Functional Graphs of Power Maps over Finite Fields: Cycle Structure, Rho-Length Anomalies, and the d-Adic Valuation

Alexander Towell*

Abstract

We study the functional graphs of power maps $f_d(x) = x^d \bmod p$ over finite fields $\mathbb{F}_p$ for $d \in \{2, 3, 5, 7\}$ and all primes p up to 10^5. For each prime, we compute the complete cycle structure: cycle lengths, rho-lengths (tail lengths before entering a cycle), number of connected components, and the fraction of cyclic versus tail points.

We prove that the mean rho-length is controlled by the d-adic valuation $v_d(p-1)$: when $d \mid p-1$, the trees in the functional graph are complete d-ary trees of height $h = v_d(p-1)$, and the mean rho-length over tail points is exactly $(h \cdot d^{h+1} - (h+1)d^h + 1)/((d-1)(d^h-1))$, which is asymptotically $h - 1/(d-1)$. Primes with anomalously large rho-lengths have $p - 1 = d^k \cdot m$ for large k and small m; for $d = 2$, the extremal prime is $p = 65537 = 2^{16} + 1$ (a Fermat prime) with $\bar{\rho} = 15$.

We verify computationally that this formula matches the measured mean rho-length exactly (to machine precision) across all 18,355 non-permutation (d, p) pairs surveyed, and that the cyclic fraction for non-permutation primes exceeds the Flajolet–Odlyzko random map prediction by a large factor, reflecting the uniform branching structure of power maps.

MSC 2020: 11T06, 37P05, 05C20.

Keywords: functional graph, power map, finite field, cycle structure, rho-length, d-adic valuation.

1 Introduction

For a prime p and integer $d \geq 2$, the *power map* $f_d \colon \mathbb{F}_p \to \mathbb{F}_p$ defined by $f_d(x) = x^d$ is a fundamental object in number theory. Its iteration underlies Pollard's rho factoring algorithm [6], discrete logarithm attacks, and the study of arithmetic dynamics over finite fields [8].

The *functional graph* of f_d is the directed graph on vertex set $\mathbb{F}_p$ with edges $x \to f_d(x)$. Since $\mathbb{F}_p$ is finite, every connected component consists of a cycle with trees hanging off cycle nodes. Two key parameters describe the structure:

- The *cycle length* $\lambda(x)$: the length of the cycle that the orbit of x eventually enters.
- The *rho-length* $\rho(x)$: the number of iterations before the orbit enters the cycle (also called the *tail length* or *pre-period*).

For a random map on N elements, the expected rho-length and cycle length are both $\Theta(\sqrt{N})$ by the analysis of Flajolet and Odlyzko [3]. Power maps, however, are not random: the algebraic structure of $\mathbb{F}_p^*$ imposes constraints that can dramatically alter these statistics.

*Department of Computer Science, Southern Illinois University Edwardsville. Email: atowell@siue.edu. ORCID: 0000-0001-6443-9897.

Prior work. Vasiga and Shallit [10] studied the functional graphs of $x \mapsto x^2$ and $x \mapsto x^2 - 2$ over $\mathrm{GF}(p)$, giving cycle length formulas and statistical measures. Chou and Shparlinski [2] extended these results to general fixed d, providing unconditional asymptotic formulas for periodic point counts averaged over primes. Somer and Křížek [9] characterized cycle decompositions of power digraphs modulo n, and Ahmad and Syed [1] determined vertex heights. Most recently, Qureshi and Reis [7] showed that the trees attached to cyclic vertices in the functional graph of x^d over an abelian group decompose as tensor products of elementary trees. Martins and Panario [4] compared polynomial functional graphs to random maps, though without isolating monomial maps x^d as a separate class.

Despite this body of work, no prior study has identified the d-adic valuation $v_d(p-1)$ as the single arithmetic parameter controlling rho-length deviations from random map predictions.

Our contributions. We conduct a systematic computational survey of the functional graphs of f_d for $d \in \{2, 3, 5, 7\}$ across all primes $p \leq 10^5$. Our main findings are:

1. **Exact rho-length formula (Theorem 4.1).** When $d \mid p-1$, the trees in $G(p,d)$ are complete d-ary trees of height $h = v_d(p-1)$, and the mean rho-length over tail points has the closed form $(hd^{h+1} - (h+1)d^h + 1)/((d-1)(d^h-1))$. We verify this exactly across all 18,355 non-permutation (d,p) pairs in the survey.

2. **Anomalous primes.** Primes with the largest rho-lengths are precisely those with $p-1 = d^k \cdot m$ for large k. The extremal case for $d = 2$ is the Fermat prime $p = 65537$.

3. **Permutation statistics.** The fraction of primes for which f_d is a permutation converges to the predicted density from Dirichlet's theorem.

