How Fast Does the Inverse Walk Approximate a Random Permutation?

Tianren Liu[*]
Peking University

Angelos Pelecanos[†]
UC Berkeley

Stefano Tessaro[‡]
University of Washington

Vinod Vaikuntanathan[§]
MIT

May 16, 2025

Abstract

For a finite field $\mathbb{F}$ of size n, the (patched) inverse permutation $\mathrm{INV} : \mathbb{F} \to \mathbb{F}$ computes the inverse of x over $\mathbb{F}$ when $x \neq 0$ and outputs 0 when $x = 0$, and the ARK_K (AddRoundKey) permutation adds a fixed constant K to its input, i.e.,

$$\mathrm{INV}(x) = x^{n-2} \quad \text{and} \quad \mathrm{ARK}_K(x) = x + K \ .$$

We study the process of alternately applying the INV permutation followed by a random linear permutation ARK_K, which is a random walk over the alternating (or symmetric) group that we call the *inverse walk*.

We show both lower and upper bounds on the number of rounds it takes for this process to approximate a random permutation over $\mathbb{F}$. We show that r rounds of the inverse walk over the field of size n with

$$r = \Theta\left(n \log n + n \log \frac{1}{\varepsilon}\right)$$

rounds generate a permutation that is ε-close (in variation distance) to a uniformly random even permutation (i.e. a permutation from the alternating group A_n). Our result is provided with explicit constants, and is tight, up to logarithmic factors.

Our result answers an open question from the work of Liu, Pelecanos, Tessaro, and Vaikuntanathan (CRYPTO 2023) by proving the t-wise independence of (a variant of) AES for t up to the square root of the field size, compared to the original result that only held for $t = 2$. It also constitutes a significant improvement on a result of Carlitz (Proc. American Mathematical Society, 1953) who showed a *reachability result*: namely, that every even permutation can be generated *eventually* by composing INV and ARK. We show a *tight convergence result*, namely a tight quantitative bound on the number of rounds to reach a random (even) permutation.

Our work brings to the forefront the view of block ciphers as random walks and uses an entirely new set of (functional analytic) tools to study their pseudorandomness, both of which we hope will prove useful in the study of block ciphers.

[*]trl@pku.edu.cn
[†]apelecan@berkeley.edu
[‡]tessaro@cs.washington.edu
[§]vinodv@mit.edu

Contents

1 Introduction

The design and analysis of block ciphers such as the Advanced Encryption Standard (AES) [DR02] is a central topic in cryptography. On the one hand, despite extensive cryptanalysis, spanning a wide range of attacks including linear [MY92] and differential [BS91] cryptanalysis, higher-order [Lai94], truncated [Knu94] and impossible [Knu98] differential attacks, interpolation [JK97] and algebraic attacks [CP02], integral cryptanalysis [KW02], biclique attacks [BKR11], there has not been a single devastating attack thus far that undermines our confidence in AES. On the other hand, the situation is unsatisfactory from a foundational perspective: indeed, it is not clear whether it is even possible to formulate a meaningful non-tautological computational hardness assumption that implies the security of AES within the classical framework of provable security.

In this work, we continue the line of research that attempts to formally prove the security of block ciphers against *restricted* classes of attacks, with a focus on *substitution permutation networks* (SPNs), an important class of block ciphers that includes AES. The guiding principle of this line of study is to gradually expand the class of attacks we consider to include a large set of known cryptanalytic paradigms. In particular, we build on a recent pair of works by Liu, Pelecanos, Tessaro, and Vaikuntanathan (LPTV) [LTV21, LPTV23] who study the t-wise independence of SPNs, a statistical security property that prevents all t-input statistical attacks including differential and linear cryptanalysis (with $t = 2$) and higher order differential attacks (with larger t).[1]

