Chapter 6

Open problems

We conclude this memoir with a discussion of open problems related to our results and methods.

6.1 Asymptotic behavior of mpe over prime powers

Our Proposition 5.1.10 states that as q ranges over an initial segment of all prime powers, the average value of $\mathrm{mpe}(q-1)$ (the maximum exponent of a prime in the full factorization of $q-1$) is bounded from above by a constant (independent of that segment). This led to the important observation that when d is fixed, then for asymptotically almost all finite fields $\mathbb{F}_q$, the complexities in Theorem 5.1.9 (2,3) are polynomial in $\log q$. Following that, at the end of Section 5.1, we raised the analogous problem restricted to powers of 2, and our computational evidence (gathered in the form of Table 5.1) leads to the following conjecture.

Conjecture 6.1.1. *The average value of* $\mathrm{mpe}(2^k-1)$*, where* k *ranges over an initial segment of* $\mathbb{N}^+$*, is always less than* 2*. Formally, this conjecture asserts that for each* $K \in \mathbb{N}^+$*, one has*

$$\frac{1}{K}\sum_{k=1}^{K} \mathrm{mpe}(2^k-1) < 2.$$

In fact, looking at Table 5.1, one might even conjecture that the said average value is always less than $3/2$, which would imply that $\mathrm{mpe}(2^k-1)=1$ (i.e., that 2^k-1 is square-free) for more than half of all $k \in \mathbb{N}^+$. We do note that the conjecture already fails when the prime 2 is replaced by 3. Indeed, since $4 = 2^2 \mid 3^k-1$ whenever $2 \mid k$, and $16 = 2^4 \mid 3^k-1$ whenever $4 \mid k$, one has

$$\frac{1}{K}\sum_{k=1}^{K} \mathrm{mpe}(3^k-1) \geq \frac{1}{2}\cdot 1 + \frac{1}{4}\cdot 2 + \frac{1}{4}\cdot 4 = 2$$

whenever $4 \mid K$. In general, there cannot be a universal constant c such that

$$\frac{1}{K}\sum_{k=1}^{K} \mathrm{mpe}(p^k-1) < c,$$

since $\mathrm{mpe}(p^k-1) \geq \mathrm{mpe}(p-1)$ for all k, and $\mathrm{mpe}(p-1)$ can be arbitrarily large (as follows, e.g., from Dirichlet's theorem on primes in arithmetic progressions, see [8, Chapter 7]). Of course, as a more general conjecture, it may be possible that for each

prime p, there is a constant c_p such that $\frac{1}{K}\sum_{k=1}^{K} \mathrm{mpe}(p^k - 1) < c_p$ for all positive integers K (and Conjecture 6.1.1 simply asserts that $c_2 = 2$ is a valid choice).

6.2 Efficient comparison of arithmetic partitions

We recall Problem 3 from the beginning of Chapter 5: given a generalized cyclotomic mapping f of $\mathbb{F}_q$ and a pair (r, l), where $r \in \mathbb{F}_q$ is f-periodic of cycle length l, the task is to obtain a compact description of the digraph isomorphism type of the connected component of Γ_f containing r, viewed as a necklace of rooted tree isomorphism types. In Problem 3, it is also assumed that a partition-tree register for f (in the sense of Definition 5.1.2) is given, and we can use this to refer to the rooted trees with their numbers $n \in \{0, 1, \ldots, N\}$ in this register, rather than spell each of them out completely.

Assuming that $r \neq 0_{\mathbb{F}_q}$ (the case "$r = 0_{\mathbb{F}_q}$" is easily dealt with separately), the main idea behind our algorithm from the proof of Theorem 5.1.9 (3) for tackling this problem is to identify the positions on the f-cycle of r that lie in a given coset C_j with the elements of $\mathbb{Z}/(l/\ell)\mathbb{Z}$, where ℓ is the associated coset cycle length (i.e., the $\bar{f}$-cycle length of i for the unique $i \in \{0, 1, \ldots, d-1\}$ such that $r \in C_i$), and to derive an arithmetic partition $\mathcal{P}^{(i_t)}$ of $\mathbb{Z}/(l/\ell)\mathbb{Z}$ for $t \in \{0, 1, \ldots, \ell - 1\}$ such that vertices in C_{i_t} on the cycle that lie in the same block of $\mathcal{P}^{(i_t)}$ have the same tree above them in Γ_f. Formally, we may express this via a labeling function $\mathrm{lab}_{i_t} : \mathcal{P}^{(i_t)} \to \{0, 1, \ldots, N\}$ that maps each block of $\mathcal{P}^{(i_t)}$ to its associated rooted tree number.

Now, the description of the connected component of Γ_f containing r obtained this way is *not* an injective encoding of its isomorphism type. That is, isomorphic connected components may end up getting different descriptions, and it appears to be a nontrivial computational problem to decide efficiently whether two given descriptions pertain to the same isomorphism type. The purpose of this section is to discuss this open problem in more detail, reducing it to some concrete questions to be answered.

