Chapter 1

Introduction

A *discrete dynamical system* is a pair (X, g), where X is a set and g is a function $X \to X$. The motivation behind this definition is to think of a complicated system that evolves in discrete time steps (such as a neural network), with X being the set of all states which the system can assume, and $g(x)$ being the successor state of $x \in X$. For this reason, one calls X the *state space* and g the *(state) transition function* of (X, g). When studying a discrete dynamical system (X, g), one is naturally interested in the behavior of g under iteration (i.e., in the function iterates g^n for $n \in \mathbb{N}_0 = \{n \in \mathbb{Z} : n \geq 0\}$). See the monograph [48] for a general introduction to discrete dynamical systems, and [48, Chapter 7] in particular for some examples of practical applications of them.

When X is finite, one also calls (X, g) a *finite dynamical system*. Some important special cases with regard to applications are when $X = \mathbb{Z}/m\mathbb{Z}$ and g is a polynomial modulo m (which is used in Pollard's rho algorithm [59]), or when $X = \mathbb{F}_q^n$ (Cartesian power of the finite field $\mathbb{F}_q$), with a particular focus on $q = 2$ in the literature (see [36, 43, 44, 54, 74]). It should be noted that one may identify $\mathbb{F}_q^n$ with $\mathbb{F}_{q^n}$ by fixing an $\mathbb{F}_q$-basis in the latter, so there is in fact no loss of generality when assuming $n = 1$ (i.e., when only considering finite fields themselves as state spaces).

A simple yet remarkable fact when X is finite is that all points $x \in X$ are *pre-periodic under* g, i.e., there exist unique smallest integers $\mathrm{pperl}_g(x) \geq 0$ and $\mathrm{perl}_g(x) \geq 1$, called the *pre-period (length)* and *period (length)* of x under g, respectively, such that $g^{\mathrm{pperl}_g(x)+\mathrm{perl}_g(x)}(x) = g^{\mathrm{pperl}_g(x)}(x)$. A different terminology frequently used in the literature for pre-periodic points is eventually (or ultimately) periodic. In the case $\mathrm{pperl}_g(x) = 0$, one says that x is *periodic under* g (or *g-periodic*); periodic points are also known as purely periodic. The subset of X consisting of all g-periodic points is denoted by $\mathrm{per}(g)$. A point in X that is not g-periodic is called *transient under* g (or g-transient). Various stochastic parameters of random functions $X \to X$ that are of interest for the study of finite dynamical systems, such as the expected pre-period and period length of a point, were determined in [23].

An important means of visualizing a discrete dynamical system (X, g), especially when X is finite, is the so-called *functional graph of* g, denoted by Γ_g. This is the directed graph with vertex set X that has an arc (directed edge) $x \to g(x)$ for each $x \in X$, and no other arcs. It is straightforward to show that a directed graph Γ with vertex set X is a functional graph (i.e., is of the form Γ_g for some $g : X \to X$) if and only if each $x \in X$ has out-degree 1 in Γ.

Particularly for *finite* functional graphs Γ_g, one can give the following precise characterization of their shape: a *connected component of* Γ_g is the induced subgraph

of Γ_g on a subset of X that is the vertex set of a connected component of the underlying undirected graph of Γ_g. Each such connected component contains a single cycle of periodic points of g. Apart from those periodic points, the connected component consists precisely of those points which eventually map to the cycle after sufficiently many iterations of g – the *iterated pre-images (under g)* of points on the cycle. For each x on the cycle, the iterated pre-images y of x such that $x = y^{\mathrm{pperl}_g(y)}$ form a directed rooted tree, with root x, that has all of its arcs oriented toward the root. Henceforth, for simplicity, whenever we say "(directed) rooted tree", it means "directed rooted tree in which all arcs are oriented toward the root". Here is a picture to illustrate the situation:

