod{p^{c+1}}$.

Proof. Multiplying both sides of the given equation by p^{k-c} , we get that

$$p^{k-c}(\vec{v}_1 + \vec{v}_2 + \cdots + \vec{v}_g) \equiv p^k \vec{w} \equiv \vec{0} \pmod{p^k}.$$

From Proposition 5.2, the sum $\sum_{i=1}^g p^{k-c} \vec{v}_i$ can only equal zero if each vector in the sum is the zero vector. This means each $p^{k-c} \vec{v}_i \equiv \vec{0} \pmod{p^k}$, and so each \vec{v}_i must have at least a factor of p^c in it. Therefore, $\vec{v}_i \equiv \vec{0} \pmod{p^c}$ for all \vec{v}_i .

Now, let $\vec{v}_i \equiv p^c \vec{v}_i'$. Then

$$p^c(\vec{v}_1' + \vec{v}_2' + \cdots + \vec{v}_g') \equiv p^c \vec{w} \pmod{p^k}.$$

We have that $\vec{w} \not\equiv \vec{0} \pmod{p}$, so $p^{k-1} \vec{w} \not\equiv \vec{0} \pmod{p^k}$. This implies

$$p^{k-1} \vec{w} \equiv p^{k-1}(\vec{v}_1' + \vec{v}_2' + \cdots + \vec{v}_g') \not\equiv \vec{0} \pmod{p^k}.$$

Therefore, at least one $p^{k-1} \vec{v}_i' \not\equiv \vec{0} \pmod{p^k}$, which means $\vec{v}_i' \not\equiv \vec{0} \pmod{p}$, so then

$$p^c \vec{v}_i' \equiv \vec{v}_i \not\equiv \vec{0} \pmod{p^{c+1}}$$

for at least one \vec{v}_i . **QED**

Proposition 5.4 essentially guarantees that, when a set of dimensionally-independent vectors sum to a vector with a certain number of factors of p , then the dimensionally-independent vectors themselves must have a certain number of factors of p . This fact ends up being crucial for Theorem 5.

5.2 Creating Dimensionally-Independent Sets

If we have a dimensionally-independent set, it's helpful to have an easy way to check whether adding a new vector to it will keep it dimensionally independent.

Chapter 5. The Cores of Linear Cellular Automata

Proposition 5.5. Given a prime-power p^k , if $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_g\}$ is dimensionally independent and $\vec{n} \not\equiv \vec{0} \pmod{p^k}$ is a vector such that $\text{span}(\vec{n}) \cap \text{span}(V) \equiv \{\vec{0}\} \pmod{p^k}$, then

$$V \cup \{\vec{n}\}$$

is a dimensionally-independent set.

Proof. To show $V \cup \{\vec{n}\}$ is dimensionally independent, we must show that, for all $\vec{v} \in V$,

$$\text{span}(\vec{v}) \cap \text{span}((V \setminus \{\vec{v}\}) \cup \{\vec{n}\}) \equiv \{\vec{0}\} \pmod{p^k}.$$

Assume there exists some vector $\vec{v}_i \in V$ that doesn't follow the above relation. Then there exists a nonnegative integer c where

$$p^c \vec{v}_i \equiv a_1 \vec{v}_1 + a_2 \vec{v}_2 + \dots + a_{i-1} \vec{v}_{i-1} + a_{i+1} \vec{v}_{i+1} + \dots + a_g \vec{v}_g + a_n \vec{n} \pmod{p^k}$$

and

$$p^c \vec{v}_i \not\equiv \vec{0} \pmod{p^k}.$$

Note that $a_n \vec{n} \not\equiv \vec{0} \pmod{p^k}$ since V is dimensionally independent. Rearranging, we see

$$a_n \vec{n} \equiv p^c \vec{v}_i - a_1 \vec{v}_1 - a_2 \vec{v}_2 - \dots - a_{i-1} \vec{v}_{i-1} - a_{i+1} \vec{v}_{i+1} - \dots - a_g \vec{v}_g \pmod{p^k}.$$

This implies

$$\text{span}(\vec{n}) \cap \text{span}(V) \not\equiv \{\vec{0}\} \pmod{p^k}.$$

This is a contradiction, so no such \vec{v}_i can exist. This means $V \cup \{\vec{n}\}$ is dimensionally independent. **QED**

Proposition 5.5 shows that, if all nonzero multiples of a vector are not in the span of a dimensionally-independent set, we can add the vector to that set and retain dimensional independence.

It's also useful to be able to say something about when adding a vector to a dimensionally-independent set does *not* retain dimensional independence.

5.2. Creating Dimensionally-Independent Sets

Proposition 5.6. For a prime-power p^k , let $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_g\}$ be a dimensionally-independent set, and let \vec{u} be a nonzero vector such that $V \cup \{\vec{u}\}$ is not dimensionally independent. If the maximum number of factors of p that any vector in V has is less than or equal to the number of factors of p in \vec{u} , then given the sum

$$p^c \vec{u} \equiv \sum_i a_i \vec{v}_i \pmod{p^k}$$

where $a_i \vec{v}_i \not\equiv \vec{0} \pmod{p^k}$, we have that $a_i \equiv 0 \pmod{p^c}$ for all a_i .

Proof. By Proposition 5.1 and 5.4, each of $a_i \vec{v}_i$ must have at least the same number of factors of p as $p^c \vec{u}$. If each of \vec{v}_i has at most the same number of factors of p as \vec{u} , then each of a_i must have at least a factor of p^c in order for Proposition 5.4 to hold. So, $a_i \equiv 0 \pmod{p^c}$ for all a_i . **QED**

Intuitively, Proposition 5.6 says that, given a vector \vec{u} that's not dimensionally independent to a dimensionally-independent set, the number of factors of p in \vec{u} can give some idea for what an expression for $p^c \vec{u}$ would look like in terms of the vectors in the dimensionally-independent set. Ultimately, it's this proposition which allows the following theorem to be shown.

Theorem 5. For a prime-power p^k , any nonempty submodule M of $\mathbb{Z}_{p^k}^L$ can be expressed as the span of a set of dimensionally-independent vectors.

Proof. If $M = \{\vec{0}\}$, then let $V = \emptyset$. We have that $\text{span}(V) = \{\vec{0}\}$, and V is dimensionally independent.

Otherwise, every submodule is closed under taking linear combinations, so every submodule of $\mathbb{Z}_{p^k}^L$ can be represented as the span of a set of vectors. Let

$$\hat{M} = \{\vec{m}_1, \vec{m}_2, \dots, \vec{m}_g\}$$

where $M = \text{span}(\hat{M})$. Without loss of generality, assume $\vec{0} \notin \hat{M}$.

The following procedure will produce a dimensionally-independent set of vectors whose span is M .

Chapter 5. The Cores of Linear Cellular Automata

Step 1. Order the vectors in \hat{M} by the number of factors of p they have (least to greatest) and store them in an ordered list S . For example, if \hat{M} were to look like

$$\hat{M} = \{p^2\vec{n}_1, p^3\vec{n}_2, \vec{n}_3, p^2\vec{n}_4, \vec{n}_5\}$$

where \vec{n}_1 to \vec{n}_5 have no factors of p , then a valid list S could look like

$$S = [\vec{n}_3, \vec{n}_5, p^2\vec{n}_1, p^2\vec{n}_4, p^3\vec{n}_2].$$

Another valid list S could be

$$S = [\vec{n}_5, \vec{n}_3, p^2\vec{n}_4, p^2\vec{n}_1, p^3\vec{n}_2]. \quad (5.3)$$

Let S_i index the i -th element in S . Let $|S|$ count the number of elements in S . As an example, for the list S specified by Equation (5.3), $S_2 = \vec{n}_3$ and $|S| = 5$.