Of course, in the actual implementation of our algorithm for Problem 3, each of the arithmetic partitions $\mathcal{P}^{(i_t)}$ mentioned above is expressed through a spanning congruence sequence, of length $\bar{m}_{i_t}$ say, and lab_{i_t} may be expressed through a function $\{\emptyset, \neg\}^{\bar{m}_{i_t}} \to \{-1, 0, 1, \ldots, N\}$, where -1 is a dummy value to be assigned to a logical sign tuple $\vec{\nu}$ if the associated set $\mathcal{B}(\mathcal{P}^{(i_t)}, \vec{\nu})$ is empty. By abuse of notation, we also call this function lab_{i_t}. We give a special name to ordered pairs such as $(\mathcal{P}^{(i_t)}, \mathrm{lab}_{i_t})$.

Definition 6.2.1. Let $m, N \in \mathbb{N}_0$ with $m > 0$. An *N-labeled arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$* is a pair $(\mathfrak{P}(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K), \mathrm{lab})$ consisting of an arithmetic partition of $\mathbb{Z}/m\mathbb{Z}$ with a fixed spanning congruence sequence and a so-called labeling function $\mathrm{lab} : \{\emptyset, \neg\}^K \to \{-1, 0, 1, \ldots, N\}$ such that $\mathrm{lab}(\vec{\nu}) = -1$ if and only if $\mathcal{B}(\mathcal{P}, \vec{\nu}) = \emptyset$.

In order to characterize when two compact descriptions obtained by the algorithm for Problem 3 represent isomorphic connected components, let us first consider the special case where $\ell = 1$. Then we only need to worry about the coset C_i and the associated N-labeled arithmetic partition $(\mathcal{P}^{(i)}, \mathrm{lab}_i)$ of $\mathbb{Z}/l\mathbb{Z}$.

We need to understand how applying a cyclic shift to the associated sequence of rooted tree isomorphism types affects the N-labeled arithmetic partition. Let $\mathfrak{t}$ be the translation $x \mapsto x + 1$ of $\mathbb{Z}/l\mathbb{Z}$. It generates a cyclic subgroup $\mathfrak{T}$ of order l of $\mathrm{Sym}(\mathbb{Z}/l\mathbb{Z})$, namely the image of the regular representation of $\mathbb{Z}/l\mathbb{Z}$ on itself. This group $\mathfrak{T}$ also acts naturally on the power set of $\mathbb{Z}/l\mathbb{Z}$ via

$$M^t := t(M) = \{y^t : y \in M\}.$$

In the same manner, this leads to an action of $\mathfrak{T}$ on the power set of the power set of $\mathbb{Z}/l\mathbb{Z}$ (i.e., an action which transforms families of subsets of $\mathbb{Z}/l\mathbb{Z}$ into other such families), and this action restricts to one on the set of all arithmetic partitions of $\mathbb{Z}/l\mathbb{Z}$. Indeed, if $\mathcal{P} = \mathfrak{P}(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K)$ is an arithmetic partition of $\mathbb{Z}/l\mathbb{Z}$, then $\mathcal{P}^{\mathfrak{t}}$, the partition of $\mathbb{Z}/l\mathbb{Z}$ obtained by shifting all blocks of $\mathcal{P}$ to the right by one unit, is just the arithmetic partition $\mathfrak{P}(x \equiv \mathfrak{b}_j + 1 \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K)$ of $\mathbb{Z}/l\mathbb{Z}$.

If we assume that $\mathcal{P}$ is N-labeled and that each block of $\mathcal{P}^t$, where $t \in \mathfrak{T}$, carries the same label as the block of $\mathcal{P}$ it is shifted from, we finally get an action of $\mathfrak{T}$ on the set of N-labeled arithmetic partitions of $\mathbb{Z}/l\mathbb{Z}$, which is useful for our characterization. Formally, this action is defined via

$$\begin{aligned}&(\mathfrak{P}(x \equiv \mathfrak{b}_j \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K), \mathrm{lab})^{\mathfrak{t}^n}\\ &\quad := (\mathfrak{P}(x \equiv \mathfrak{b}_j + n \pmod{\mathfrak{a}_j} : j = 1, 2, \ldots, K), \mathrm{lab}).\end{aligned}$$

In the special case "$\ell = 1$" we are currently discussing, applying a right cyclic shift of n units to the cyclic sequence of rooted tree isomorphism types associated with the N-labeled arithmetic partition $(\mathcal{P}^{(i)}, \mathrm{lab}_i)$ corresponds to replacing $(\mathcal{P}^{(i)}, \mathrm{lab}_i)$ by $(\mathcal{P}^{(i)}, \mathrm{lab}_i)^{\mathfrak{t}^n}$. This motivates the following definition.

Definition 6.2.2. Let $m, N \in \mathbb{N}_0$ with $m > 0$, and let $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ be N-labeled arithmetic partitions of $\mathbb{Z}/m\mathbb{Z}$. These partitions are *equivalent* if there is a $t \in \mathfrak{T}$ such that $(\mathcal{P}', \mathrm{lab}')^t = (\mathcal{P}, \mathrm{lab})$. In that case, the smallest positive integer n such that $(\mathcal{P}', \mathrm{lab}')^{\mathfrak{t}^n} = (\mathcal{P}, \mathrm{lab})$ is the *translation number of* $(\mathcal{P}, \mathrm{lab})$ *and* $(\mathcal{P}', \mathrm{lab}')$. If $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ are not equivalent, then we define their translation number to be ∞.