Censored AES. The first of these works [LTV21] showed the 2-wise independence of the AES construction with many (over 9000) rounds. In an attempt to show comparable results with a lower number of rounds, the second work [LPTV23] looked instead at SPNs with uniformly chosen *random* and secret S-boxes, and proved t-wise independence of a construction called AES*, first introduced by Baignères and Vaudenay [BV05], which differs from AES in that it uses such random S-boxes. They also suggest a generic way to instantiate their results with the concrete AES S-box, i.e., the patched inverse permutation $\mathrm{INV} : x \mapsto x^{2^8-2}$ over the binary extension field $\mathbb{F}_{2^8}$. They observe in particular that a key-alternating cipher obtained by iterating INV, alternated with adding a random sub-key between each two sequential calls of INV, converges pretty quickly to being a *pairwise independent* permutation. We refer to this construction as the "INV KAC." Replacing the random S-box in their AES* result with the INV-KAC, they obtain in particular the following result, which they cast in terms of a construction they call "censored AES", which is essentially AES with some of the mixing layers removed.

Theorem 1.1 ([LPTV23], Theorem 7)**.** *192-round censored AES is 2^{-128}-close to pairwise independent.*

It is natural to conjecture that the actual AES cipher is not less secure than its censored counterpart, i.e., additional mixing layers only help, and so one can conjecture that the bound extends to 192-round AES, hence improving the bound of [LTV21] under this conjecture.

Their result however only applies to pairwise independence. This calls the question of whether the INV-KAC can be proved to be t-wise independent, in order to prove similar results for $t > 2$. In [LTV21], the authors sketch a proof for why this cipher needs at least linear in the size of the field many rounds to converge to a 4-wise independent permutation, but we note that this is not necessarily a major limitation in a context where the domain itself has small size, i.e, 256 elements. This argument relies on a similar idea as the interpolation attack by Jakobsen and Knudsen [JK97].

[1] We note here two caveats of their results: the first is that they assume the round keys are independent; and secondly, their proof works for many rounds of SPN/AES, although this has been improved in subsequent works.

Our contributions. In this work, we prove almost matching lower and upper bounds on the number of rounds for the INV KAC to reach t-wise independence. We note here that since the cipher composes an AddRoundKey (henceforth ARK) operation with an INV operation, each round of the cipher generates a permutation with a fixed parity (the parity depends on the size of the field). Thus, we can only hope for the cipher to converge to the alternating group (as opposed to the symmetric group) which means that t has to be at most $n-2$, for n being the size of our field.

Theorem (Lower bound, Theorem 4.2)**.** *An r-round* INV *KAC over the field of size n requires at least $r \geq \frac{(1-\varepsilon)n}{4} - \frac{1}{2}$ rounds to reach ε-close to a 4-wise independent permutation.*

Theorem (Upper bound, Theorem 3.1)**.** *An r-round* INV *KAC over the field of size n with*

$$r = O\left(n \log n + n \log \frac{1}{\varepsilon}\right)$$

rounds generates a permutation that is ε-close to a uniformly random even permutation (equivalently a uniformly random permutation from the alternating group A_n).

We provide explicit constants in the formal version of our result, Theorem 3.1. We also provide an asymptotically worse result in Theorem 3.2 with slightly better constants which is stronger than Theorem 3.1 for fields of size below a fixed threshold.

The proofs of our theorems view each round of the INV KAC as a step in a random walk over the alternating group, starting from the identity permutation. In each step you apply a random one of the n many permutations $\Pi_K(x) = \mathrm{INV}(x+K)$ where the addition is over the underlying field. Thus, a random walk of length r in this graph is equivalent to a composition of r many random permutations $\Pi_{K_1}, \Pi_{K_2}, \ldots, \Pi_{K_r}$, where the randomness is over the choice of the round keys $K_1, K_2, \ldots, K_r$. The problem then reduces to bounding the mixing time of this random walk over the alternating group A_n (here, n is the size of the field which, in the case of AES, is 2^8).

A prior version of our upper bound theorem first appeared in the appendix of [LPTV24] with the polynomially worse bound $r \lesssim O(n^2 \log n)$, and only for fields of characteristic 2. In this paper we use the comparison method to bound the three functional inequalities[2] of the underlying random walk, and are able to obtain asymptotically tighter bounds on the mixing time, compared to only using the spectral gap. We include an additional comparison result that works for fields of odd characteristic.

Having the upper bound, one can now extend the random S-box results of [LPTV23] to the concrete AES S-box. In particular, by following the same proof as in Theorem 7 of [LPTV23], but for larger t, we get the following corollary.