Finally, initialise V to be the empty set.

Step 2. Set $\vec{s} = S_1$.

Step 3. Check to see whether $V \cup \{\vec{s}\}$ is a dimensionally-independent set. If it is, replace V with $V \cup \{\vec{s}\}$ and proceed to step 5. Otherwise, continue to step 4.

Step 4. If $V \cup \{\vec{s}\}$ is not a dimensionally-independent set, then for some smallest p^c , $0 \leq c < k$, we have that

$$p^c\vec{s} \equiv \sum_i a_i \vec{v}_i \pmod{p^k}$$

for vectors $\vec{v}_i \in V$ and nonzero constants $a_i \in \mathbb{Z}_{p^k}$ where $a_i \vec{v}_i \not\equiv \vec{0} \pmod{p^k}$. Because of the ordering of S , the vectors in V are guaranteed to have at most the same number of factors of p in them as \vec{s} . This means we can make use of Proposition 5.6 and rewrite the sum for $p^c\vec{s}$ as

$$p^c\vec{s} \equiv p^c \sum_i b_i \vec{v}_i \pmod{p^k}$$

for nonzero constants $b_i \in \mathbb{Z}_{p^k}$ where $b_i \vec{v}_i \not\equiv \vec{0} \pmod{p^k}$. Define \vec{x} to be

$$\vec{x} \equiv \vec{s} - \sum_i b_i \vec{v}_i \pmod{p^k}.$$

5.2. Creating Dimensionally-Independent Sets

If \vec{x} happens to be the zero vector, discard it and proceed to step 5.

By construction, $p^c \vec{x} \equiv \vec{0} \pmod{p^k}$. As well, for all $0 \leq f < c$, we have that

$$\begin{aligned} p^f \vec{x} &\equiv p^f \vec{s} - p^f \sum_i b_i \vec{v}_i \pmod{p^k} \\ \implies p^f \vec{x} &\notin \text{span}(V) \end{aligned}$$

since $p^f \vec{x}$ is the sum of a vector in $\text{span}(V)$ and a vector not in $\text{span}(V)$ (we know that $p^f \vec{s} \notin \text{span}(V)$ since $p^c \vec{s}$ is the first nonzero multiple of \vec{s} in $\text{span}(V)$).

By Proposition 5.5, this means the set $V \cup \{\vec{x}\}$ is dimensionally independent. Also, $\text{span}(V \cup \{\vec{s}\}) \equiv \text{span}(V \cup \{\vec{x}\})$ since

$$\vec{x} \equiv \vec{s} - \sum_i b_i \vec{v}_i \pmod{p^k} \implies \vec{s} \equiv \vec{x} + \sum_i b_i \vec{v}_i \pmod{p^k},$$

so any linear combination made with the vectors in $V \cup \{\vec{s}\}$ can be made with the vectors in $V \cup \{\vec{x}\}$, and vice versa.

If the number of factors of p in \vec{x} is at most the same as the number of factors of p in \vec{s} , replace V with $V \cup \{\vec{x}\}$ and proceed to step 5.

Otherwise, if $\vec{s} = S_i$, insert \vec{x} into S somewhere greater than the i -th position so that the ordering of S is preserved (i.e. the vectors in S are still ordered from least number of factors of p to greatest). Note that we're not removing the current \vec{s} from S ; we're increasing the number of vectors in S by inserting \vec{x} . The list S , then, can be imagined as a queue of sorts, with the current \vec{s} representing which element in the queue we're currently using.

The instruction above ensures that the vector \vec{x} can always be created in step 4; each vector \vec{s} will always have at least the same number of factors of p as any vector in V .

Step 5. If $\vec{s} = S_i$ where $i < |S|$, set $\vec{s} = S_{i+1}$ and return to step 3. Otherwise, continue to step 6.

Note that \vec{s} is set to each vector in S only once, and so each vector in S can lead to the insertion of only one extra vector into S (via step 4). However, the extra vector inserted into S must necessarily have a greater number of factors of p than the vector that led to its insertion. Eventually, if vectors continue to be added, the new inserted vectors will be unable to lead to more inserted vectors since they'll have so many factors of p that any vectors with a greater number of factors of p will be congruent to the zero vector (which

Chapter 5. The Cores of Linear Cellular Automata

step 4 will discard). Therefore, this process is guaranteed to eventually proceed to step 6; it is impossible to indefinitely insert vectors into S .

Step 6. V is now a dimensionally-independent set of vectors whose span is M . **QED**

Theorem 5 establishes another similarity between linear and dimensional independence: every subspace of \mathbb{Z}_p^L has a set of linearly-independent vectors that span it, while every submodule of $\mathbb{Z}_{p^k}^L$ has a set of dimensionally-independent vectors that span it. This gives us a guaranteed way to represent any submodule of $\mathbb{Z}_{p^k}^L$, allowing us to more easily work with arbitrary submodules.

The following gives an example of using Theorem 5 to find a set of dimensionally-independent vectors that spans a given submodule.

Example 5.1. Let $M = \text{span}\left(\left\{\begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix}, \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix}, \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix}\right\}\right) \bmod 25$. We want to find a set of dimensionally-independent vectors whose span is M . For this, we'll use the procedure defined by Theorem 5.

In summary, the steps we'll use are as follows:

1. Order the vectors in our set by ascending order of factors of p present. Store them in an ordered list S . As well, set $V = \emptyset$.
2. Set our “pointer” vector \vec{s} to be the first vector in this list.
3. Check whether $V \cup \{\vec{s}\}$ is dimensionally independent. If it is, replace V with $V \cup \{\vec{s}\}$ and proceed to step 5. Otherwise, proceed to step 4.
4. Find the smallest positive integer c such that $p^c \vec{s} \equiv p^c \sum_i b_i \vec{v}_i \pmod{p^k}$ for vectors $\vec{v}_i \in V$ and nonzero coefficients $b_i \in \mathbb{Z}_{p^k}$ where $b_i \vec{v}_i \not\equiv \vec{0} \pmod{p^k}$. Define \vec{x} as

$$\vec{x} \equiv \vec{s} - \sum_i b_i \vec{v}_i \pmod{p^k}.$$

If $\vec{x} \equiv \vec{0} \pmod{p^k}$, proceed to step 5. Else, if the number of factors of p in \vec{x} is at most the same as the number of factors of p in \vec{s} , replace V with $V \cup \{\vec{x}\}$ and proceed to step 5. Otherwise, insert \vec{x} into S at an index greater than the index of the vector to which \vec{s} points (so that the ordering of S by ascending number of factors of p is preserved).

5.2. Creating Dimensionally-Independent Sets

5. Set \vec{s} to point to the next vector in S (the one with an index one greater than the one it's currently pointing to) and return to step 3. If no such next vector exists, continue to step 6.
6. The set V is now a dimensionally-independent set whose span is the same as M .

Using these steps, we can find a dimensionally-independent set V whose span is the same as the set M .

Step 1. We order our list of vectors that span M from least number of factors of 5 to greatest. This gives us the list

$$S = \left[\begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}, \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix}, \begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix}, \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix} \right].$$

As well, we initialise $V = \emptyset$.

Step 2. We set $\vec{s} = \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}$, which is the first vector in S .

Step 3-1. Because $V = \emptyset$, we know that $V \cup \{\vec{s}\}$ is a dimensionally-independent set. Thus, we replace V with $V \cup \{\vec{s}\} = \left\{ \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \right\}$.

Step 5-1. We set \vec{s} to the next vector in S , which is $\begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix}$.