We observe that translation numbers are not symmetric in their two arguments. Rather, if $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ are equivalent, then the translation number of $(\mathcal{P}', \mathrm{lab}')$ and $(\mathcal{P}, \mathrm{lab})$ is the difference of the stabilizer order

$$|\mathrm{Stab}_{\mathfrak{T}}((\mathcal{P}, \mathrm{lab}))| = |\{t \in \mathfrak{T} : (\mathcal{P}, \mathrm{lab})^t = (\mathcal{P}, \mathrm{lab})\}| = |\mathrm{Stab}_{\mathfrak{T}}((\mathcal{P}', \mathrm{lab}'))|$$

and the translation number of $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$. For example, if $(\mathcal{P}, \mathrm{lab})^{\mathfrak{t}^3} = (\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')^{\mathfrak{t}} = (\mathcal{P}, \mathrm{lab})$, then $(\mathcal{P}, \mathrm{lab})^{\mathfrak{t}^2} = (\mathcal{P}', \mathrm{lab}')$.

As we mentioned just before Definition 6.2.2, in case $\ell = 1$, two N-labeled arithmetic partitions of the form $(\mathcal{P}^{(i)}, \mathrm{lab}_i)$ and $(\mathcal{P}^{(j)}, \mathrm{lab}_j)$ correspond to isomorphic connected components of Γ_f if and only if they are equivalent. We thus pose the following algorithmic problem.

Problem 6.2.3. *Find an efficient algorithm which, for given $m, N \in \mathbb{N}_0$ with $m > 0$ and N-labeled arithmetic partitions $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ of $\mathbb{Z}/m\mathbb{Z}$, computes the translation number of $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$.*

We note that by our convention on translation numbers of inequivalent N-labeled arithmetic partitions, such an algorithm could in particular be used to efficiently decide whether $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ are equivalent in the first place. Moreover, it could be used to determine the stabilizer in $\mathfrak{T}$ of a given N-labeled arithmetic partition $(\mathcal{P}, \mathrm{lab})$, because that stabilizer is generated by $\mathfrak{t}^n$, where n is the translation number of $(\mathcal{P}, \mathrm{lab})$ with itself.

We now return from "$\ell = 1$" to the general case. By the details of our identification of the l positions on the f-cycle of r with the elements in ℓ disjoint copies of $\mathbb{Z}/(l/\ell)\mathbb{Z}$ (see Section 5.2.3), it is not hard to see that in general, applying a right cyclic shift by n units to the sequence of rooted tree isomorphism types associated with the sequence $((\mathcal{P}^{(i_t)}, \mathrm{lab}_{i_t}))_{t=0,1,\ldots,\ell-1}$ corresponds to replacing that sequence with

$$\begin{aligned}&(((\mathcal{P}^{(i_j)}, \mathrm{lab}_{i_j}))_{j=0,1,\ldots,\ell-1})^{\mathfrak{t}^n}\\ &\quad := ((\mathcal{P}^{(i_j)}, \mathrm{lab}_{i_j})^{\mathfrak{t}^{n''+1}})_{j=\ell-n',\ell-n'+1,\ldots,\ell-1} \diamond ((\mathcal{P}^{(i_j)}, \mathrm{lab}_{i_j})^{\mathfrak{t}^{n''}})_{j=0,1,\ldots,\ell-n'-1},\end{aligned}$$

where $n' := n \bmod \ell$ and $n'' := (n - n')/\ell$. This defines an action of $\mathfrak{T}$ on the set of all length ℓ sequences of N-labeled arithmetic partitions of $\mathbb{Z}/(l/\ell)\mathbb{Z}$, and as for $\ell = 1$, we call two such sequences *equivalent* if they can be mapped to each other under this action. In order to decide in general whether two descriptions produced by our algorithm for Problem 3 correspond to isomorphic connected components, we need to decide whether these descriptions are equivalent in this more general sense. However, it turns out that this can be done efficiently if we have an algorithm as in Problem 6.2.3. Let us explain why.

We assume that $(\mathcal{X}_k, \mathrm{lab}^{(k)})_{k=0,1,\ldots,\ell-1}$ and $(\mathcal{Y}_k, \mathrm{Lab}^{(k)})_{k=0,1,\ldots,\ell-1}$ are length ℓ sequences of N-labeled arithmetic partitions of $\mathbb{Z}/m\mathbb{Z}$ for some $m \in \mathbb{N}^+$. They are equivalent if and only if there are $n' \in \{0, 1, \ldots, \ell - 1\}$ and $n'' \in \mathbb{Z}$ such that $((\mathcal{Y}_k, \mathrm{Lab}^{(k)}))_{k=0,1,\ldots,\ell-1}$ is equal to

$$((\mathcal{X}_t, \mathrm{lab}^{(t)})^{\mathfrak{t}^{n''+1}})_{t=\ell-n',\ell-n'+1,\ldots,\ell-1} \diamond ((\mathcal{X}_t, \mathrm{lab}^{(t)})^{\mathfrak{t}^{n''}})_{t=0,1,\ldots,\ell-n'-1}. \tag{6.1}$$