4. **Deviation from random maps.** The cyclic fraction of power maps exceeds the Flajolet–Odlyzko random map prediction by a large factor (roughly $60\times$ at $p \approx 50{,}000$ for $d = 2$), reflecting the uniform branching structure.

2 Definitions and notation

Let p be an odd prime and $d \geq 2$ an integer.

Definition 2.1. The *power map* $f_d \colon \mathbb{F}_p \to \mathbb{F}_p$ is $f_d(x) = x^d$. Its *functional graph* $G(p,d)$ is the directed graph on vertex set $\{0, 1, \ldots, p-1\}$ with edges $x \to x^d \bmod p$.

Since $|\mathbb{F}_p|$ is finite, each weakly connected component of $G(p,d)$ contains a unique cycle with rooted trees attached to cycle vertices. We write $\rho(x)$ for the *rho-length* (tail length) of x: the smallest $k \geq 0$ such that $f_d^k(x)$ lies on a cycle. Points with $\rho(x) = 0$ are *cyclic*; all others are *tail points*. We write $\bar{\rho}$ for the mean of $\rho(x)$ over all tail points.

Definition 2.2. The *d-adic valuation* of a positive integer n, written $v_d(n)$, is the largest integer $k \geq 0$ such that $d^k \mid n$.

Three basic facts govern the structure of $G(p,d)$:

1. f_d is a permutation of $\mathbb{F}_p^*$ if and only if $\gcd(d, p-1) = 1$.

2. The number of fixed points of f_d in $\mathbb{F}_p$ is $\gcd(d-1, p-1) + 1$ (including $x = 0$).

3. The image $f_d(\mathbb{F}_p^*)$ is the subgroup of d-th powers in $\mathbb{F}_p^*$, which has size $(p-1)/\gcd(d, p-1)$.

3 Computational results

We computed the full functional graph of f_d for $d \in \{2, 3, 5, 7\}$ and every prime $p \leq 10^5$. Table 1 summarizes the survey.

Table 1: Survey summary for $d \in \{2, 3, 5, 7\}$, all primes $p \leq 10^5$.

d	Primes	Permutation	Non-perm.	Mean $\bar{\rho}$	Mean cyclic frac.
2	9,591	0 (0%)	9,591	1.330	0.334
3	9,590	4,806 (50.1%)	4,784	1.126	0.250
5	9,589	7,202 (75.1%)	2,387	1.043	0.166
7	9,588	7,995 (83.4%)	1,593	1.013	0.126

The permutation fractions are consistent with the predicted densities: for prime d, the fraction of primes p with $\gcd(d, p-1) = 1$ is $1 - 1/(d-1)$ by Dirichlet's theorem, giving 0% for $d = 2$, 50% for $d = 3$, 75% for $d = 5$, and $83.\overline{3}$% for $d = 7$.

Rho-length by d-adic valuation. For $d = 2$, Table 2 shows the mean rho-length stratified by $v_2(p-1)$. The predicted values from Theorem 4.1 match the observed means to within 10^{-3} across all 9,591 primes.

Table 2: Mean rho-length over tail points, by $v_2(p-1)$, for $d = 2$. The predicted column is the exact formula from Theorem 4.1. All 9,591 primes $p \leq 10^5$ match to machine precision.

$v_2(p-1)$	Count	Observed $\bar{\rho}_{\text{tail}}$	Predicted
1	4,808	1.0000	1.0000
2	2,399	1.6667	1.6667
3	1,196	2.4286	2.4286
4	589	3.2667	3.2667
5	299	4.1613	4.1613
6	154	5.0952	5.0952
7	75	6.0551	6.0551
8	32	7.0314	7.0314
9–16	39	(all match exactly)	

The exact match is not a coincidence: the theorem's formula is a deterministic function of d and $v_d(p-1)$, independent of the cofactor $m = (p-1)/d^{v_d(p-1)}$. The tree structure depends only on the d-adic part of $p - 1$.

Anomalous primes. The three primes with largest mean rho-length for each d are shown in Table 3. In every case, $p - 1$ has a large d-adic valuation, confirming that the d-adic structure of $p - 1$ is the sole driver of rho-length anomalies.

4 Mean rho-length and the d-adic valuation

We now prove the main structural result. Let $h = v_d(p-1)$ denote the d-adic valuation of $p - 1$: the largest integer k such that $d^k \mid p - 1$.