Step 3-2. The vectors $\begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix}$ are not dimensionally independent since

$$10 \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} 10 \\ 20 \\ 10 \end{bmatrix} \equiv 5 \begin{bmatrix} 2 \\ 4 \\ 2 \end{bmatrix} \pmod{25}, \quad (5.4)$$

and so $\text{span} \left(\begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \right) \cap \text{span} \left(\begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix} \right) \neq \{\vec{0}\}$. Thus, $V \cup \left\{ \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix} \right\}$ is not a dimensionally-independent set.

Step 4-1. Using Congruence (5.4), we have that

$$5 \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix} \equiv 5(2) \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \pmod{25},$$

Chapter 5. The Cores of Linear Cellular Automata

so we define \vec{x} to be

$$\begin{aligned}\vec{x} &\equiv \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix} - 2 \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \\ &\equiv \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix} - \begin{bmatrix} 2 \\ 14 \\ 2 \end{bmatrix} \\ &\equiv \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix} \pmod{25}.\end{aligned}$$

We see that \vec{x} has a factor of 5 in it while \vec{s} doesn't, so we insert \vec{x} into S so as not to disturb the ordering of S . Our list S becomes

$$S = \left[\begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}, \begin{bmatrix} 7 \\ 4 \\ 12 \end{bmatrix}, \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix}, \begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix}, \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix} \right].$$

Step 5-2. We set \vec{s} to the next vector in S , which is $\begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix}$.

Step 3-3. We see that the set $V \cup \left\{ \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix} \right\}$ is a dimensionally-independent set since $\text{span} \left(\begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \right) \cap \text{span} \left(\begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix} \right) \equiv \{\vec{0}\}$, and so we replace V with $V \cup \left\{ \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix} \right\} = \left\{ \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix} \right\}$.

Step 5-3. We set \vec{s} to the next vector in S , which is $\begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix}$.

Step 3-4. The set $V \cup \left\{ \begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix} \right\}$ is not dimensionally independent since $15 \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix}$.

Step 4-2. From above, we see that

$$1 \begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix} \equiv 1(15) \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \pmod{25},$$

so we define \vec{x} to be

$$\vec{x} \equiv \begin{bmatrix} 15 \\ 5 \\ 15 \end{bmatrix} - 15 \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix} \equiv \vec{0} \pmod{25}.$$

Since \vec{x} is the zero vector, we discard it.

Step 5-4. We set \vec{s} to be the next vector in S , which is $\begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix}$.

Step 3-5. If we try to create $\begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix}$, or any non-multiple-of- p multiple of it, as a linear combination of the vectors in V , we'll find that we can't.² Thus, $V \cup \left\{ \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix} \right\}$ is a dimensionally-independent set, so we replace V with $V \cup \left\{ \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix} \right\} \equiv \left\{ \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix}, \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix} \right\}$.

²We only need to check non-multiples-of- p of the vector since any multiple which is a multiple of p will necessarily reduce the span of the vector, which goes against the intent of the process.

5.3. Prime-Power Cores

Step 6. V is now a set of dimensionally-independent vectors that spans M . Thus, we have that

$$M \equiv \text{span} \left(\left\{ \begin{bmatrix} 1 \\ 7 \\ 1 \end{bmatrix}, \begin{bmatrix} 5 \\ 15 \\ 10 \end{bmatrix}, \begin{bmatrix} 5 \\ 10 \\ 15 \end{bmatrix} \right\} \right).$$

◇

With Theorem 5 establishing a general structure that all submodules have, we can use this to begin thinking about the potential structure of LCA cores.

5.3 Prime-Power Cores

Given an LCA of the form $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ for some prime-power p^k ($k > 1$), one way we might try to understand its core is by understanding the core of the related LCA $(\mathbb{Z}_{p^{k-1}}, \mathbb{Z}_{p^{k-1}}^L, \mathbf{A})$ and using the relationship between prime-power moduli to somehow extrapolate information about the other. One of the most direct ways we have to relate LCAs is through embed vectors, which create a bijective mapping between LCAs of differing prime-power moduli. The following proposition makes use of embed vectors to create a relationship between the cores of LCAs with differing prime-power moduli.

Proposition 5.7. Given the LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ for some prime-power p^k , we have that

$$\vec{v} \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A}) \iff \phi(\vec{v}) \in \mathcal{K}_{\mathbb{Z}_{p^{k+1}}^L}(\mathbf{A}),$$

where ϕ is the embedding bijection from $\mathbb{Z}_{p^k}^L$ to $p\mathbb{Z}_{p^{k+1}}^L$.

Proof. First, we'll show $\vec{v} \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A}) \implies \phi(\vec{v}) \in \mathcal{K}_{\mathbb{Z}_{p^{k+1}}^L}(\mathbf{A})$.

If $\vec{v} \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$, then there must exist some smallest positive integer c where

$$\mathbf{A}^c \vec{v} \equiv \vec{v} \pmod{p^k}$$

since \vec{v} has a transient length of zero. The mapping ϕ creates a bijection onto $p\mathbb{Z}_{p^{k+1}}^L$, so

$$\phi(\mathbf{A}^c \vec{v}) \equiv \phi(\vec{v}) \pmod{p^{k+1}}.$$

Chapter 5. The Cores of Linear Cellular Automata

Since ϕ preserves matrix multiplication, we have that

$$\phi(\mathbf{A}^c \vec{v}) \equiv \mathbf{A}^c \phi(\vec{v}) \equiv \phi(\vec{v}) \pmod{p^{k+1}}.$$

From this, we see that $\phi(\vec{v})$ has a transient length of zero modulo p^{k+1} , and so it must be in $\mathcal{K}_{\mathbb{Z}_{p^{k+1}}^L}(\mathbf{A})$.

Next, we'll show $\phi(\vec{v}) \in \mathcal{K}_{\mathbb{Z}_{p^{k+1}}^L}(\mathbf{A}) \implies \vec{v} \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$.

If $\phi(\vec{v}) \in \mathcal{K}_{\mathbb{Z}_{p^{k+1}}^L}(\mathbf{A})$, then there must exist some smallest positive integer c where

$$\mathbf{A}^c \phi(\vec{v}) \equiv \phi(\vec{v}) \pmod{p^{k+1}}$$

since $\phi(\vec{v})$ has a transient length of zero. The mapping ϕ preserves matrix multiplication, so

$$\mathbf{A}^c \phi(\vec{v}) \equiv \phi(\mathbf{A}^c \vec{v}) \equiv \phi(\vec{v}) \pmod{p^{k+1}}.$$

The mapping ϕ creates a bijection (as above), so ϕ^{-1} can be applied to the above equation:

$$\mathbf{A}^c \vec{v} \equiv \vec{v} \pmod{p^k}.$$

We see that \vec{v} has a transient length of zero modulo p^k , and so it must be in $\mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$. **QED**

Proposition 5.7 gives us a way to relate the cores of LCAs with differing prime-power moduli using embed vectors. However, embed vectors aren't the only tool we have to relate LCAs. The following proposition uses lift vectors to relate different cores.