Corollary 1.2. *Assuming $t < 2^{(0.499-1/(4k))b}$, $\Theta\left(b2^b \cdot \min\{k, \log t\}\right)$-round censored SPN with k b-bit blocks, the AES S-box, and a maximal-branch-number linear mixing is $2^{-\Theta(kb)}$-close to t-wise independent.*

The Censoring Conjecture. If one believes that the mixing layers are useful for AES to achieve pseudorandomness, then it is natural to expect that removing a large fraction of them should only hurt the convergence to t-wise independence. This leads us to conjecture that r-round AES is not less secure than r-round censored AES.

Conjecture 1.3 (Censoring conjecture)**.** *For any t, there exists a fixed constant $r_{\text{thresh}} > 0$, such that for any number of rounds $r > r_{\text{thresh}}$, r-round AES is at least as close to t-wise independence as any r-round censored AES.*

[2] We will refer to the log-Sobolev constant, the modified log-Sobolev constant, and the spectral gap as the functional inequalities of a Markov chain.

We introduce a small round threshold r_{thresh} to allow the block cipher to enter a "generic" state, and avoid any non-typical behavior that may be happening during the first few rounds.

It is worth noting that a censoring result of a similar flavor has appeared in the Markov chain literature by Peres and Winkler [PW11]. Their setting is limited to a very specific family of Markov chains (Glauber dynamics in a monotone spin system, when the starting position is an extremal configuration) that do not include the behavior of a typical block cipher. Nevertheless, this censoring result has found many applications, and fully understanding how censoring affects the mixing time of a Markov chain remains an open problem.

A Mathematical Motivation. In 1963, Carlitz [Car53, Car63, Zie13] proved that the group of all permutations of $\mathbb{F}_q$ is generated by the permutations induced by degree-one polynomials and INV. Our result extends Carlitz's theorem in two ways. First, we show that if we restrict our attention to degree-one polynomials whose linear coefficient is 1, then we still generate at least the group of all even permutations[3]. Second, while Carlitz shows that every permutation can be reached via a composition of degree-one polynomials and INV, it does not tell us anything about how many steps it takes to do so: we provide an almost-tight quantitative bound on the number of steps (operations) needed to reach a random permutation. We remark that the notion of Carlitz rank of a permutation, namely the number of "Carlitz steps" one needs to take to get the permutation, has been studied in the literature; see [AcMT09, Top14, IW18].

In reverse, the study of this mathematical problem brings to bear sophisticated techniques from both finite field theory (cf. [Car53, Car63] and followups) and Markov chain theory to the practical problem of proving block cipher security.

Related work. We believe that the INV KAC is an important cipher design to study for two reasons. First, many substitution-permutation networks (SPNs) like the Advanced Encryption Standard (AES) [DR02] use the INV as their (non-linear) S-box permutation. Thus studying the INV "in isolation" may provide insight into its strengths and weaknesses. Additionally, the INV KAC is a block cipher that we can understand almost exactly. Furthermore, our analysis of the INV KAC, perhaps in conjunction with generalizations of Carlitz's result to general power maps [Sta98], may be useful in the analysis of other KACs, such as MiMC [AGR+16], which uses the cube instead of the inverse.

Our techniques. Our upper bound result follows from two separate lemmas, one for fields of characteristic 2 and one for fields of odd characteristic. In both cases, when dealing with a field of size n, we employ the comparison method of Sinclair, Diaconis and Saloff-Coste [Sin92, DSC93] to establish the lower bound of $\beta_n^{\text{INV}} \gtrsim \frac{1}{n}$ on the modified log-Sobolev constant of the Markov chain P_n^{INV} (described above) on the alternating group A_n in which every step corresponds to one round of the INV KAC. (For a definition of the modified log Sobolev constant and its implication for mixing times, the reader is referred to Section 2.2.) Such a bound on the modified log-Sobolev constant implies a mixing time bound of $O(n \log n + n \log \frac{1}{\varepsilon})$ and thus the same number of rounds is required for the INV KAC to approximate a uniformly random even permutation.