To check whether this is the case, we assume that n' is fixed (in the worst case, we need to try out $\ell \in O(d)$ values for n'). For $t = 0, 1, \ldots, \ell - 1$, we compute the number

$$\mathfrak{b}_t := \begin{cases} \text{translation number of } (\mathcal{Y}_t, \mathrm{Lab}^{(t)}) \text{ and } (\mathcal{X}_{\ell-n'+t}, \mathrm{lab}^{(\ell-n'+t)}), & \text{if } t < n', \\ \text{translation number of } (\mathcal{Y}_t, \mathrm{Lab}^{(t)}) \text{ and } (\mathcal{X}_{t-n'}, \mathrm{lab}^{(t-n')}), & \text{otherwise} \end{cases}$$

using the algorithm from Problem 6.2.3. If any of these numbers is ∞, then the chosen value of n' does not work. Otherwise, we compute $\mathfrak{a}_t$, the group order of the stabilizer of $(\mathcal{Y}_t, \mathrm{Lab}^{(t)})$ in $\mathfrak{T}$ (which here is a cyclic group of order m), using the said algorithm. The question is whether there exists $n'' \in \mathbb{Z}$ such that

$$\begin{aligned} n'' + 1 &\equiv \mathfrak{b}_0 \pmod{\mathfrak{a}_0}, \\ n'' + 1 &\equiv \mathfrak{b}_1 \pmod{\mathfrak{a}_1}, \\ &\vdots \\ n'' + 1 &\equiv \mathfrak{b}_{n'-1} \pmod{\mathfrak{a}_{n'-1}}, \\ n'' &\equiv \mathfrak{b}_{n'} \pmod{\mathfrak{a}_{n'}}, \\ n'' &\equiv \mathfrak{b}_{n'+1} \pmod{\mathfrak{a}_{n'+1}}, \\ &\vdots \\ n'' &\equiv \mathfrak{b}_{\ell-1} \pmod{\mathfrak{a}_{\ell-1}} \end{aligned}$$

because these congruences characterize when $((\mathcal{Y}_t, \mathrm{Lab}^{(t)}))_{t=0,1,\ldots,\ell-1}$ is equal to (6.1). Viewing this as a system of m-congruences in the single variable n'', the existence of n'' can easily be decided using Proposition 2.2.1.

We conclude this section by noting that the algorithm from Problem 6.2.3 can be used to decide whether two (N-labeled) arithmetic partitions are equal, i.e., have the same (labeled) blocks. This is because two N-labeled arithmetic partitions $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ of $\mathbb{Z}/m\mathbb{Z}$ are equal if and only if the translation number of $(\mathcal{P}, \mathrm{lab})$ and $(\mathcal{P}', \mathrm{lab}')$ equals the translation number of $(\mathcal{P}, \mathrm{lab})$ with itself. Moreover, $\mathcal{P}$ and $\mathcal{P}'$ are equal if and only if $(\mathcal{P}, \mathbf{0})$ and $(\mathcal{P}', \mathbf{0})$ are equal, where (in each of the two cases) $\mathbf{0}$ denotes the constantly zero labeling function. Still, the algorithmic problem of verifying whether two given arithmetic partitions of $\mathbb{Z}/m\mathbb{Z}$ are equal is interesting in its own right, and it may admit an efficient algorithmic solution even if Problem 6.2.3 does not, so we pose it separately.

Problem 6.2.4. *Find an efficient algorithm which, for given $m \in \mathbb{N}^+$ and (spanning m-congruence sequences of) arithmetic partitions $\mathcal{P}$ and $\mathcal{Q}$ of $\mathbb{Z}/m\mathbb{Z}$, decides whether $\mathcal{P} = \mathcal{Q}$.*

6.3 More problems concerning asymptotic growth rates

Let $m = p_1^{v_1} \cdots p_K^{v_K}$ be a positive integer with its factorization displayed. We recall from Definition 1.5 (3) that the minimal number of spanning m-congruences for an arithmetic partition $\mathcal{P}$ of $\mathbb{Z}/m\mathbb{Z}$ is called the (arithmetic) complexity of $\mathcal{P}$ and denoted by $\mathrm{AC}(\mathcal{P})$. In Remark 3.3.6, we observed that the trivial partition $\mathcal{T}_m$ of $\mathbb{Z}/m\mathbb{Z}$, all of whose blocks are singletons, satisfies $\mathrm{AC}(\mathcal{T}_m) \leq \sum_{j=1}^{K} p_j^{v_j} - K$. While this bound is equal to $m - 1$ when m is a prime power, we also observed in Remark 3.3.6 that the bound is asymptotically equivalent to $\log^2 m/(2 \log\log m)$ when m is a primorial, which leads to the question whether the actual complexity of $\mathcal{T}_m$ can be significantly smaller than that (on a suitable infinite class of values for m).

Question 6.3.1. *Is it true that*

$$\liminf_{m\to\infty} \frac{\mathrm{AC}(\mathcal{T}_m)}{\log^2 m/\log\log m} > 0?$$