Proposition 5.8. For a given prime-power p^k , if $\vec{v} \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$, then a lift vector $\vec{\ell} \in \mathbb{Z}_{p^{k+n}}^L$ of \vec{v} exists such that

$$\vec{\ell} \in \mathcal{K}_{\mathbb{Z}_{p^{k+n}}^L}(\mathbf{A}).$$

Proof. We have that $\vec{v} \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$, so there must exist a smallest positive integer c such that

$$\mathbf{A}^c \vec{v} \equiv \vec{v} \pmod{p^k}.$$

5.3. Prime-Power Cores

Lifting this congruence to a higher-power modulus gives

$$\mathbf{A}^c \vec{v} \equiv \vec{v} + p^k \vec{u} \pmod{p^{k+n}}, \quad \vec{u} \in \mathbb{Z}_{p^n}^L.$$

Taking this new vector and multiplying it by \mathbf{A}^c , we get

$$\mathbf{A}^c(\vec{v} + p^k \vec{u}) \equiv \vec{v} + p^k(\vec{u} + \mathbf{A}^c \vec{u}) \pmod{p^{k+n}}.$$

In general, if we continue taking these vectors and multiplying them by \mathbf{A}^c , we'll get a sequence of vectors where each term looks like

$$\vec{v} + p^k \vec{w}_i \pmod{p^{k+n}}, \quad \vec{w}_i \in \mathbb{Z}_{p^n}^L.$$

There are only a finite number of vectors in $\mathbb{Z}_{p^n}^L$, so eventually $\vec{w}_x \equiv \vec{w}_y$ for some integers x and y , $x > y$. Let \vec{W} be the first such \vec{w}_x where this occurs. Then, if we let

$$\vec{\ell} \equiv \vec{v} + p^k \vec{W} \pmod{p^{k+n}},$$

then $\vec{\ell}$ is guaranteed to be in $\mathcal{K}_{\mathbb{Z}_{p^{k+n}}^L}(\mathbf{A})$ by construction. The vector $\vec{\ell}$ also happens to be a lift of \vec{v} since

$$\vec{\ell} \equiv \vec{v} \pmod{p^k}.$$

QED

While Proposition 5.8 doesn't explicitly give us a lift vector for a given vector in a related core, it does guarantee such a vector exists, which is more than enough for our purposes.

With Propositions 5.7 and 5.8 established, we can show that the cores of LCAs with prime-power moduli must necessarily be spanned by a set of linearly-independent vectors by making use of the above relationships between the cores of LCAs with differing prime-power moduli.

Theorem 6. *For any LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ where p^k is a prime-power, $\mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$ can be represented as the span of a set of linearly-independent vectors.*

Chapter 5. The Cores of Linear Cellular Automata

Proof. If $\mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A}) = \{\vec{0}\}$, then let $M = \emptyset$. We have that $\text{span}(M) = \{\vec{0}\}$, and M is linearly independent.

Otherwise, let $M = \{\vec{m}_1, \vec{m}_2, \dots, \vec{m}_g\}$ be a set of dimensionally-independent vectors where $\text{span}(M) = \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$. Such a set is guaranteed to exist by Theorem 5. Note that each \vec{m}_i must necessarily be in $\mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$.

For each vector $\vec{m}_i \in M$, the following argument shows that $\vec{m}_i \not\equiv \vec{0} \pmod{p}$:

Let $\vec{m}_i \equiv p^{t_i} \vec{n}_i$ where $t_i < k$ and $\vec{n}_i \not\equiv \vec{0} \pmod{p}$. By repeated application of Proposition 5.7, we know that

$$\vec{n}_i \in \mathcal{K}_{\mathbb{Z}_{p^{k-t_i}}^L}(\mathbf{A}).$$

By Proposition 5.8, there exists a lift vector $\vec{\ell}_i$ of \vec{n}_i such that

$$\vec{\ell}_i \in \mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A}).$$

We note that

$$p^{t_i} \vec{\ell}_i \equiv p^{t_i} (\vec{n}_i + p^{k-t_i} \vec{z}) \equiv p^{t_i} \vec{n}_i + p^k \vec{z} \equiv p^{t_i} \vec{n}_i \equiv \vec{m}_i \pmod{p^k}$$

for some vector $\vec{z} \in \mathbb{Z}_{p^k}^L$ since $\vec{\ell}_i$ is a lift of \vec{n}_i .

Since $\vec{\ell}_i$ is in $\mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$, it must also be in $\text{span}(M)$. This means $\vec{\ell}_i$ can be represented as a linear combination of vectors in M :

$$\vec{\ell}_i \equiv \sum_j a_j \vec{m}_j \pmod{p^k}$$

where each $a_j \in \mathbb{Z}_{p^k}$ and each $a_j \vec{m}_j \not\equiv \vec{0} \pmod{p^k}$. This means

$$\vec{m}_i \equiv p^{t_i} \vec{\ell}_i \equiv p^{t_i} \sum_j a_j \vec{m}_j \pmod{p^k}.$$

The set M is dimensionally independent, so

$$\vec{m}_i \equiv p^{t_i} \sum_j a_j \vec{m}_j \equiv p^{t_i} a_i \vec{m}_i \pmod{p^k}.$$

5.3. Prime-Power Cores

The only way this congruence can be true is if $t_i = 0$. Therefore, $\vec{m}_i \equiv \vec{n}_i \pmod{p^k}$, and $\vec{n}_i \not\equiv \vec{0} \pmod{p}$.

With this argument, we can be sure that no vector in M has a multiple of p attached to it. By Proposition 5.3, then, M is guaranteed to be a set of linearly-independent vectors. **QED**

Theorem 6 shows that the core of any LCA with a prime-power modulus is always a free module, meaning it has a basis. This means any LCA with a prime-power modulus has an update matrix that, when restricted to the LCA's core, can be represented as an invertible matrix.

Chapter 5. The Cores of Linear Cellular Automata

Chapter 6

The Existence of Maximal Vectors

From Section 3.1, we know that no vector in an LCA can have a cycle length greater than the cycle length of the LCA's update matrix. The update matrix's cycle length is a sort of “maximal” cycle length that vectors can achieve within an LCA. We create the designation of *maximal vector* to represent this.

Definition 6.1. For an LCA $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$, a *maximal vector* is any vector $\vec{v} \in \mathbb{Z}_N^L$ whose cycle length under \mathbf{A} equals the cycle length of \mathbf{A} .

From proposition 3 in “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]), it's known that for any LCA with a prime modulus, a maximal vector exists. Our interest in this chapter will be proving that a maximal vector always exists for LCAs with a prime-power modulus.

We may wonder whether the Primary Decomposition Theorem or the Minimal Polynomial Theorem will come in handy when working with maximal vectors. After all, the proof that maximal vectors always exist in LCAs with prime moduli relies heavily on the Primary Decomposition Theorem, and the Minimal Polynomial Theorem is a statement specifically about the multiplicative orders of vectors and matrices; they both seem like good tools to employ. However, there's another theorem pertaining to the cycle spaces of vectors that happens to be the most useful for our purposes: the Cyclic Decomposition Theorem.

Theorem 7 (Cyclic Decomposition Theorem). *For p prime, let $V \subseteq \mathbb{Z}_p^L$ be an invariant vector space under \mathbf{A} , where $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$ is a matrix. There exist nonzero vectors \vec{v}_1 to \vec{v}_r with*

Chapter 6. The Existence of Maximal Vectors

minimal annihilating polynomials $m_1(x)$ to $m_r(x)$ in $\mathbb{Z}_p[X]$, respectively, such that

$$V = \bigoplus_{i=1}^r \mathcal{S}_{\vec{v}_i},$$

where $m_{i+1}(x) \mid m_i(x)$ for $1 \leq i < r$, and $m_1(x)$ is the minimal polynomial of \mathbf{A} . As well, the number r and the polynomials $m_1(x)$ to $m_r(x)$ are uniquely determined by the conditions of the theorem.¹

At a high level, the Cyclic Decomposition Theorem allows us to break up an LCA's configuration space into a direct sum of cycle spaces, one of which has a generating vector which is a maximal vector. Like the Primary Decomposition Theorem and the Minimal Polynomial Theorem, the Cyclic Decomposition Theorem only applies directly to LCAs with a prime modulus, so we can't immediately apply it to LCAs with a prime-power modulus to prove the existence of maximal vectors. However, we *can* use it to show a relatively surprising result in the prime case which allows us to prove what we want to prove.