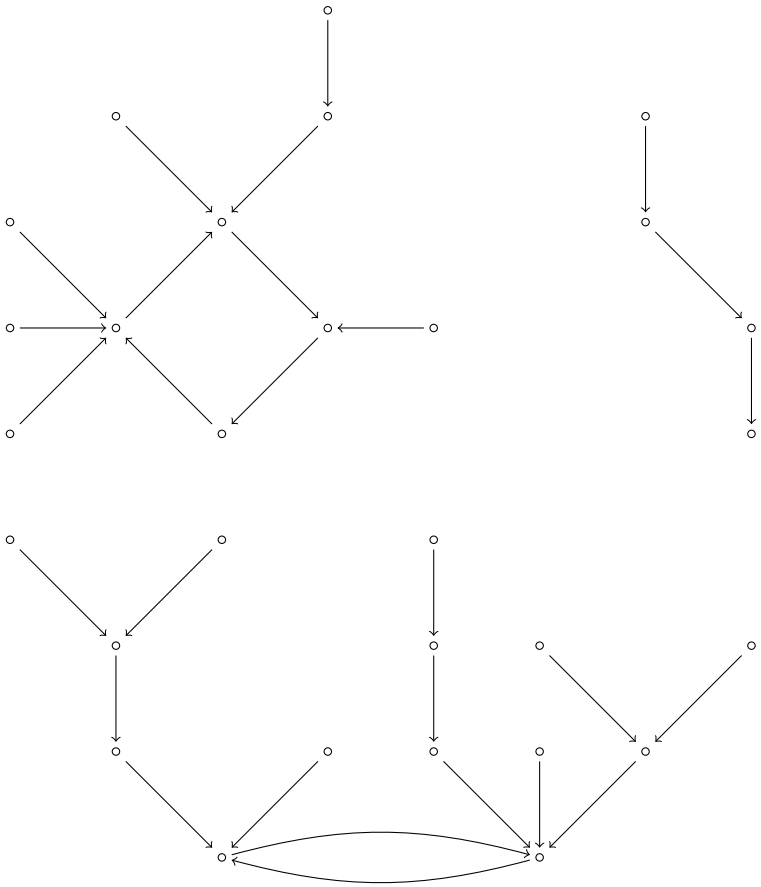

Conversely, each finite digraph of the shape described above is a functional graph, as it is readily verified that all vertices in it have out-degree 1. The study of finite dynamical systems may be understood as the study of finite functional graphs. In this context, it is also noteworthy that in case ψ_1 and ψ_2 are permutations of a finite set, we have $\Gamma_{\psi_1} \cong \Gamma_{\psi_2}$ if and only if ψ_1 and ψ_2 are of the same *cycle type*, i.e.,

they have the same number of cycles of each given length. Formally, the cycle type of a permutation ψ of X, denoted by $\mathrm{CT}(\psi)$, is defined as the unique monomial in $\mathbb{Q}[x_n : n \in \mathbb{N}^+]$, where $\mathbb{N}^+ = \{n \in \mathbb{Z} : n \geq 1\}$, in which the degree of each variable x_n is the number of ψ-cycles of length equal to n. For example, if $X = \{1, 2, \ldots, 9\}$ and $\psi = (1, 2, 3)(4, 5)(6, 7)(8)(9)$, then

$$\mathrm{CT}(\psi) = x_1^2 x_2^2 x_3.$$

Cycle types (and the related notion of cycle indices) are well studied in combinatorics, and studying isomorphism types of functional graphs may be seen as a natural generalization of this to arbitrary functions on finite sets.

Functional graphs of certain classes of functions on finite fields received considerable attention recently, see the papers [33, 50, 57, 61, 62, 71–73] and references therein. A survey of this topic can be found in [49]. Additionally, the papers [18, 58] do not deal explicitly with functional graphs, but with the iteration of functions on finite fields, and their results could be reformulated in terms of functional graphs. In this memoir, we contribute to this line of research by investigating functional graphs of so-called *generalized cyclotomic mappings* in the following sense.

Definition 1.1. Let q be a prime power, and let $d \mid q - 1$. A *generalized cyclotomic mapping of $\mathbb{F}_q$ of index d* is a function $f : \mathbb{F}_q \to \mathbb{F}_q$ with $f(0) = 0$ such that the restriction of f to each coset C_i of the unique index d subgroup of $\mathbb{F}_q^*$ agrees with a monomial function

$$x \mapsto a_i x^{r_i}.$$

More specifically, let ω be a primitive element of $\mathbb{F}_q$ (i.e., a generator of the cyclic multiplicative group $\mathbb{F}_q^*$), and let C be the index d subgroup of $\mathbb{F}_q^*$. The d cosets of C in $\mathbb{F}_q^*$ are of the form

$$C_i = \omega^i C$$

for $i = 0, 1, \ldots, d - 1$. The general form of an index d generalized cyclotomic mapping f of $\mathbb{F}_q$ is

$$f(x) = \begin{cases} 0, & \text{if } x = 0, \\ a_0 x^{r_0}, & \text{if } x \in C = C_0, \\ a_1 x^{r_1}, & \text{if } x \in C_1, \\ \vdots & \vdots \\ a_{d-1} x^{r_{d-1}}, & \text{if } x \in C_{d-1}, \end{cases} \tag{1.1}$$

where $a_i \in \mathbb{F}_q$ and $r_i \in \{0, 1, \ldots, q - 2\}$ for $i = 0, 1, \ldots, d - 1$. These functions are interesting because they generalize monomial mappings (which constitute the special

case $d = 1$) while still being relatively well controlled. From an abstract algebraic point of view, it is noteworthy that monomial functions

$$\mathbb{F}_q^* \to \mathbb{F}_q^*, \quad x \mapsto ax^r,$$

where $a \neq 0_{\mathbb{F}_q}$ necessarily, are *affine maps* of the multiplicative group $\mathbb{F}_q^*$, in the sense that they are compositions of a group endomorphism of $\mathbb{F}_q^*$ (viz., the power function $x \mapsto x^r$) with a (multiplicative) translation

$$x \mapsto ax$$

by a fixed group element a (see also Definition 2.1.15). Hence, generalized cyclotomic mappings in which all coefficients a_i from (1.1) are non-zero may be viewed as "coset-wise affine" functions, and we explore the idea of generalizing the methods and results from this memoir to other (possibly non-abelian) groups in Section 6.4. In this context, we also note that the celebrated Collatz function $g : \mathbb{Z} \to \mathbb{Z}$, given by the formula

$$g(x) = \begin{cases} x/2, & \text{if } x \in 2\mathbb{Z}, \\ 3x + 1, & \text{if } x \in 2\mathbb{Z} + 1, \end{cases}$$

is also a coset-wise affine function, of its respective domain of definition group $\mathbb{Z}$. The reason why we are able to develop a theory for understanding the behavior of generalized cyclotomic mappings under iteration in this memoir (while the analogous task for the Collatz function is wide open) is because generalized cyclotomic mappings preserve the associated partition of $\mathbb{F}_q$ into the cosets C_i and the singleton set $\{0_{\mathbb{F}_q}\}$ (and, relatedly, they form a semigroup under function composition) – see also the distinction between the two concepts introduced in Definition 6.4.1 (2,3).