We observe that $\mathcal{T}_m$ is the unique (arithmetic) partition of $\mathbb{Z}/m\mathbb{Z}$ that achieves the maximum possible number of blocks, m. One may ask more generally for nontrivial bounds that relate the arithmetic complexity of an arithmetic partition $\mathcal{P}$ of $\mathbb{Z}/m\mathbb{Z}$ with its number of blocks. Trivially, the number of distinct blocks of $\mathcal{P}$ is at most $2^{\mathrm{AC}(\mathcal{P})}$, and this bound is attained if $\mathrm{AC}(\mathcal{P}) \in \{0, 1\}$. In fact, for any given value $k \in \mathbb{N}_0$, there is an $m \in \mathbb{N}^+$ and an arithmetic partition $\mathcal{P}$ of $\mathbb{Z}/m\mathbb{Z}$ such that $\mathrm{AC}(\mathcal{P}) = k$ and $\mathcal{P}$ has 2^k distinct blocks: simply let m be the k-th primorial $p_k\# = p_1 p_2 \cdots p_k$, where p_j denotes the j-th smallest prime number, and

$$\mathcal{P} := \mathfrak{P}(x \equiv 0 \pmod{p_j} : j = 1, 2, \ldots, k).$$

However, once $\mathrm{AC}(\mathcal{P})$ becomes sufficiently large with respect to m, the number of blocks of $\mathcal{P}$ falls behind $2^{\mathrm{AC}(\mathcal{P})}$; at latest, this happens once $\mathrm{AC}(\mathcal{P}) > \log_2 m$, because the block count of $\mathcal{P}$ cannot be larger than m. In the example we just gave, where $m = p_k\#$, we have $\mathrm{AC}(\mathcal{P}) = k \sim (\log m/\log\log m)$, which motivates the following open problem.

Problem 6.3.2. *Either find functions $f, g : [0, \infty) \to [0, \infty)$ such that*

(1) $(\log x/\log\log x) \lesssim f(x) \in o(\log x)$,

(2) $g(x) \in o(2^x)$, *and*

(3) *for every positive integer m and every arithmetic partition $\mathcal{P}$ of $\mathbb{Z}/m\mathbb{Z}$ with $\mathrm{AC}(\mathcal{P}) \geq f(m)$, the number of blocks of $\mathcal{P}$ is at most $g(\mathrm{AC}(\mathcal{P}))$,*

or prove that such functions do not exist.

In Section 5.3.2, we took note of a potential obstacle to using tree necklace lists to give a general, efficient algorithm for deciding whether the functional graphs of

two given generalized cyclotomic mappings of $\mathbb{F}_q$ are isomorphic. Namely, it could be that even when their index is fixed, generalized cyclotomic mappings f have too many distinct isomorphism types of connected components in their functional graphs. Specifically, we pose the following problem.

Problem 6.3.3. *Prove or disprove that for every $d \in \mathbb{N}^+$, there is a constant $c = c(d)$ such that for every prime power q and every index d generalized cyclotomic mapping f of $\mathbb{F}_q$, the number of distinct isomorphism types of connected components of Γ_f is in $O(\log^c q)$.*

We observe that for $d = 1$, all rooted trees above *non-zero* f-periodic points are isomorphic; see Theorem 2.1.5, noting that $\Gamma_{f_{|\mathbb{F}_q^*}} \cong \Gamma_{A_0}$, unless f is constantly zero, in which case the statement in question is vacuously true. Therefore, for $d = 1$, Problem 6.3.3 is equivalent to proving or disproving that the number of distinct cycle lengths of an affine map of $\mathbb{Z}/(q-1)\mathbb{Z}$, where q ranges over all prime powers, is bounded from above by some fixed polynomial in $\log q$. Even this appears to be an open problem, in spite of Remark 5.3.2.4.

In view of Proposition 5.3.1.4, we accept that an arbitrarily small but positive asymptotic fraction of prime powers q needs to be excluded in order for the algorithms from Section 5.3.2 to be efficient. In this context, we note the following problem related to Problem 6.3.3.

Problem 6.3.4. *For $d \in \mathbb{N}^+$ and $c > 0$, we denote by $\varepsilon(d, c)$ the asymptotic proportion of all prime powers q with $d \mid q-1$ for which there exists an index d generalized cyclotomic mapping f of $\mathbb{F}_q$ such that Γ_f has more than $\log^c q$ distinct isomorphism types of connected components. Prove or disprove that as d is fixed and $c \to \infty$, one has $\varepsilon(d, c) \to 0$.*

Of course, a proof of the assertion in Problem 6.3.3 also yields a proof of the assertion in Problem 6.3.4. On the other hand, if the assertion in Problem 6.3.3 is false (which the authors believe is the case), the one in Problem 6.3.4 may still be true, thus allowing one to solve the isomorphism problem for functional graphs of generalized cyclotomic mappings of a fixed index efficiently in "most" cases (only having to exclude an arbitrarily small positive asymptotic fraction of cases if one accepts a suitably large degree in the poly-log bound).

6.4 Extension to other coset-wise affine functions

An index d generalized cyclotomic mapping f of $\mathbb{F}_q$, given in cyclotomic form (1.1), such that all a_i and r_i are non-zero restricts to a function $\mathbb{F}_q^* \to \mathbb{F}_q^*$, which is "coset-wise affine" in the sense that its restriction to any given coset C_i of the index d subgroup C of $\mathbb{F}_q^*$ maps to another coset $C_{\bar{f}(i)}$ via an affine map of the cyclic group C

(here, we are using the general, group-theoretic sense of the word "affine map", as in Definition 2.1.15). That we can split f up into such smaller, easy to handle parts is crucial for the approach of understanding Γ_f presented in this memoir.

In this section, we aim to generalize this idea. More specifically, we replace $\mathbb{F}_q^*$ by some group G (usually, but not necessarily finite), and C by a subgroup H of G. We consider the following two notions of coset-wise affine functions.

Definition 6.4.1. Let G be a group, and let H be a subgroup of G.