First, though, we'll explicitly state an observation about maximal vectors that will make our following arguments clearer.

Proposition 6.1. If $\vec{v} \in \mathbb{Z}_N^L$ is a maximal vector under the invertible matrix $\mathbf{A} \in \mathbb{Z}_N^{L \times L}$, then $\mathbf{A}^n \vec{v}$ is also a maximal vector.

Proof. Let ω be the cycle length of \vec{v} under \mathbf{A} . Assume $\mathbf{A}^n \vec{v}$ is not a maximal vector. Then there exists some number $\alpha < \omega$ such that

$$\mathbf{A}^\alpha (\mathbf{A}^n \vec{v}) \equiv \mathbf{A}^n (\mathbf{A}^\alpha \vec{v}) \equiv \mathbf{A}^n \vec{v} \pmod{N}.$$

Because \mathbf{A}^{-1} exists, this means

$$\mathbf{A}^\alpha \vec{v} \equiv \vec{v} \pmod{N}.$$

This contradicts the fact that \vec{v} is a maximal vector. So, if \vec{v} is a maximal vector, $\mathbf{A}^n \vec{v}$ must also be a maximal vector. **QED**

¹This theorem is adapted from “Linear algebra” (Hoffman and Kunze [2]). Our version has been altered to better fit within the context of this thesis.

Proposition 6.1 is fairly straightforward. If a vector is a maximal vector, then all of its iterates are maximal, too, as they're all in the same cycle; it'll take a maximal vector's iterate the same number of iterations to iterate back to itself. Referring back to Figure 2.2, this fact should be apparent.

Now, let's use the Cyclic Decomposition Theorem to show that a particularly nice basis must exist for the configuration spaces of LCAs with prime moduli.

Proposition 6.2. Given an invertible matrix $\mathbf{A} \in \mathbb{Z}_p^{L \times L}$, there exists a basis for \mathbb{Z}_p^L of maximal vectors under the matrix \mathbf{A} .

Proof. By the Cyclic Decomposition Theorem, there exist vectors \vec{v}_1 to \vec{v}_r such that

$$\mathbb{Z}_p^L = \bigoplus_{i=1}^r \mathcal{S}_{\vec{v}_i}. \quad (6.1)$$

Let $m_i(x)$ be the minimal annihilating polynomial of \vec{v}_i . Proposition 3.17 tells us that $\dim(\mathcal{S}_{\vec{v}_i}) = \deg(m_i(x))$. Thus, we have that

$$\dim(\mathbb{Z}_p^L) = \sum_{i=1}^r \deg(m_i(x))$$

since the Cyclic Decomposition Theorem gives us a *direct* sum for our vector space.

Also by the Cyclic Decomposition Theorem, we know that $m_1(x)$ is the minimal polynomial for \mathbf{A} . By the Minimal Polynomial Theorem, then, \vec{v}_1 must be a maximal vector. Thus, the set $B = \{\vec{v}_1, \mathbf{A}\vec{v}_1, \mathbf{A}^2\vec{v}_1, \dots, \mathbf{A}^{\deg(m_1(x))-1}\vec{v}_1\}$ is a linearly-independent set of $\deg(m_1(x))$ maximal vectors by Propositions 6.1 and 3.17.

Next, for $1 < i \leq r$, define the set X_i as

$$X_i = \bigcup_{j=0}^{\deg(m_i(x))-1} \{\vec{v}_1 + \mathbf{A}^j \vec{v}_i\}.$$

Each X_i forms a linearly-independent set. To see why, assume otherwise. Then, there exists

Chapter 6. The Existence of Maximal Vectors

constants a_0 to $a_{\deg(m_i(x))-1}$ in \mathbb{Z}_p , at least one of which is nonzero, such that

$$\sum_{j=0}^{\deg(m_i(x))-1} a_j (\vec{v}_1 + \mathbf{A}^j \vec{v}_i) \equiv \vec{0} \pmod{p}.$$

This sum can be broken into two vectors:

$$\left(\sum_{j=0}^{\deg(m_i(x))-1} a_j \vec{v}_1 \right) + \left(\sum_{j=0}^{\deg(m_i(x))-1} a_j \mathbf{A}^j \vec{v}_i \right) \equiv \vec{0} \pmod{p}.$$

The first vector is in $\mathcal{S}_{\vec{v}_1}$, while the second vector is in $\mathcal{S}_{\vec{v}_i}$. Equation (6.1) tells us that $\mathcal{S}_{\vec{v}_1} \cap \mathcal{S}_{\vec{v}_i} = \{\vec{0}\}$ for $1 < i \leq r$, and so both vectors must be zero for the sum to equal zero. This means

$$\sum_{j=0}^{\deg(m_i(x))-1} a_j \mathbf{A}^j \vec{v}_i \equiv \vec{0} \pmod{p}.$$

The vectors \vec{v}_i to $\mathbf{A}^{\deg(m_i(x))-1} \vec{v}_i$ are linearly independent (as they form a basis for $\mathcal{S}_{\vec{v}_i}$), and so the only way this sum can equal zero is if the constants are zero. This contradicts the fact that at least one of our constants a_0 to $a_{\deg(m_i(x))-1}$ is nonzero. Thus, our assumption that X_i is a linearly-dependent set must be false.

Furthermore, because $\mathcal{S}_{\vec{v}_1} \cap \mathcal{S}_{\vec{v}_i} = \{\vec{0}\}$ for $1 < i \leq r$, and because the cycle length of any vector in \mathbb{Z}_p^L under \mathbf{A} must divide the cycle length of \vec{v}_1 , proposition 2 in “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]) guarantees that every vector in X_i is a maximal vector.

Now, consider the set

$$\mathcal{B} = B \cup X_2 \cup X_3 \cup \cdots \cup X_r.$$

We can show that \mathcal{B} is linearly independent. To see why, assume otherwise. Then, there exist constants $a_{1,0}$ to $a_{r,\deg(m_r(x))-1}$, at least one of which is nonzero, such that

$$\left(\sum_{j=0}^{\deg(m_1(x))-1} a_{1,j} \mathbf{A}^j \vec{v}_1 \right) + \left(\sum_{i=2}^r \sum_{j=0}^{\deg(m_i(x))-1} a_{i,j} (\vec{v}_1 + \mathbf{A}^j \vec{v}_i) \right) \equiv \vec{0} \pmod{p}. \quad (6.2)$$

To ease with notation, define $\vec{\sigma}_i$ as

$$\vec{\sigma}_i = \sum_{j=0}^{\deg(m_i(x))-1} a_{i,j} \mathbf{A}^j \vec{v}_i.$$

Then, we can split the left side of Equation (6.2) into r different vectors:

$$\left(\vec{\sigma}_1 + \sum_{i=2}^r \sum_{j=0}^{\deg(m_i(x))-1} a_{i,j} \vec{v}_1 \right) + (\vec{\sigma}_2) + (\vec{\sigma}_3) + \cdots + (\vec{\sigma}_r) \equiv \vec{0} \pmod{p}. \quad (6.3)$$

The k -th vector in this sum will be an element of $\mathcal{S}_{\vec{v}_k}$. From Equation (6.1), we know the cycle spaces $\mathcal{S}_{\vec{v}_1}$ to $\mathcal{S}_{\vec{v}_r}$ have pairwise intersections of $\{\vec{0}\}$, and so the only way the above sum can equal zero is if each of the r vectors are zero. Then, for each $1 < i \leq r$, we must have that

$$\vec{\sigma}_i \equiv \sum_{j=0}^{\deg(m_i(x))-1} a_{i,j} \mathbf{A}^j \vec{v}_i \equiv \vec{0} \pmod{p}.$$

We proved above that the only way this sum can equal zero is if the relevant constants are zero. Thus, Equation (6.3) reduces to

$$\vec{\sigma}_1 \equiv \sum_{j=0}^{\deg(m_1(x))-1} a_{1,j} \mathbf{A}^j \vec{v}_1 \equiv \vec{0} \pmod{p}.$$

Once again, this sum can only equal zero if the relevant constants are zero. However, this causes a contradiction since, at this point, *all* of our constants $a_{1,0}$ to $a_{r,\deg(m_r(x))-1}$ are zero, even though one of them must be nonzero. Thus, our assumption that \mathcal{B} is a linearly-dependent set must be false. So, \mathcal{B} is a linearly-independent set.