We observe that a given generalized cyclotomic mapping of $\mathbb{F}_q$ may have several possible indices, and that *every* function $f : \mathbb{F}_q \to \mathbb{F}_q$ with $f(0) = 0$ is a generalized cyclotomic mapping of $\mathbb{F}_q$ of index $q - 1$, though the study of generalized cyclotomic mappings is mostly focused on small values of d. For $d = q - 1$, known methods of handling generalized cyclotomic mappings, such as [15, Algorithm 1], are essentially the trivial brute-force approaches. Apart from [15], generalized cyclotomic mappings were also studied in [80, 81, 87]. An important special case is when $r_i = r$ for all i; then one speaks of an *r-th order cyclotomic mapping of $\mathbb{F}_q$ of index d*, and those functions were studied, e.g., in [24, 56, 78, 79, 82].

Our goal in this memoir is to develop algorithms that answer fundamental questions concerning the structure of the functional graph Γ_f of a given index d generalized cyclotomic mapping f of $\mathbb{F}_q$, specified in the form (1.1). For example, let ω be any fixed primitive element of $\mathbb{F}_{256}$, and consider the following index 5 generalized

cyclotomic mapping f of $\mathbb{F}_{256}$.

$$f(x) = \begin{cases} 0, & \text{if } x = 0, \\ \omega^5 x^9, & \text{if } x \in C_0, \\ x^3, & \text{if } x \in C_1, \\ x^{17}, & \text{if } x \in C_2, \\ \omega^3 x^{34}, & \text{if } x \in C_3, \\ \omega^4 x^9, & \text{if } x \in C_4. \end{cases} \tag{1.2}$$

The functional graph Γ_f has 256 vertices, and one can understand its structure by drawing it, which we do below. In this drawing, a vertex labeled $n \in \{0, 1, \ldots, 254\}$ corresponds to the field element $\omega^n \in \mathbb{F}_q^*$ (in particular, the label 0 corresponds to the field element $\omega^0 = 1_{\mathbb{F}_{256}}$), whereas the vertex representing the field element 0 is labeled by $0_{\mathbb{F}_{256}}$. It turns out that Γ_f has four connected components, and in one of them (the fourth one in our order of drawing), the rooted trees glued to three particular f-periodic points on the unique cycle in that connected component are relatively large and thus drawn separately; we mark those rooted trees with Δ_k for $k \in \{1, 2, 3\}$ in the schematic drawing of the corresponding connected component.

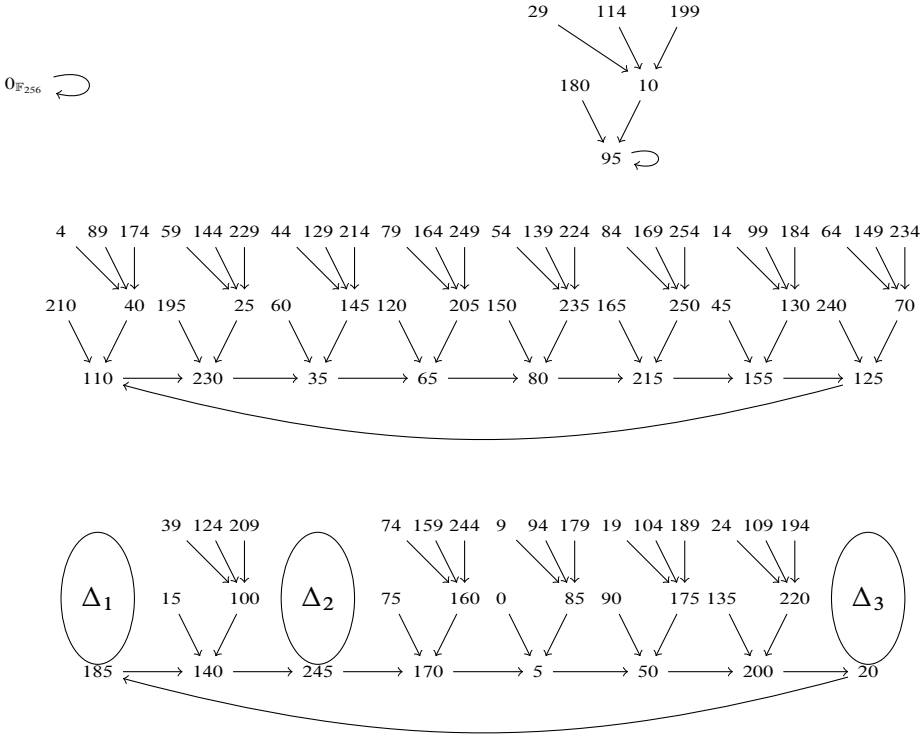

The rooted tree Δ_1 looks as follows.