Table 3: Top three anomalous primes by mean rho-length, each d.

d	p	$p-1$ factorization	v_d	$\bar{\rho}$	Predicted
2	65,537	2^{16}	16	15.00	15.00
2	40,961	$2^{13} \cdot 5$	13	12.00	12.00
2	86,017	$2^{12} \cdot 3 \cdot 7$	12	11.00	11.00
3	39,367	$2 \cdot 3^9$	9	8.50	8.50
3	52,489	$2^3 \cdot 3^8$	8	7.50	7.50
3	87,481	$2^3 \cdot 3^7 \cdot 5$	7	6.50	6.50
5	62,501	$2^2 \cdot 5^6$	6	5.75	5.75
5	37,501	$2^2 \cdot 3 \cdot 5^5$	5	4.75	4.75
5	97,501	$2^2 \cdot 3 \cdot 5^4 \cdot 13$	4	3.75	3.75
7	72,031	$2 \cdot 3 \cdot 5 \cdot 7^4$	4	3.83	3.83
7	28,813	$2^2 \cdot 3 \cdot 7^4$	4	3.83	3.83
7	14,407	$2 \cdot 3 \cdot 7^4$	4	3.83	3.83

Theorem 4.1. *Let p be prime, let d be a prime with $d \mid p-1$, and let $h = v_d(p-1)$. Write $\bar{\rho}$ for the mean rho-length over all tail points in the functional graph of $f_d(x) = x^d$ on $\mathbb{F}_p$. Then each tree rooted at a nonzero cyclic point has $d-1$ branches at the root and d branches at every other internal node, with height h and $d^h - 1$ tail points. The number of tail points at depth j (i.e., with $\rho(x) = j$) is $(d-1) \cdot d^{j-1}$ for $j = 1, \ldots, h$. The mean rho-length satisfies*

$$\bar{\rho} \;=\; \frac{h \cdot d^{h+1} - (h+1) \cdot d^h + 1}{(d-1)(d^h - 1)}.$$

In particular, $\bar{\rho} = h - \frac{1}{d-1} + O(d^{-h})$ as $h \to \infty$.

Proof. Since d is prime and $d \mid p-1$, every element of $\mathrm{Im}(f_d) \cap \mathbb{F}_p^*$ (the d-th powers) has exactly d preimages under f_d in $\mathbb{F}_p^*$. The cyclic points in $\mathbb{F}_p^*$ are the d^h-th powers, forming a subgroup of size $(p-1)/d^h$. Each cyclic point c is the root of a tree whose structure we now describe.

At depth 1, the point c has d preimages under f_d, one of which is cyclic (the predecessor on the cycle). The remaining $d-1$ preimages are tail points at depth 1. Each tail point at depth $j < h$ has exactly d preimages (since it lies in $\mathrm{Im}(f_d^{h-j})$, the subgroup of d^{h-j}-th powers, and $d \mid p-1$ ensures full branching). At depth h, the preimages are elements of $\mathbb{F}_p^* \setminus \mathrm{Im}(f_d)$ (the non-d-th-powers), which have no further preimages.

Thus the tree has $d-1$ branches at the root, d branches at each internal node, and height h. The number of tail points at depth j is $(d-1) \cdot d^{j-1}$, giving $\sum_{j=1}^{h} (d-1)d^{j-1} = d^h - 1$ tail points per tree.

The mean rho-length over all tail points (which is the same for every tree, by the vertex-transitivity of $\mathbb{F}_p^*$ under translations) is

$$\bar{\rho} = \frac{\sum_{j=1}^{h} j \cdot (d-1) \cdot d^{j-1}}{\sum_{j=1}^{h} (d-1) \cdot d^{j-1}} = \frac{\sum_{j=1}^{h} j \cdot d^{j-1}}{(d^h - 1)/(d-1)}.$$

Using the standard identity $\sum_{j=1}^{h} j \cdot x^{j-1} = \frac{hx^{h+1} - (h+1)x^h + 1}{(x-1)^2}$ at $x = d$, the numerator becomes

$\frac{h \cdot d^{h+1} - (h+1)d^h + 1}{(d-1)^2}$ and the denominator is $(d^h - 1)/(d-1)$, so

$$\bar{\rho} = \frac{h \cdot d^{h+1} - (h+1)d^h + 1}{(d-1)(d^h - 1)}. \qquad \square$$

Remark 4.2. When $\gcd(d, p-1) < d$ but $\gcd(d, p-1) > 1$, the trees are no longer complete d-ary trees, but the mean rho-length is still controlled by $v_d(p-1)$. We verify this computationally in Section 3. When $\gcd(d, p-1) = 1$, the power map is a permutation and $\bar{\rho} = 0$.

5 Comparison with random maps

For a uniform random map on N elements, Flajolet and Odlyzko [3] showed that the expected fraction of cyclic points is $\sqrt{\pi/(2N)}$, which tends to zero as $N \to \infty$. Power maps deviate sharply from this prediction.

For $d = 2$, the fraction of cyclic points in $G(p, 2)$ is $(p-1)/2^{v_2(p-1)} + 1$ divided by p, which converges to a positive constant (approximately 1/3 on average over primes) rather than to zero. The Flajolet–Odlyzko prediction for $p \approx 50{,}000$ is $\sqrt{\pi/10^5} \approx 0.006$, while the observed mean cyclic fraction is 0.334 — a factor of roughly 60.