Using the same comparison method, we can also lower bound two other functional inequalities of the P_n^{INV} Markov chain, namely the log-Sobolev constant $\alpha_n^{\text{INV}} \gtrsim \frac{1}{n \log n}$ and the spectral gap $\lambda_n^{\text{INV}} \gtrsim \frac{1}{n}$. The latter two bounds in conjunction imply a different, asymptotically worse mixing time of $O(n \log^2 n + n \log \frac{1}{\varepsilon})$ for P_n^{INV}. However, this bound comes with slightly better constant factors and it is stronger than the former mixing time bound whenever the size of the field n is smaller than a fixed constant threshold. We include both bounds in our paper for completeness and because they can be derived using the same

[3] In some cases this is the best we can do, e.g. over $\mathbb{F}_7$, since both the INV and ARK are even permutations. For other fields (including fields of characteristic 2), we can generate the entire symmetric group since we combine both even and odd permutations.

techniques. For the rest of this discussion, we restrict our attention to the former bound that uses the modified log-Sobolev constant.

At a high level, the comparison method requires one to construct a multi-commodity flow between pairs of vertices that are connected in a Markov chain, by using paths over the P_n^{INV} edges (which correspond to composing INV and ARK operations). Furthermore, this flow must have low congestion, that is, we would like the flow passing through every edge in P_n^{INV} to be roughly uniform. The main technical ingredient of this work is constructing these flows and showing that they have low congestion.

For technical reasons, the comparison method cannot be applied to bound the modified log-Sobolev constant, but only the log-Sobolev constant and the spectral gap. To overcome this barrier, we use ideas that have recently appeared in [LMR+24] to modify the comparison method and use the same multi-commodity flow but a slightly different definition of congestion that depends more quadratically instead of linearly on the lengths of the paths in the flow. Since the flows we construct in this paper have paths of constant length, the additional factors do not cause an issue for us.

For fields of characteristic 2, we first compare this Markov chain to the chain $P_n^{\mathrm{2cyc},0}$, in which a random step corresponds to transposing a uniformly random element with the element 0, and leaving all other elements intact. The key component in such a comparison is to show how to generate transpositions of the form $(0, i)$ for any non-zero i using ARK and INV. To do this, we extend the proof of Carlitz [Car63]. In turn, the log-Sobolev constant of $P_n^{\mathrm{2cyc},0}$ can be bounded by a standard comparison to the well-studied random transpositions walk P_n^{2cyc} where each step is a transposition of the form (i, j) that swaps i and j and leaves the rest of the domain intact [DS81, LY98, FOW18, Sal20].

For fields of odd characteristic, we compare the P_n^{INV} Markov chain to the chain $P_n^{\mathrm{3cyc,ap}}$, in which a random step corresponds to applying a 3-cycle chosen uniformly at random from a specific subset of all 3-cycles (that involve elements from an arithmetic progression). Then, we bound the log-Sobolev constant of $P_n^{\mathrm{3cyc,ap}}$ by comparing it to the also well-studied Markov chain P_n^{3cyc} that applies a random 3-cycle at every step [Goe04, STY22]. We do this by showing how one can combine a constant number of these arithmetic-progression 3-cycles to generate an arbitrary 3-cycle over the n elements.

One may wonder whether there exists a common upper bound approach for fields of characteristic both odd and even. We remark that our comparison approach in the characteristic-2 case follows in the footsteps of [Car63] and shows how to perform a transposition. Our exact path construction turns out to be very different from Carlitz's, as the ARK operation does not implement any degree-one polynomial (as Carlitz requires) but is restricted to linear *shifts*. Instead, we show that we can implement the following set of transpositions

$$\left\{\left(u + v(uv+1)^{-1}, u + (v+1)(uv+u+1)^{-1}\right)\right\}_{u,v}$$

for most values of $u, v \in \mathbb{F}_n$, using the following 7 rounds (8 random keys) of the INV KAC:

$$\begin{aligned}\mathrm{ARK}_u \circ \mathrm{INV} \circ \mathrm{ARK}_u \circ \mathrm{INV} \circ \mathrm{ARK}_{v+1} \circ \mathrm{INV} \circ \mathrm{ARK}_1 \circ \\ \circ \mathrm{INV} \circ \mathrm{ARK}_1 \circ \mathrm{INV} \circ \mathrm{ARK}_v \circ \mathrm{INV} \circ \mathrm{ARK}_u \circ \mathrm{INV} \circ \mathrm{ARK}_u\,.\end{aligned}$$