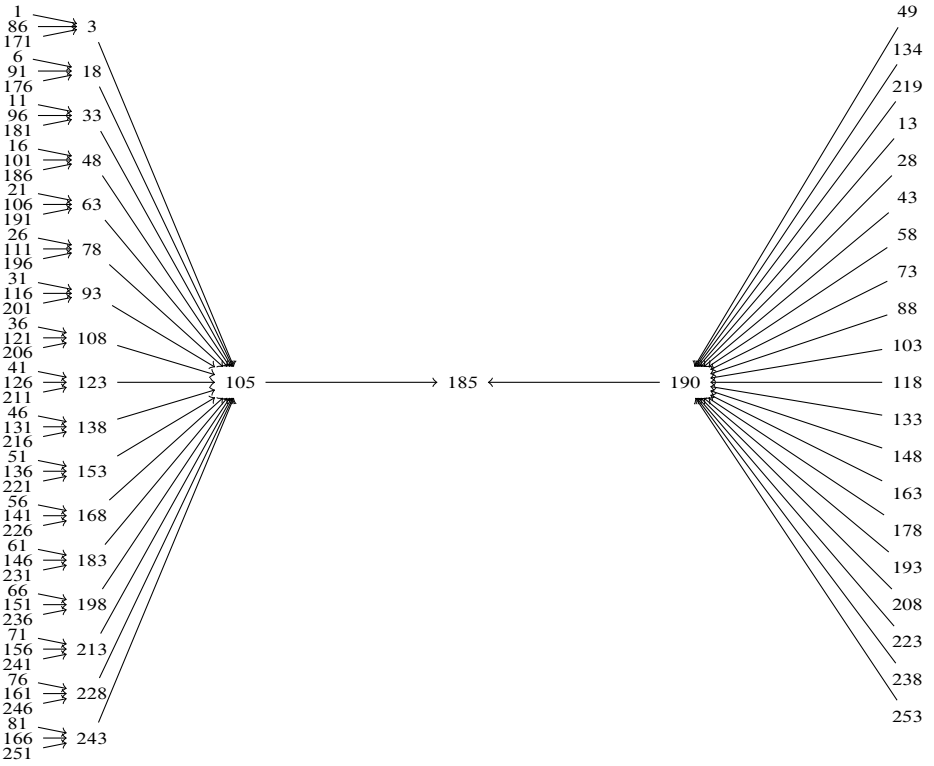

The rooted tree Δ_2 looks as follows.

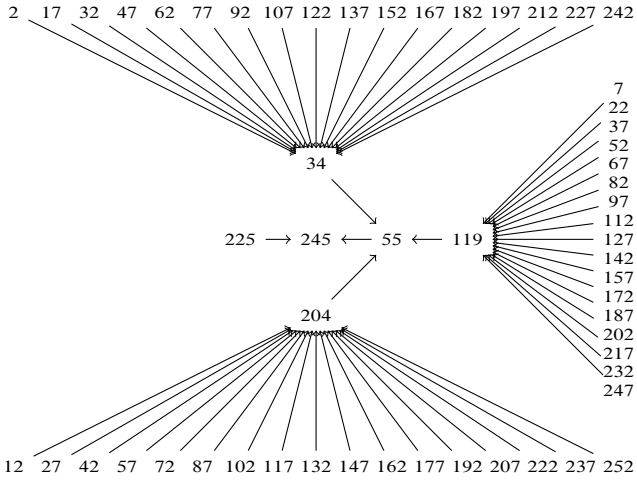

The rooted tree Δ_3 looks as follows.

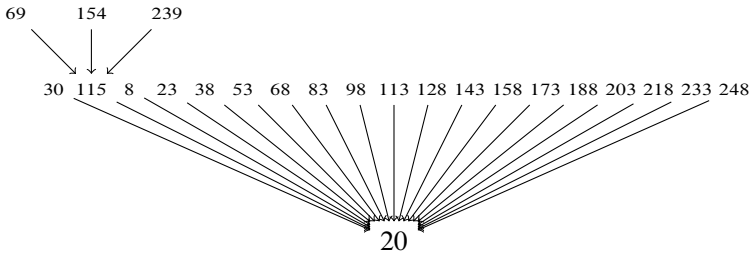

Of course, this approach of understanding Γ_f by drawing it becomes intractable for large values of q, as its complexity is at least linear in q (i.e., exponential in $\log q$). The aim of our algorithms is to obtain an understanding of the structure of Γ_f without needing to draw it vertex by vertex. A detailed complexity analysis of those algorithms, which we carry out in Chapter 5, shows that for asymptotically almost every finite field and fixed index d, our algorithms have implementations with polynomial complexity (in $\log q$) on quantum computers, and implementations with subexponential complexity on classical computers. In the remainder of this introduction, we discuss the main ideas underlying our algorithms. We also note that we revisit the example (1.2) in Section 4.2, where we derive the structure of its functional graph with our methods.

The first step in understanding the functional graph Γ_g of any function $g : X \to X$ is to obtain a suitable parametrization of the connected components of Γ_g. The following notion is helpful in that regard.

Definition 1.2. Let X be a finite set, and $g : X \to X$. A *cycle representatives and lengths list* (or *CRL-list* for short) *of* g is a (finite) set $\mathcal{L} \subseteq X \times \mathbb{N}^+$ with the following properties:

(1) The first entries of the ordered pairs in $\mathcal{L}$ form a system of representatives for the cycles of g on its periodic points.

(2) If $(r, l) \in \mathcal{L}$, then l is the cycle length of r under g.