This discrepancy arises because power maps have highly non-random indegree distributions: every element in $\mathrm{Im}(f_d)$ has exactly $\gcd(d, p-1)$ preimages among nonzero elements, while $(p-1)(1 - 1/\gcd(d, p-1))$ elements have zero preimages. This uniform branching structure produces complete d-ary trees, in contrast to the Poisson-distributed tree shapes of random maps.

6 Method

For each prime p and exponent d, we compute the full functional graph by evaluating $f_d(x) = x^d \bmod p$ for all $x \in \{0, \ldots, p-1\}$ using Python's built-in modular exponentiation (`pow(x, d, p)`), which uses exact integer arithmetic. We then determine each element's rho-length and cycle length via Floyd's cycle detection algorithm, which requires $O(1)$ space per starting point. The total computation for one (d, p) pair takes $O(p)$ time and $O(p)$ space.

All results are verified against three algebraic invariants:

1. The number of fixed points equals $\gcd(d-1, p-1) + 1$.
2. The image size equals $(p-1)/\gcd(d, p-1) + 1$.
3. The total number of cyclic plus tail points equals p.

The full survey ($d \in \{2, 3, 5, 7\}$, all primes $p \leq 10^5$) required approximately 75 minutes of single-threaded computation on an Intel Core i7-12700H (2.3 GHz, 32 GB RAM). Computational analysis performed with AI assistance (Claude, Anthropic); all results verified by independent algebraic invariant checks.

7 Data availability

This paper is archived with DOI `10.5281/zenodo.19476511`. Source code, survey data (JSON), and verification scripts are available at `https://github.com/queelius/open-problems`. The survey data files contain, for each (d, p): the mean rho-length, maximum rho-length, number of cyclic points, number of tail points, number of connected components, and cycle length distribution.

8 Concluding remarks

The mean rho-length formula in Theorem 4.1 applies when d is prime and $d \mid p-1$. Several natural extensions remain open.

1. **Composite d.** When d is composite, the branching factor at each tree level may vary (since $\gcd(d^j, p-1) \neq d^{\min(j,h)}$ in general). A generalization of the formula involving the full sequence $\gcd(d, p-1), \gcd(d^2, p-1), \ldots$ should be possible using the tensor product tree decomposition of Qureshi and Reis [7].

2. **Extension fields.** The functional graph of x^d over $\mathbb{F}_{p^n}$ for $n > 1$ has a richer structure because $\mathbb{F}_{p^n}^*$ is cyclic of order $p^n - 1$. The d-adic valuation of $p^n - 1$ controls the tree height, but the interaction with the Galois structure of $\mathbb{F}_{p^n}$ introduces additional symmetries.

3. **Implications for Pollard's rho.** The Pollard rho factoring algorithm iterates $x \mapsto x^2 + c$ (not $x \mapsto x^d$), but the structural similarity suggests that tail lengths in Pollard iterations may also be governed by the 2-adic valuation of the group order, rather than by random map heuristics alone.

References

[1] U. Ahmad and H. Syed. On the heights of power digraphs modulo n. *Czech. Math. J.*, 62(2):541–556, 2012.

[2] W.-S. Chou and I. E. Shparlinski. On the cycle structure of repeated exponentiation modulo a prime. *J. Number Theory*, 107(2):345–356, 2004.

[3] P. Flajolet and A. M. Odlyzko. Random mapping statistics. *Advances in Cryptology – EUROCRYPT '89*, LNCS 434, pp. 329–354. Springer, 1990.

[4] R. Martins and D. Panario. On the heuristic of approximating polynomials over finite fields by random mappings. *Int. J. Number Theory*, 12(7):1987–2016, 2016.

[5] R. Martins, D. Panario, and C. Qureshi. A survey on iterations of mappings over finite fields. In *Radon Ser. Comput. Appl. Math.*, vol. 23, pp. 135–172. De Gruyter, 2019.

[6] J. M. Pollard. A Monte Carlo method for factorization. *BIT Numerical Mathematics*, 15(3):331–334, 1975.

[7] C. Qureshi and L. Reis. On the functional graph of the power map over finite groups. *Discrete Math.*, 346(6):113373, 2023.

[8] J. H. Silverman. *The Arithmetic of Dynamical Systems*. Graduate Texts in Mathematics, vol. 241. Springer, 2007.

[9] L. Somer and M. Křížek. On a connection of number theory with graph theory. *Czech. Math. J.*, 54(2):465–485, 2004.

[10] T. Vasiga and J. Shallit. On the iteration of certain quadratic maps over GF(p). *Discrete Math.*, 277(1–3):219–240, 2004.