(1) An *affine function* $H \to G$ is a function $H \to G$ of the form $h \mapsto h^\varphi b$ for some group homomorphism $\varphi : H \to G$ and some $b \in G$.

(2) A function $f : G \to G$ is called *H-coset-wise affine in the wide sense* if for every right coset $C = Hr_C$ of H in G, there is an affine function $A_C : H \to G$ such that $f(hr_C) = A_C(h)$ for all $h \in H$.

(3) A function $f : G \to G$ is called *H-coset-wise affine in the narrow sense* if for every right coset $C = Hr_C$ of H in G, there is an affine map A_C of H and a $t_C \in G$ such that $f(hr_C) = A_C(h)t_C$ for all $h \in H$.

In contrast to H-coset-wise affine functions in the narrow sense, an H-coset-wise affine function in the wide sense does not need to map each right coset of H to a single such coset. This makes it hard to study the behavior of H-coset-wise affine functions in the wide sense under iteration. The most celebrated example of this is the Collatz function g, corresponding to $G = \mathbb{Z}$ and $H = 2\mathbb{Z}$ and given by the coset-wise affine formula

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

On the other hand, any function $g : \mathbb{Z} \to \mathbb{Z}$ such that g agrees, on each coset $k + n\mathbb{Z}$ of the index n subgroup $n\mathbb{Z}$, with an affine function

$$x \mapsto a^{(k)}x + b^{(k)}$$

with *integer* coefficients $a^{(k)}$ and $b^{(k)}$, is $n\mathbb{Z}$-coset-wise affine in the narrow sense and thus amenable to the ideas mentioned in the first paragraph of this section. Henceforth, we restrict our attention to H-coset-wise affine functions in the narrow sense, which we simply call *H-coset-wise affine functions* for short.

An important special case is when $G = (\mathbb{F}_q, +)$ and H is an $\mathbb{F}_p$-subspace of G, for which this class of functions was already considered in [14]. We expect our approach for understanding functional graphs of generalized cyclotomic mappings to work mostly analogously for H-coset-wise affine functions, with one big caveat: in the proof of Lemma 2.2.2 (our "master lemma"), we made essential use of the equivalence of statements (1) and (3) in Proposition 2.2.1. In the more general group-theoretic context of the current section, this equivalence needs to be replaced by the following property of the group H.

Definition 6.4.2. Let H be a group. We say that H is *pairwise congruence-consistent* if any given system of congruences over H,

$$\begin{aligned} x &\equiv h_1 \pmod{N_1}, \\ x &\equiv h_2 \pmod{N_2}, \\ &\vdots \\ x &\equiv h_K \pmod{N_K}, \end{aligned}$$

where $h_1, h_2, \ldots, h_K \in H$ and $N_1, N_2, \ldots, N_K$ are normal subgroups of H, is consistent if and only if each pair of congruences in the system is consistent.

The equivalence of statements (1) and (3) in Proposition 2.2.1 can be reformulated as "Finite cyclic groups are pairwise congruence-consistent." In the special case "$G = (\mathbb{F}_q, +)$" mentioned above, the group H is of the form $\mathbb{F}_p^n$, i.e., it is a finite elementary abelian p-group. If our approach is to work completely analogously for that case, we would need that finite elementary abelian groups are pairwise congruence-consistent. However, that is not the case, as the following result shows (noting that all abelian groups are nilpotent).

Theorem 6.4.3. *Let G be a finite nilpotent group. The following are equivalent:*

(1) *G is pairwise congruence-consistent.*

(2) *G is cyclic.*

We prove Theorem 6.4.3 at the end of this section. Before doing so, we make two more comments.

Firstly, we observe that non-cyclic finite pairwise congruence-consistent groups exist; there are both non-solvable examples, such as any non-abelian finite simple group (for trivial reasons), and solvable examples, such as $\mathrm{AGL}_1(q) = \mathbb{F}_q \rtimes \mathbb{F}_q^*$ for any prime power q. To see that the latter kind of groups are pairwise congruence-consistent, we note that $\mathrm{AGL}_1(q)$ has $\mathbb{F}_q$ as its unique minimal, nontrivial normal subgroup, so any congruence over $\mathrm{AGL}_1(q)$ has an associated congruence over the cyclic group $\mathbb{F}_q^*$ (i.e., a congruence in the classical, number-theoretic sense) such that the solution set of the congruence over $\mathrm{AGL}_1(q)$ is the full pre-image, under the canonical projection $\mathrm{AGL}_1(q) \to \mathbb{F}_q^*$, of the solution set of the associated number-theoretic congruence. In particular, if any pair of congruences in a given system of congruences over $\mathrm{AGL}_1(q)$ is consistent, the same holds true for the associated system over $\mathbb{F}_q^*$, whence that system is consistent by Proposition 2.2.1, and so the original system over $\mathrm{AGL}_1(q)$ must also have a solution.

Secondly, we pose the following two open problems, which are motivated by Theorem 6.4.3 and the discussion leading to it.

Problem 6.4.4. *Classify the finite groups that are pairwise congruence-consistent.*

Problem 6.4.5. *For important classes of finite groups that are* not *contained in the class of finite pairwise congruence-consistent groups (such as the class of finite (elementary) abelian groups), devise efficient algorithms that decide whether a given system of congruences over a group in that class is consistent.*

In order to prove Theorem 6.4.3, we first consider the following property of groups.