Remark 1.3. When g is a function on a finite set, it is easy to determine the cycle type of the restriction $g_{|\mathrm{per}(g)}$ from any CRL-list $\mathcal{L}$ of g. Namely,

$$\mathrm{CT}(g_{|\mathrm{per}(g)}) = \prod_{(r,l)\in\mathcal{L}} x_l.$$

A CRL-list of g can thus be seen as a refinement of $\mathrm{CT}(g_{|\mathrm{per}(g)})$.

We recall from above that each connected component of Γ_g contains precisely one cycle of g on its periodic points. This means that a CRL-list of g also gives a parametrization of the connected components of g via representative vertices, along

with the basic information how long the cycle of each representative is. Let us give some more details on how to obtain $\mathfrak{L}$ when $X = \mathbb{F}_q$ and g is an index d generalized cyclotomic mapping f of $\mathbb{F}_q$.

We already introduced the notation $C_i = \omega^i C$ for $i \in \{0, 1, \ldots, d-1\}$ to denote the cosets of C in $\mathbb{F}_q^*$. Let us additionally set $C_d := \{0_{\mathbb{F}_q}\}$. Then the sets C_i for $i = 0, 1, \ldots, d$ form a partition of $\mathbb{F}_q$ that is preserved by f in the sense that f maps blocks of this partition to other such blocks (not necessarily surjectively). In other words, there is a unique function $\bar{f} : \{0, 1, \ldots, d\} \to \{0, 1, \ldots, d\}$, which we call *induced by* f, such that $f(C_i) \subseteq C_{\bar{f}(i)}$ for each $i = 0, 1, \ldots, d$. We note in particular that $\bar{f}(d) = d$, and that $\bar{f}^{-1}(\{d\}) = \{d\}$ unless at least one of the coefficients a_i in (1.1) is 0.

Setting $s := (q-1)/d = |C|$, we may view each coset $C_i = \omega^i C$ for $i = 0, 1, \ldots, d-1$ as a copy of the cyclic group $\mathbb{Z}/s\mathbb{Z}$ (with underlying set $\{0, 1, \ldots, s-1\}$ and modular addition as its group operation) via the bijection $\iota_i : \mathbb{Z}/s\mathbb{Z} \to C_i$, $x \mapsto \omega^{i+dx}$. As such, f may be viewed as a function that maps between copies of $\mathbb{Z}/s\mathbb{Z}$ (as well as a unique singleton block). More specifically, if $i \in \{0, 1, \ldots, d-1\}$ and $a_i \neq 0$, and if we write $a_i = \omega^{e_i}$, then we have $d \mid e_i + r_i i - \bar{f}(i)$ necessarily, and f maps $\omega^{i+dx} \in C_i$ to

$$a_i(\omega^{i+dx})^{r_i} = \omega^{e_i + r_i i + r_i dx} = \omega^{\bar{f}(i) + d\cdot(\frac{e_i + r_i i - \bar{f}(i)}{d} + r_i x)} \in C_{\bar{f}(i)}.$$

This means that under the identifications of C_i and $C_{\bar{f}(i)}$ with $\mathbb{Z}/s\mathbb{Z}$, the restriction of f to C_i corresponds to the affine function $A_i : x \mapsto r_i x + (e_i + r_i i - \bar{f}(i))/d$ of $\mathbb{Z}/s\mathbb{Z}$.

In summary, f consists essentially of affine functions mapping between copies of $\mathbb{Z}/s\mathbb{Z}$, though some copies of $\mathbb{Z}/s\mathbb{Z}$ may also be constantly mapped into $C_d = \{0\}$ by f, in case the corresponding coefficient a_i is 0. We note that if all a_i are non-zero, then the way $f_{|\mathbb{F}_q^*}$ preserves the partition of $\mathbb{F}_q^*$ into the cosets C_i for $i \in \{0, 1, \ldots, d-1\}$ is analogous to the way the elements of the imprimitive permutational wreath product $\mathrm{Sym}(C) \wr \mathrm{Sym}(d)$ (where $\mathrm{Sym}(X)$ and $\mathrm{Sym}(n)$ denote the symmetric group on the set X and on $\{0, 1, \ldots, n-1\}$, respectively) preserve this partition. In fact, the definition of an imprimitive permutational wreath product naturally extends to one of an *imprimitive wreath product of transformation semigroups* such that $f_{|\mathbb{F}_q^*}$ is an element of the imprimitive wreath product of C^C (the transformation semigroup of all functions $C \to C$) with $\{0, 1, \ldots, d-1\}^{\{0,1,\ldots,d-1\}}$, and $\bar{f}_{|\{0,1,\ldots,d-1\}}$ is the projection of f to $\{0, 1, \ldots, d-1\}^{\{0,1,\ldots,d-1\}}$. Wreath products of transformation semigroups have been studied before and play a central role in algebraic automata theory, though the notion used in that theory is the natural generalization of *primitive* permutational wreath products [34, pp. 55f.].

In any case, these ideas allow us to easily reduce the determination of a CRL-list $\mathfrak{L}$ of f to the determination of CRL-lists of affine functions on $\mathbb{Z}/s\mathbb{Z}$ – see

Section 3.1 for the details of this. CRL-lists of affine functions of finite cyclic groups are determined in Section 2.3.

The remainder of our algorithmic approach is concerned with understanding, for each given $(r, l) \in \mathcal{L}$, the isomorphism type of the connected component of Γ_f containing r. We recall from above that this connected component is essentially obtained by glueing certain directed rooted trees to the vertices on the cycle. Let us introduce the following precise notation.