By construction, the set \mathcal{B} has $\sum_{i=1}^r \deg(m_i(x))$ elements, which is the dimension of \mathbb{Z}_p^L . Thus, \mathcal{B} must form a basis for \mathbb{Z}_p^L . As well, each vector in \mathcal{B} is a maximal vector by construction. **QED**

Proposition 6.2 is an example of how powerful the Cyclic Decomposition Theorem can be. Much like the Primary Decomposition Theorem, the fact that we get a *direct* sum of cycle spaces for our configuration space allows us to more easily understand how vectors

Chapter 6. The Existence of Maximal Vectors

between the cycle spaces interact with each other. This, in turn, grants us the ability to construct vectors with particular properties—in our case, maximal vectors that are linearly independent.

With Proposition 6.2, we can use a few of our results from Section 3.3 to prove that maximal vectors must always exist for LCAs with prime-power moduli.

Theorem 8. *Given an invertible matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$, a maximal vector exists within the LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ for odd prime powers p^k .*

Proof. Since \mathbf{A} is invertible modulo p^k , its determinant must be invertible modulo p^k . By Proposition 3.12, the determinant of \mathbf{A} must also be invertible modulo p , and so \mathbf{A} must be invertible modulo p .

Because \mathbf{A} is invertible modulo p , Proposition 6.2 guarantees that a basis of maximal vectors under \mathbf{A} exists for \mathbb{Z}_p^L modulo p . We'll call this basis \mathcal{B} .

Now, assume that, modulo p , ω is the multiplicative order of \mathbf{A} . Because each $\vec{b} \in \mathcal{B}$ is a maximal vector, we know ω is also the multiplicative order of \vec{b} modulo p . By Proposition 3.15, there are two possibilities for \mathbf{A} :

Case 1: For all prime-power moduli p^ℓ for positive integers ℓ , the multiplicative order of \mathbf{A} modulo p^ℓ will be ω . In this case, the maximal vector under \mathbf{A} modulo p guaranteed by proposition 3 of “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]) will remain a maximal vector under all prime-power moduli p^ℓ (as the multiplicative order of a vector cannot decrease as the power of a prime-power modulus increases—see Example 2.1). Thus, a maximal vector will exist for all prime-power moduli of the form p^ℓ .

Case 2: For some prime-power moduli p^ℓ , the multiplicative order of \mathbf{A} increases from ω to $p\omega$. In this case, we have that

$$\mathbf{A}^\omega \equiv \mathbf{I} + p^{\ell-1}\mathbf{B} \pmod{p^\ell}$$

for some nonzero matrix $\mathbf{B} \in \mathbb{Z}_p^{L \times L}$. By Proposition 3.16, the multiplicative order of any $\vec{b} \in \mathcal{B}$ modulo p^ℓ is either ω or $p\omega$. Calculating $\mathbf{A}^\omega \vec{b}$, we see

$$\mathbf{A}^\omega \vec{b} \equiv \vec{b} + p^{\ell-1}\mathbf{B}\vec{b} \pmod{p^\ell}.$$

It's possible that $\vec{b} \in \ker(p^{\ell-1}\mathbf{B})$, in which case the multiplicative order of \vec{b} is ω . However, by Proposition 3.13, \mathcal{B} is still a linearly-independent set modulo p^ℓ , and so it forms a basis for $\mathbb{Z}_{p^\ell}^L$. This means that at least one vector $\vec{\beta} \in \mathcal{B}$ must not be in the kernel of $p^{\ell-1}\mathbf{B}$. Otherwise, if no such vector existed, then $p^{\ell-1}\mathbf{B}$ would annihilate all the vectors of our basis, meaning it would necessarily be the zero matrix. We know $p^{\ell-1}\mathbf{B} \neq \mathbf{0}$ since $\mathbf{A}^\omega \neq \mathbf{I}$ by assumption, so such a vector $\vec{\beta}$ must exist.

Therefore,

$$\begin{aligned} \mathbf{A}^\omega \vec{\beta} &\equiv \vec{\beta} + p^{\ell-1}\mathbf{B}\vec{\beta} \pmod{p^\ell} \\ &\neq \vec{\beta}. \end{aligned}$$

The multiplicative order of $\vec{\beta}$ cannot possibly be ω from the above equation. Then, by Proposition 3.16, it must be $p\omega$, meaning it is a maximal vector modulo p^ℓ . As well, by Proposition 3.14, the matrix \mathbf{B} in the expression for $\mathbf{A}^\omega \pmod{p^\ell}$ will be the same matrix \mathbf{B} in the expressions for $\mathbf{A}^{p^x\omega} \pmod{p^{\ell+x}}$ for positive integers x , so all higher powers of the prime as the modulus will have that

$$\begin{aligned} \mathbf{A}^{p^x\omega} \vec{\beta} &\equiv \vec{\beta} + p^{\ell+x-1}\mathbf{B}\vec{\beta} \pmod{p^{\ell+x}} \\ &\neq \vec{\beta}. \end{aligned}$$

So long as $\vec{\beta}$ is a maximal vector modulo $p^{\ell+x-1}$, the above equation shows that it must also be a maximal vector modulo $p^{\ell+x}$ by applying Proposition 3.16 (since Proposition 3.15 guarantees that $p^{x+1}\omega$ is the multiplicative order of \mathbf{A} modulo $p^{\ell+x}$). Therefore, by applying the above reasoning up to modulo p^k , $\vec{\beta}$ must also be a maximal vector modulo p^k . **QED**

Using Theorem 6, we can immediately extend Theorem 8 to *any* LCA, not just those with invertible update matrices. For an LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$, the trick is to let $\hat{\mathbf{A}}$ be the restriction of \mathbf{A} to $\mathcal{K}_{\mathbb{Z}_{p^k}^L}(\mathbf{A})$. By Theorem 6, $\hat{\mathbf{A}}^{-1}$ exists, and so Theorem 8 can be applied to the core of the LCA specifically.

Corollary 6.1. Given an arbitrary matrix $\mathbf{A} \in \mathbb{Z}_{p^k}^{L \times L}$, a maximal vector exists within the LCA $(\mathbb{Z}_{p^k}, \mathbb{Z}_{p^k}^L, \mathbf{A})$ for odd prime powers p^k .

Chapter 6. The Existence of Maximal Vectors

Notice that the crux of our argument in the proof for Theorem 8 isn't so much an intrinsic property of LCAs with prime-power moduli, but rather an intrinsic property of LCAs with *prime* moduli (Proposition 6.2). As has been the case throughout this thesis, the key to understanding the prime-power case is to first understand the prime case, then use the relations between the prime and prime-power case to make conclusions.

In any case, Theorem 8 shows that, like the prime case, LCAs with prime-power moduli will always have at least one vector whose multiplicative order is as big as possible (i.e. equal to the multiplicative order of the update matrix) and is a multiple of every other vector's multiplicative order. This extends proposition 3 from “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]) to include LCAs with prime-power moduli rather than just prime moduli. Despite the inherent differences between the prime and prime-power case, this theorem establishes yet another dynamical similarity between them.