Definition 6.4.6. Let G be a group. We say that G has the *pairwise coset-intersection property* (or *PCIP* for short) if the following holds: for any positive integer m and any sequence $(C_1, C_2, \ldots, C_m)$ of (left or right) cosets of subgroups of G, if $C_j \cap C_k \neq \emptyset$ for all $1 \leq j < k \leq m$, then $\bigcap_{j=1}^{m} C_j \neq \emptyset$. A group satisfying the PCIP is also called a *PCIP-group* for short.

The following proposition is immediate from observing that the solution set of the congruence $x \equiv g \pmod{N}$ over the group G is the coset $gN = Ng$ of N.

Proposition 6.4.7. *Let G be a group. The following are equivalent.*

(1) *G is pairwise congruence-consistent.*

(2) *For any positive integer m and any sequence $(C_1, C_2, \ldots, C_m)$ of cosets of normal subgroups of G, if $C_j \cap C_k \neq \emptyset$ for all $1 \leq j < k \leq m$, then $\bigcap_{j=1}^{m} C_j \neq \emptyset$.*

Proposition 6.4.7 has two important consequences.

Corollary 6.4.8. *The following hold.*

(1) *Every PCIP-group is pairwise congruence-consistent.*

(2) *An abelian group satisfies the PCIP if and only if it is pairwise congruence-consistent.*

In view of Corollary 6.4.8, the following result proves Theorem 6.4.3 for *abelian* groups.

Theorem 6.4.9. *Let G be a finite group. The following are equivalent.*

(1) *G satisfies the PCIP.*

(2) *G is cyclic.*

Proof. The implication "(2)⇒(1)" holds by Proposition 2.2.1, so we focus on "(1)⇒(2)".

It is not hard to show that all subgroups and quotients of a PCIP-group are PCIP-groups themselves. A minimal counterexample to the implication "(1)⇒(2)" would thus be a finite, non-cyclic group G all of whose proper subgroups are cyclic. These groups were classified by Miller and Moreno [52] to be one of the following.

(1) $\mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$, where p is a prime;

(2) the quaternion group Q_8; or

(3) the metacyclic group

$$\mathbb{Z}/q^n\mathbb{Z} \ltimes \mathbb{Z}/p\mathbb{Z} = \langle x, y : x^p = y^{q^n} = 1, y^{-1}xy = x^r \rangle,$$

where $r \equiv 1 \pmod{q}$ and $r^q \equiv 1 \pmod{p}$, but $r \not\equiv 1 \pmod{p}$, since otherwise, the group is cyclic or isomorphic to $\mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$.

It suffices to show that none of these groups G satisfies the PCIP, which we do now, by specifying three subsets $C_j \subseteq G$ for $j = 1, 2, 3$, each of which is a left and right coset of some subgroup of G, such that the C_j intersect pairwise while $\bigcap_{j=1}^{3} C_j = \emptyset$.

- For groups of the first type, where the elements are pairs (x, y) with $x, y \in \mathbb{F}_p$, let

$$C_1 := \{(x, y) \in \mathbb{F}_p^2 : x = 1\}, \quad C_2 := \{(x, y) \in \mathbb{F}_p^2 : y = 1\},$$

and

$$C_3 := \{(x, y) \in \mathbb{F}_p^2 : x + y = 1\}.$$

- For $Q_8 = \{\pm 1, \pm i, \pm j, \pm k\}$, let

$$C_1 := \langle i \rangle j = j \langle i \rangle = \{\pm j, \pm k\}, \quad C_2 := \langle j \rangle i = i \langle j \rangle = \{\pm i, \pm k\},$$

and

$$C_3 := \langle k \rangle j = j \langle k \rangle = \{\pm i, \pm j\}.$$

- We note that the central quotient of a group of the third type is of the same form but with $n = 1$, which we may thus assume without loss of generality. Let

$$C_1 = \langle x \rangle y = y \langle x \rangle = \{y, yx, yx^2, \ldots, yx^{p-1}\},$$
$$C_2 = \langle y \rangle = \{1, y, y^2, \ldots, y^{q-1}\},$$

and

$$C_3 = \langle yx \rangle = \{1, yx, y^2x^{1+r}, y^3x^{1+r+r^2}, \ldots, y^{q-1}x^{1+r+r^2+\cdots+r^{q-2}}\}.$$

Since the y-exponent is 1 in all elements of C_1, we have $C_1 \cap C_2 = \{y\}$ and $C_1 \cap C_3 = \{yx\}$. Moreover, $C_2 \cap C_3 = \{1\}$, since C_2 and C_3 are distinct subgroups of prime order. It follows that $C_1 \cap C_2 \cap C_3 = \emptyset$, as required. ∎

Proof of Theorem 6.4.3. As in the proof of Theorem 6.4.9, the implication "(2)⇒(1)" is clear by Proposition 2.2.1, so we focus on "(1)⇒(2)".

Let G be a finite nilpotent group that is pairwise congruence-consistent. If G is abelian, then G must be cyclic by Corollary 6.4.8 (2) and Theorem 6.4.9, so we

assume (aiming for a contradiction) that G is non-abelian. Then for some prime p, the (unique) Sylow p-group S_p of G is non-abelian. Because G is the direct product of its Sylow subgroups, we find that S_p is a quotient of G. Moreover, it is not hard to prove that quotients of pairwise congruence-consistent groups are themselves pairwise congruence-consistent. Therefore, S_p is pairwise congruence-consistent.