Definition 1.4. Let Γ be a finite functional graph with vertex set X, and let g be the unique function $X \to X$ such that $\Gamma = \Gamma_g$. For each $x \in X$, we define $\mathrm{Tree}_\Gamma(x)$, the so-called *tree above x in Γ*, as follows.

(1) If x is g-transient, we define $\mathrm{Tree}_\Gamma(x)$ as the induced subgraph of Γ on the set

$$\{x\} \cup \{y \in \mathrm{V}(\Gamma) : y \neq x, \text{ and } g^k(y) = x \text{ for some } k = k(y) \geq 1\}.$$

(2) If x is g-periodic, we define $\mathrm{Tree}_\Gamma(x)$ as the induced subgraph of Γ on the set

$$\{x\} \cup \bigcup \{\mathrm{V}(\mathrm{Tree}_\Gamma(y)) : y \text{ is } g\text{-transient and } g(y) = x\},$$

with the convention that in case $g(x) = x$, the loop at x is deleted from $\mathrm{Tree}_\Gamma(x)$.

With this definition, $\mathrm{Tree}_{\Gamma_g}(x)$ is defined for all $x \in X = \mathrm{V}(\Gamma_g)$, and for periodic vertices x, those are the trees that need to be glued to the cycles of g in order to obtain the full connected components of Γ_g.

Necklaces are a well-studied concept in combinatorics. Let us consider vertex-labeled, directed graphs that consist of a single, directed cycle (let us call such a graph a *necklace graph*). Intuitively, one may think of the vertices as beads on a necklace (in the common sense of the word), and the vertex labels represent colors of those beads. An *isomorphism of vertex-labeled digraphs* is a digraph isomorphism preserving vertex labels, and a *necklace* is an isomorphism class of necklace graphs under isomorphism of vertex-labeled digraphs. If $\vec{\mathfrak{x}} = (\mathfrak{x}_0, \mathfrak{x}_1, \ldots, \mathfrak{x}_{L-1})$ is a length-L sequence with entries from a set $\mathfrak{X}$, then we denote by $[\vec{\mathfrak{x}}] = [\mathfrak{x}_0, \mathfrak{x}_1, \ldots, \mathfrak{x}_{L-1}]$ the orbit of $\vec{\mathfrak{x}}$ under the natural action of the cyclic group $\mathbb{Z}/L\mathbb{Z}$ on $\mathfrak{X}^L$. Hence, $[\vec{\mathfrak{x}}]$ consists of those length-L sequences over $\mathfrak{X}$ that can be obtained from $\vec{\mathfrak{x}}$ through cyclic shifts. We also call $[\vec{\mathfrak{x}}]$ the *cyclic sequence associated with* $\vec{\mathfrak{x}}$. We observe that two necklace graphs are isomorphic if and only if their sequences of vertex labels are cyclically equivalent, whence in combinatorics, a necklace is often simply defined as a cyclic sequence (cyclic equivalence class of strings).

The connected components of a functional graph Γ_g of a function $g : X \to X$, where X is a finite set, may be viewed as necklace graphs. Indeed, we take the unique

directed cycle contained in a given connected component as the underlying digraph of the associated necklace graph. The label of a vertex x on that cycle is defined as the rooted tree isomorphism type of $\mathrm{Tree}_{\Gamma_g}(x)$. For example, if we denote by

- $\mathfrak{I}_0$ the digraph isomorphism type of the trivial rooted tree (consisting of a single vertex without arcs);
- $\mathfrak{I}_1$ the most common rooted tree isomorphism type above a periodic vertex in the functional graph of the exemplary generalized cyclotomic mapping f of $\mathbb{F}_{256}$ defined in (1.2) above (i.e., a rooted tree of height 2, where the root has in-degree 2 and one of the two neighbors of the root has in-degree 0, the other has in-degree 3);
- $\mathfrak{I}_2$ the digraph isomorphism type of Δ_3 in the above example (the seemingly chaotic numbering for the $\mathfrak{I}_j$ is chosen such that it matches with Table 4.4 in Section 4.2);
- $\mathfrak{I}_3$ the digraph isomorphism type of Δ_1;
- $\mathfrak{I}_4$ the digraph isomorphism type of Δ_2;

then the four connected components of the example above may be identified with necklace graphs corresponding to the following cyclic sequences of rooted tree isomorphism types (in order of drawing):

- $[\mathfrak{I}_0]$;
- $[\mathfrak{I}_1]$;
- $[\mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1]$;
- $[\mathfrak{I}_3, \mathfrak{I}_1, \mathfrak{I}_4, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_1, \mathfrak{I}_2]$.

With the above convention of identifying connected components of functional graphs with certain necklace graphs, two digraphs that are connected components of finite functional graphs are isomorphic as digraphs if and only if they are isomorphic as necklace graphs (i.e., they represent the same necklace of rooted tree isomorphism types). This means that in order to understand the connected components of Γ_g, we need to understand the associated cyclic sequences of rooted tree isomorphism types.

We note that if the goal is to understand the (undirected graph) isomorphism type of the underlying undirected graph of a connected component of a functional graph, an analogous approach can be used. One needs to replace necklace graphs by *bracelet graphs* (undirected, vertex-labeled cycle graphs), cyclic sequences by *dihedral sequences* (orbits of the natural action of the dihedral group of order $2L$ on $\mathfrak{X}^L$, where the generating reflection acts by writing the sequence in reverse order), and necklaces by *bracelets* (isomorphism classes of bracelet graphs).