Chapter 7

Conclusion

In this thesis, we built on numerous previous results regarding the dynamics of finite linear cellular automata (LCAs), mainly those contained in “Linear cellular automata” (Patterson [5]) and “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]). Via Theorem 3, the Minimal Polynomial Theorem, a tool for algebraically determining cycle lengths and transient lengths of vectors in a prime-modulus LCA, was extended to certain LCAs with prime-powered moduli, allowing for faster computation of these values in a greater number of cases.

Theorem 5 imposed a certain structure onto the submodules of LCAs with prime-powered moduli, stating that a spanning set of “dimensionally-independent vectors” always exists (much like a basis always exists for a subspace). Using this theorem, we were able to show that the cores of LCAs—that is, the largest set of vectors where the update matrix acts as an invertible transformation—must always have a basis for prime-power moduli via Theorem 6. This result helps drastically simplify computations involving the core of an LCA since having a basis allows us to represent the action of the LCA (the update matrix) as another matrix when restricted to the core, and matrices are easier to work with than more general operators.

We also showed that a maximal vector always exists for LCAs with invertible update matrices modulo a prime power (via Theorem 8), which extended proposition 3 from “Dynamics of finite linear cellular automata over \mathbb{Z}_N ” (Mendivil and Patterson [4]) and gave a better understanding as to the structure of possible cycle lengths within an LCA (as all cycle

Chapter 7. Conclusion

lengths must divide the cycle length of some “maximal” vector).

One topic which could be expanded upon in future research is the cycle converting matrix (CCM). The connection between CCMs and the annihilating polynomials of update matrices suggests some deep structure to these matrices that could be used to gain further insights into the dynamics of arbitrary LCAs. Already, via some of the propositions provided in this thesis, we can quickly deduce information regarding the possible cycle lengths of vectors within a given LCA. However, based on the observed behaviour of CCMs, it seems that, in certain cases, CCMs can give far more information about an LCA than what our propositions guarantee (such as the exact number of vectors in an LCA with a given cycle length). We were unable to characterise these specific cases, and so this wasn’t discussed in great detail throughout the thesis. Future research could be dedicated to describing the conditions under which CCMs give more information about an LCA. As well, our discussion of CCMs was mainly limited to the prime modulus case, and so the behaviour of CCMs in the prime-power case is another potential avenue for future work.

Similar to CCMs, our discussion of cycle spaces throughout this thesis was mainly focused on the prime modulus case. In the prime-power case, cycle spaces are no longer subspaces, and so their behaviour is more difficult to describe. However, with the idea of dimensionally-independent vectors, it’s possible that their behaviour can still be described in some capacity. Our research simply did not go in this direction, and so this is a potentially-fruitle direction in which future work can go.

Something of interest we noted during our research was how the different iterates of vectors in a prime-powered LCA, say modulo p^k , “mapped” between different LCAs with different prime-powered moduli, say modulo p^ℓ . As an example, how would the iterates of a vector \vec{v} under a matrix \mathbf{A} modulo 5^2 relate to the iterates of \vec{v} under \mathbf{A} modulo 5^3 ? Via the concept of lifts and embed vectors, there’s clearly a relationship between the two different moduli (as the iterates in the higher-power modulus must reduce down to the iterates in the lower-power modulus via modular reduction), but this exact relationship is difficult to explain generally. In some cases, the iterates modulo the higher modulus will form a cycle much longer than the corresponding iterates modulo the smaller modulus, whereas other times the higher-modulus iterates will form cycles of the same length, but there will be multiple of them. We were unable to find an exact characterisation of this behaviour, though it seems

likely that a particular structure to this behaviour does exist, probably related to the LCA's update matrix.

Chapter 7. Conclusion

Appendix A

The Chinese Remainder Theorem

Throughout this thesis, we've concerned ourselves primarily with LCAs of prime and prime-power moduli. As mentioned briefly in Chapter 2, focusing on these two cases is sufficient to understand the behaviour of *all* LCAs, as whenever the modulus is a general composite number (i.e. a product of one or more primes or prime powers), we can break it up into its coprime, prime/prime-power factors and use the corresponding LCAs with each coprime factor as its modulus to understand the original LCA. The tool that allows us to do this is the Chinese Remainder Theorem.

There are many different ways to state the Chinese Remainder Theorem, each one helpful in its own area of math. For us, we'll use the statement most often used in the realm of abstract algebra.

Theorem 9 (Chinese Remainder Theorem). *For pairwise coprime integers n_1 to n_r , let*

$$\mathcal{Z} = \mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \cdots \times \mathbb{Z}_{n_r}$$

be the direct product of the rings \mathbb{Z}_{n_1} to \mathbb{Z}_{n_r} , where $N = n_1 \times n_2 \times \cdots \times n_r$. The mapping

$$\rho : \mathbb{Z}_N \rightarrow \mathcal{Z}, \quad \rho(x) = (x \bmod n_1, x \bmod n_2, \cdots, x \bmod n_r)$$

defines a ring isomorphism

$$\mathbb{Z}_N \cong \mathcal{Z}.$$

Appendix A. The Chinese Remainder Theorem

The statement of the Chinese Remainder Theorem is quite abstract, but its interpretation is straightforward. If we're given the ring \mathbb{Z}_N for some $N = n_1 n_2 \cdots n_r$, where n_1 to n_r are pairwise coprime integers, then the behaviour of any element in \mathbb{Z}_N can be understood by looking at the behaviour of a set of corresponding elements in \mathbb{Z}_{n_1} to \mathbb{Z}_{n_r} . This allows us to completely avoid computation in \mathbb{Z}_N and work only in the rings \mathbb{Z}_{n_1} to \mathbb{Z}_{n_r} !

For us, the Chinese Remainder Theorem allows us to take LCAs of the form $(\mathbb{Z}_N, \mathbb{Z}_N^L, \mathbf{A})$ for $N = p_1^{n_1} p_2^{n_2} \cdots p_r^{n_r}$ (where each p_i is a unique prime and each n_i is a positive integer) and understand the behaviour of their vectors (and update matrix, if we wish) by instead analysing the related LCAs $(\mathbb{Z}_{p_1^{n_1}}, \mathbb{Z}_{p_1^{n_1}}^L, \mathbf{A})$ to $(\mathbb{Z}_{p_r^{n_r}}, \mathbb{Z}_{p_r^{n_r}}^L, \mathbf{A})$. In this way, we can focus our attention solely on understanding the prime and prime-power modulus case for LCAs without sacrificing our grasp on the general composite modulus case.

As a comparison, recall the ϕ bijection introduced in Chapter 2. This mapping described a connection between vectors in $\mathbb{Z}_{p^k}^L$ and vectors in $p\mathbb{Z}_{p^{k+1}}^L$ (for p^k a prime power). For any vector in $\mathbb{Z}_{p^k}^L$, there was a corresponding vector in $p\mathbb{Z}_{p^{k+1}}^L$ whose behaviour exactly matched it. The ρ mapping in the Chinese Remainder Theorem behaves similarly. For any vector in \mathbb{Z}_N^L , there's a corresponding *set* of vectors in $\mathbb{Z}_{p_1^{n_1}}^L$ to $\mathbb{Z}_{p_r^{n_r}}^L$ (using the notation from above) that, together, mimic the behaviour of the original vector. While the ρ isomorphism isn't as simple as a multiplication by p like with ϕ , it can be thought of in the same way: it allows us to take a vector and transform it into a different form which behaves in exactly the same way, even if it looks vastly different.