But S_p is a non-abelian finite p-group, whence Burnside's basis theorem implies that its Frattini quotient $S_p/\Phi(S_p)$ is the finite elementary abelian p-group $\mathbb{F}_p^n$, where $n > 1$ is the minimal size of a generating set of S_p. Using again that the property of being pairwise congruence-consistent is preserved under passing to quotients, it follows that $\mathbb{F}_p^n$ is pairwise congruence-consistent. Since $\mathbb{F}_p^n$ is abelian, Corollary 6.4.8 thus implies that $\mathbb{F}_p^n$ satisfies the PCIP, which contradicts Theorem 6.4.9. ∎

6.5 Generalization to transformation graphs

This memoir is concerned with functional graphs, which are natural visualizations of individual functions on a set and are useful for understanding the long-term behavior of discrete dynamical systems. As a generalization, one may consider the situation where a dynamical system does not evolve deterministically, but for some $n \in \mathbb{N}^+$, each system state x has n possibilities (possibly with repetitions) for its successor state, occurring with different probabilities and represented by the values $g_1(x), g_2(x), \ldots, g_n(x)$ of functions $g_j : X \to X$. Let us set $\mathcal{G} := \{g_1, g_2, \ldots, g_n\}$. A first step toward studying the behavior of such a system is to understand the so-called *transformation graph* $\mathrm{TRAG}(X, \mathcal{G})$, which is defined as the edge-labeled digraph with vertex set X whose arcs are of the form $x \xrightarrow{g_j} g_j(x)$ for $x \in X$ and $j = 1, 2, \ldots, n$. The terminology "transformation graph" is from Annexstein, Baumslag and Rosenberg's paper [7]. The concept is also closely related to operands, which are actions of semigroups on sets [19, Section 11.1] (and in analogy to the terminology "group action graph" from [7], one could also call transformation graphs "operand graphs"), and to deterministic finite automata [35, Section 2.2.1]. In fact, $\mathrm{TRAG}(X, \mathcal{G})$ is like a deterministic finite automaton with state set X and input symbol set $\mathcal{G}$, but without declared start and accept states.

We observe that except for the edge labels, $\mathrm{TRAG}(X, \{g\})$ is the same as the functional graph Γ_g in our notation. As noted in [7, beginning of Section 2.1], one may also consider a simple (i.e., no multiple arcs $x \to x'$ for given $x, x' \in X$), unlabeled version of $\mathrm{TRAG}(X, \mathcal{G})$, which we denote by $\mathrm{STRAG}(X, \mathcal{G})$. Of course, for a given set X, any digraph with vertex set X in which each vertex has positive outdegree (including possibly ∞) is of the form $\mathrm{STRAG}(X, \mathcal{G})$ for a suitable non-empty $\mathcal{G} \subseteq X^X$.

It would be interesting to know whether the methods developed in our memoir could be extended to deal with graphs of the form $\mathrm{TRAG}(\mathbb{F}_q, \mathcal{F})$ and $\mathrm{STRAG}(\mathbb{F}_q, \mathcal{F})$,

where $\mathcal{F}$ is a set of generalized cyclotomic mappings of $\mathbb{F}_q$, say of a common, small index d (which also covers some cases where the index is not uniform, because if f_j for $j = 1, 2, \ldots, n$ is a generalized cyclotomic mapping of $\mathbb{F}_q$ of index d_j, then each f_j also has index $\mathrm{lcm}(d_1, d_2, \ldots, d_n)$). Specifically, we pose the following problems.

Problem 6.5.1. *For a given prime power q and set $\mathcal{F}$ of index d generalized cyclotomic mappings of $\mathbb{F}_q$, devise efficient algorithms (say of runtime polynomial in $\log q$ for fixed d) that find*

(1) *a compact parametrization of the connected components of* $\mathrm{TRAG}(\mathbb{F}_q, \mathcal{F})$ *(equivalently, of* $\mathrm{STRAG}(\mathbb{F}_q, \mathcal{F})$*) by representative vertices, and*

(2) *a compact description of the isomorphism type of a connected component of* $\mathrm{TRAG}(\mathbb{F}_q, \mathcal{F})$*, respectively of* $\mathrm{STRAG}(\mathbb{F}_q, \mathcal{F})$*, given by a vertex in the image of the parametrization from point (1).*

To the authors' knowledge, this is an open problem even for $d = 1$ and $|\mathcal{F}| = 2$ (i.e., when considering transformation graphs that are each based on two monomial functions over $\mathbb{F}_q$).

Problem 6.5.2. *For some classes of sets of index d generalized cyclotomic mappings of $\mathbb{F}_q$, devise efficient algorithms to decide, for sets $\mathcal{F}_1, \mathcal{F}_2$ in such a class, whether* $\mathrm{TRAG}(\mathbb{F}_q, \mathcal{F}_1) \cong \mathrm{TRAG}(\mathbb{F}_q, \mathcal{F}_2)$*, respectively,* $\mathrm{STRAG}(\mathbb{F}_q, \mathcal{F}_1) \cong \mathrm{STRAG}(\mathbb{F}_q, \mathcal{F}_2)$.