Let us next explain our approach for understanding the digraph isomorphism types of the connected components of Γ_g via necklaces of rooted tree isomorphism types in case $X = \mathbb{F}_q$ and g is an index d generalized cyclotomic mapping f of $\mathbb{F}_q$.

For this, we first need to understand the rooted tree above a given (periodic) point. The basic idea is to construct a certain partition $\mathcal{P}_i$ of each coset C_i that "controls" the isomorphism types of rooted trees above the vertices in each of its blocks. Dealing with entire blocks of vertices at once is crucial to ensure that the complexities of our algorithms are not at least linear in the number of vertices q like many general-purpose algorithms for handling graph isomorphism, including Babai's breakthrough quasi-polynomial algorithm from [9].

In order to sketch how the said partition $\mathcal{P}_i$ of C_i is constructed, we need to introduce some more concepts. For a given positive integer m, we define the notion of an *m-(in)congruence* to be an (in)congruence of the form $\nu(x \equiv \mathfrak{b} \pmod{\mathfrak{a}})$, where $\nu \in \{\emptyset, \neg\}$ is a "logical sign", and $\mathfrak{a}, \mathfrak{b}$ are integers with $\mathfrak{a} \geq 1$ and $\mathfrak{a} \mid m$. We subsume these two notions under the name *m-congruential condition*, or *m-CC* for short. Next, we consider the concept of an *arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$*, see point (2) of the following definition.

Definition 1.5. Let m be a positive integer. We identify the elements of $\mathbb{Z}/m\mathbb{Z}$ with their standard representatives in $\{0, 1, \ldots, m-1\}$.

(1) Let $x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j}$ for $j = 1, 2, \ldots, K$ be m-congruences. There is a unique partition of $\mathbb{Z}/m\mathbb{Z}$, which we denote by

$$\mathfrak{P}(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K),$$

such that each block of this partition is the solution set modulo m of a system of m-CCs of the form

$$\begin{aligned} &\nu_1(x \equiv \mathfrak{b}_1 \pmod{\mathfrak{a}_1}),\\ &\nu_2(x \equiv \mathfrak{b}_2 \pmod{\mathfrak{a}_2}),\\ &\quad\vdots\\ &\nu_K(x \equiv \mathfrak{b}_K \pmod{\mathfrak{a}_K}), \end{aligned} \tag{1.3}$$

where $\nu_1, \nu_2, \ldots, \nu_K \in \{\emptyset, \neg\}$ are logical signs.

(2) A partition $\mathcal{P}$ of $\mathbb{Z}/m\mathbb{Z}$ is called *arithmetic* if it is of the form

$$\mathfrak{P}(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K)$$

for a suitable non-negative integer K and suitable m-congruences

$$x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j}$$

for $j = 1, 2, \ldots, K$. If $\mathcal{P} = \mathfrak{P}(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K)$, then we also say that $\mathcal{P}$ is the arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$ *spanned by* the congruences $x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j}$ for $j = 1, 2, \ldots, K$.

(3) When $\mathcal{P}$ is an arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$, then the smallest value of $K \in \mathbb{N}_0$ such that $\mathcal{P}$ is spanned by K suitably chosen m-congruences is called the *(arithmetic) complexity of* $\mathcal{P}$, written $\mathrm{AC}(\mathcal{P})$.

(4) When $\mathcal{P}$ is an arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$ and a sequence of spanning congruences $(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j})_{j=1,2,\ldots,K}$ has been fixed for $\mathcal{P}$, then we denote for each $\vec{\nu} = (\nu_1, \nu_2, \ldots, \nu_K) \in \{\emptyset, \neg\}^K$ by $\mathcal{B}(\mathcal{P}, \vec{\nu})$ the unique subset of $\mathbb{Z}/m\mathbb{Z}$ that is the solution set of the system (1.3) (this solution set is a block of $\mathcal{P}$ as long as it is non-empty).

Remark 1.6. There are significantly fewer arithmetic partitions of $\mathbb{Z}/m\mathbb{Z}$ than there are partitions in total. Indeed, the total number of (set) partitions of $\mathbb{Z}/m\mathbb{Z}$ is the Bell number B_m, which satisfies

$$B_m \sim \frac{1}{\sqrt{m}}\left(\frac{m}{W(m)}\right)^{m+\frac{1}{2}} \exp\left(\frac{m}{W(m)} - m - 1\right)$$

as $m \to \infty$, where $W(m) \sim \log m$ is the Lambert W function (see [47, Section 1.14, Problem 9]). In particular, as $m \to \infty$,

$$\log B_m \sim -\frac{1}{2}\log m + \left(m + \frac{1}{2}\right)(\log m - \log W(m)) + \frac{m}{W(m)} - m - 1 \sim m \log m.$$

On the other hand, every arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$ is spanned by a selection of congruences of the form $x \equiv \mathfrak{b} \pmod{\mathfrak{a}}$, where $\mathfrak{a}$ ranges over the positive divisors of m, and $\mathfrak{b} \in \{0, 1, \ldots, \mathfrak{a} - 1\}$. Because the total number of such congruences is $\sigma(m)$ (the sum of all positive divisors of m), it follows that the number of arithmetic partitions of $\mathbb{Z}/m\mathbb{Z}$ is at most $2^{\sigma(m)}$, and so its natural logarithm is at most

$$\log 2 \cdot \sigma(m) \leq \log 2(\mathrm{e}^{\gamma} + \varepsilon) m \log\log m,$$

where γ denotes the Euler–Mascheroni constant and the second bound follows from a result of Robin [63].