Let's look at an example of how the Chinese Remainder Theorem is used to understand the behaviour of vectors within a composite modulus LCA. For this, we'll use the LCA $\mathcal{A} = (\mathbb{Z}_{60}, \mathbb{Z}_{60}^2, [\begin{smallmatrix} 58 & 9 \\ 51 & 18 \end{smallmatrix}])$ and the vector $\vec{v} \equiv [\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}]$. Rather than explicitly compute the iterates of \vec{v} within \mathcal{A} , we'll look at a set of corresponding LCAs and use the Chinese Remainder Theorem to construct the iterates.

The modulus for \mathcal{A} is 60, and we know that $60 = 2^2 \cdot 3 \cdot 5$. So, let's consider the three

LCAs

$$\begin{aligned}\mathcal{A}_{2^2} &= (\mathbb{Z}_{2^2}, \mathbb{Z}_{2^2}^2, [\begin{smallmatrix} 58 & 9 \\ 51 & 18 \end{smallmatrix}]) , \\ \mathcal{A}_3 &= (\mathbb{Z}_3, \mathbb{Z}_3^2, [\begin{smallmatrix} 58 & 9 \\ 51 & 18 \end{smallmatrix}]) , \\ \mathcal{A}_5 &= (\mathbb{Z}_5, \mathbb{Z}_5^2, [\begin{smallmatrix} 58 & 9 \\ 51 & 18 \end{smallmatrix}]) .\end{aligned}$$

For each of our three new LCAs, the update matrix can be reduced modulo the respective modulus since all computations within the LCAs will be reduced, too. So, we can rewrite our three LCAs as

$$\begin{aligned}\mathcal{A}_{2^2} &= (\mathbb{Z}_{2^2}, \mathbb{Z}_{2^2}^2, [\begin{smallmatrix} 2 & 1 \\ 3 & 2 \end{smallmatrix}]) , \\ \mathcal{A}_3 &= (\mathbb{Z}_3, \mathbb{Z}_3^2, [\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix}]) , \\ \mathcal{A}_5 &= (\mathbb{Z}_5, \mathbb{Z}_5^2, [\begin{smallmatrix} 3 & 4 \\ 1 & 3 \end{smallmatrix}]) .\end{aligned}$$

Already, we can see that computation within these three LCAs will be much simpler than computation within \mathcal{A} . For LCAs with extremely large composite moduli, computation in these corresponding simpler LCAs may be the only way to reasonably compute anything (due to limits on the size of integers within a computer program, for instance).

Now, we'll compute the sequence of iterates for \vec{v} within these three simpler LCAs.

For \mathcal{A}_{2^2} : $[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}], [\begin{smallmatrix} 3 \\ 1 \end{smallmatrix}], [\begin{smallmatrix} 3 \\ 3 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 3 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}], [\begin{smallmatrix} 3 \\ 1 \end{smallmatrix}], \dots$

For \mathcal{A}_3 : $[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}], [\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}], \dots$

For \mathcal{A}_5 : $[\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}], [\begin{smallmatrix} 2 \\ 4 \end{smallmatrix}], [\begin{smallmatrix} 2 \\ 4 \end{smallmatrix}], [\begin{smallmatrix} 2 \\ 4 \end{smallmatrix}], [\begin{smallmatrix} 2 \\ 4 \end{smallmatrix}], [\begin{smallmatrix} 2 \\ 4 \end{smallmatrix}], \dots$

In \mathcal{A}_3 and \mathcal{A}_5 , \vec{v} ends up iterating to a fixed point, while in \mathcal{A}_{2^2} , \vec{v} has a cycle length of 4 and a transient length of 0. How can we use this information to construct the iterates of \vec{v} within \mathcal{A} ? Perhaps unsurprisingly, the key is to use the Chinese Remainder Theorem. For each group of iterates in our list above (i.e. each column of iterates), the Chinese Remainder Theorem guarantees an isomorphism between them and a vector in \mathbb{Z}_{60}^2 , the configuration space of \mathcal{A} . Because our groups of iterates correspond to iterates of \vec{v} within the simpler

Appendix A. The Chinese Remainder Theorem

LCAs, the vectors we obtain via the isomorphism will be the iterates of \vec{v} within \mathcal{A} .

For any group of iterates in our list above (i.e. any column), if \vec{a}_{2^2} is the vector from \mathcal{A}_{2^2} , \vec{a}_3 the one from \mathcal{A}_3 , and \vec{a}_5 the one from \mathcal{A}_5 , then to get our corresponding vector \vec{x} from \mathcal{A} , we must solve the following system of congruences:

$$\begin{aligned}\vec{x} &\equiv \vec{a}_{2^2} \pmod{2^2}, \\ \vec{x} &\equiv \vec{a}_3 \pmod{3}, \\ \vec{x} &\equiv \vec{a}_5 \pmod{5}.\end{aligned}$$

A systematic method for solving such systems is given in the proof of the Chinese Remainder Theorem in “A Friendly Introduction to Number Theory” (Silverman [6]) (chapter 11, pages 79-80). We’ll skip over the actual calculations and simply list the solutions for each column in the list of iterates above.¹

$$\text{For } \mathcal{A}: \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 7 \\ 9 \end{bmatrix}, \begin{bmatrix} 7 \\ 39 \end{bmatrix}, \begin{bmatrix} 37 \\ 39 \end{bmatrix}, \begin{bmatrix} 37 \\ 9 \end{bmatrix}, \begin{bmatrix} 7 \\ 9 \end{bmatrix}, \dots$$

These are the iterates of \vec{v} within \mathcal{A} . For such a small example, all this work was hardly necessary, but it demonstrates how computation within LCAs with general composite moduli is entirely unnecessary. By factoring the modulus into prime and prime-power factors, we can work solely with LCAs with prime and prime-power moduli. Thus, results regarding LCAs with prime and prime-power moduli have been prioritized in this thesis as the corresponding results for the general composite case can always be extrapolated using the Chinese Remainder Theorem.

¹As an aside, this example perfectly demonstrates what it means for two things to be isomorphic. On one hand, we have the iterates of \vec{v} within \mathcal{A} . On the other hand, we have groups of iterates of \vec{v} within simpler LCAs. The Chinese Remainder Theorem guarantees that these two sides are isomorphic, and so their behaviours must mirror each other. In essence, the two sides are the same, just under different representations. The isomorphism describes how to translate between the two representations.

Bibliography

- [1] Rory Alisdair Dow. *Algebraic Methods for Finite Linear Cellular Automata*. PhD thesis, University of London, 1996.
- [2] Kenneth Hoffman and Ray Kunze. *Linear algebra*. Prentice Hall, Englewood Cliffs, New Jersey, 2nd edition, 1971.
- [3] Nathan Jacobson. *Lectures in abstract algebra: II. linear algebra*. Springer New York, New York, New York, 1st edition, 1953.
- [4] Franklin Mendivil and Donald Patterson. Dynamics of finite linear cellular automata over \mathbb{Z}_N . *Rocky Mountain Journal of Mathematics*, 42(2):695–709, 2012. doi: <https://doi.org/10.1216/RMJ-2012-42-2-695>.
- [5] Donald Patterson. Linear cellular automata. Bachelor’s thesis, Acadia University, 2008.
- [6] Joseph H. Silverman. *A Friendly Introduction to Number Theory*. Pearson India Education Services Pvt. Ltd, Uttar Pradesh, India, 4th edition, 2019.
- [7] Lawrence E. Spence Stephen H. Friedberg, Arnold J. Insel. *Linear algebra*. Prentice Hall, Englewood Cliffs, New Jersey 07632, 2nd edition, 1989.
- [8] Stephen Wolfram. *A New Kind of Science*. Wolfram Media, 2002. ISBN 1579550088. URL <https://www.wolframscience.com>.