Let us return to our index d generalized cyclotomic mapping f of $\mathbb{F}_q$. We recall that $s = (q-1)/d$ denotes the order (size) of the index d subgroup C of $\mathbb{F}_q^*$. The aforementioned partitions $\mathcal{P}_i$ of the cosets C_i are constructed as arithmetic partitions of $\mathbb{Z}/s\mathbb{Z}$, with which C_i is to be identified via the bijection ι_i introduced above. They have the property that for vertices $x, y \in C_i$ chosen from a common block $\mathcal{B}(\mathcal{P}_i, \vec{\nu})$ of $\mathcal{P}_i$, one has $\mathrm{Tree}_{\Gamma_f}(x) \cong \mathrm{Tree}_{\Gamma_f}(y)$, and this common isomorphism type is denoted by $\mathrm{Tree}_i(\mathcal{P}_i, \vec{\nu})$. The constructions of the partitions $\mathcal{P}_i$ and of the associated rooted tree isomorphism types $\mathrm{Tree}_i(\mathcal{P}_i, \vec{\nu})$, which are carried out in detail in Section 3.3, are based on two crucial tools:

- the elementary result Lemma 2.2.2 from Section 2.2; and

- an explicit understanding, developed in Section 3.2 but also based on earlier theory developed in Section 2.1, of the structures of rooted trees in the induced subgraph Γ_{per} of Γ_f on the union of $\{0_{\mathbb{F}_q}\}$ with all cosets C_i, where i is $\bar{f}$-periodic.

Once the $\mathcal{P}_i$ and $\mathrm{Tree}_i(\mathcal{P}_i, \vec{\nu})$ have been constructed explicitly, in order to understand the isomorphism type of the connected component of Γ_f with representative periodic vertex r, one needs to understand how the cycle moves through the various blocks of the respective coset partitions. Of course, if the cycle length l of r under f is small, one can just enumerate the points on the cycle by brute force, check in which blocks they lie and spell out the corresponding cyclic sequence of rooted trees; this is what we do at the end of the example in Section 4.2. However, if l is large, then one can obtain a more concise description of the cyclic rooted tree sequence via a certain tuple of arithmetic partitions, the blocks of which represent intersections of the cycle of r with blocks of the involved arithmetic partitions $\mathcal{P}_i$. For details on this, see Section 3.4, which builds on Section 2.4.

Here is an overview of our approach for understanding Γ_f.

(1) Determine the induced function $\bar{f}\colon\{0,1,\ldots,d\}\to\{0,1,\ldots,d\}$, and rewrite f into a collection of affine functions that map between $d+1$ sets C_i, each of the form $\mathbb{Z}/s\mathbb{Z}$ or $\{0\}$.

(2) Compute a CRL-list $\mathcal{L}$ for f as specified in Section 3.1, which is based on the results for affine maps of finite cyclic groups from Section 2.3.

(3) For each $i \in \{0,1,\ldots,d-1\}$, compute the arithmetic partition $\mathcal{P}_i$ and associated rooted tree isomorphism types $\mathrm{Tree}_i(\mathcal{P}_i, \vec{\nu})$, as well as the isomorphism type of $\mathrm{Tree}_{\Gamma_f}(0_{\mathbb{F}_q})$, as specified in Section 3.3. This requires the theory developed in Sections 2.1 and 3.2.

(4) For each $(r,l) \in \mathcal{L}$, understand the associated cyclic sequence of rooted tree isomorphism types along the cycle of r under f, either

- by listing elements on the cycle of r by brute force, then looking up in which blocks of the relevant arithmetic partitions they lie, or
- by following the approach from Section 3.4, which relies on Section 2.4.

In Chapter 5, where we give a detailed algorithmic complexity analysis, we treat the procedures described in steps (2–4) each as a separate algorithm to be analyzed. We note that in general, f has too many cycles in order for it to be possible to spell out a CRL-list of f element-wise if the procedure is to be efficient (i.e., subexponential in $\log q$); one can, however, obtain a concise parametrization of a CRL-list of f efficiently. Likewise, the approach described in point (4) can be carried out for each given pair (r,l) individually in an efficient manner for asymptotically almost all finite fields $\mathbb{F}_q$, but it is not clear in general how to obtain a "global" understanding of Γ_f efficiently. In fact, the number of distinct isomorphism types of connected components of Γ_f might be superpolynomial in $\log q$ even for fixed d (cf. Prob-

lem 6.3.3), so one would first need to come up with a compact way of parametrizing those isomorphism types. Still, as we will see in Section 5.3, for some special cases of generalized cyclotomic mappings f of $\mathbb{F}_q$, there are algorithms for describing Γ_f as a whole which are efficient for all or at least for "most" q (in an asymptotic density sense). In particular, in those cases, it can be efficiently decided whether the functional graphs of two given generalized cyclotomic mappings of $\mathbb{F}_q$ are isomorphic.

Chapter 6 concludes the memoir with a list of open problems for further research. For the reader's convenience, an extensive index of the notation and terminology appearing in this memoir is given in Tables A.1 and A.2 in Appendix A.