The resulting multi-commodity flow does not give an optimal lower bound for the log-Sobolev constant since the P_n^{INV} edges corresponding to ARK_1 suffer from higher congestion. To improve the congestion, we "spread out" the flow through these edges by demonstrating a randomized way to generate $\mathrm{ARK}_1 \circ \mathrm{INV} \circ \mathrm{ARK}_1$ using 6 random keys that depend on the random variable $w \in \mathbb{F}_n$:

$$\begin{aligned}\mathrm{ARK}_{w+1} \circ \mathrm{INV} \circ \mathrm{ARK}_{w^{-1}} \circ \mathrm{INV} \circ \mathrm{ARK}_w \circ \mathrm{INV} \circ \\ \circ \mathrm{ARK}_{w^{-1}} \circ \mathrm{INV} \circ \mathrm{ARK}_w \circ \mathrm{INV} \circ \mathrm{ARK}_{w^{-1}+1}\,.\end{aligned}$$

One additional tweak to the first and last round keys is required to obtain a low-congestion set of P_n^{INV} paths that transpose the element 0 with an almost uniformly random non-zero element i and complete the comparison with $P_n^{\mathrm{2cyc},0}$.

Additional challenges come up when we move to fields of odd characteristic. This is because when the size of the field n is a prime congruent to 3 modulo 4, then the INV operation transposes every element except $\{-1,0,1\}$[4]. Thus it computes an even permutation, since it consists of $\frac{n-3}{2}$ transpositions. Moreover, the ARK_K operation is also an even permutation, since it is either the identity, or an odd-sized cycle for all values of $K \in \mathbb{F}_n$. We conclude that one cannot construct a transposition, which is an odd permutation, by composing INV and ARK. This is why our construction for the odd characteristic goes through 3-cycles, which are an even permutation. Our starting point is the observation that the following sequence of operations

$$\mathrm{ARK}_{-u-2v^{-1}} \circ \mathrm{INV} \circ \mathrm{ARK}\, v \circ \mathrm{INV} \circ \mathrm{ARK}_{-2v^{-1}} \circ \mathrm{INV} \circ \mathrm{ARK}_v \circ \mathrm{ARK}_u$$

generates the 3-cycle $(-u, -u-2v^{-1}, -u-v^{-1})$ for any u, and $v \neq 0$. This is precisely the subset of 3-cycles whose elements are terms of an arithmetic progression, as $(-u-v^{-1})-(-u) = (-u-2v^{-1}) - (-u-v^{-1})$. Furthermore, since every ARK operation in the above sequence has an almost uniformly random key, the resulting multi-commodity flow has low congestion and allows us to compare P_n^{INV} with $P_n^{\mathrm{3cyc,ap}}$.

It is worth noting that obtaining tight mixing time bounds typically happens using group-theoretic arguments [DS81] or entropy-related methods [LY98, FOW18, Sal20] for Markov chains that enjoy rich structure and symmetry, such as the random transpositions. The comparison method has been successfully used to transfer these tight bounds to other Markov chains that do not contain enough structure or symmetry to be analyzed using the more advanced mathematical techniques.

This paper is yet another example of this phenomenon, designing paths to relate the inverse walk to walks with more structure. Moreover, we have numerical simulations that support that the paths used in our arguments are not unnecessarily complicated. Our computations have shown that the paths from Lemmas 3.8 and 3.15 have the shortest possible length.

For the lower bound, we formalize an argument of [Nyb93, LTV21], which relies on the following observation: applying a sequence of ARK's and INV's to some input results in a rational function that is described by 3 coefficients (which coefficients depend on the secret keys of the cipher), unless one of the intermediate inputs to the INV becomes 0. Since the secret round keys are random, this happens with probability $\frac{1}{n}$ per round. Thus, unless the number of rounds scales linearly with n, 4 inputs will be enough to distinguish the value of these coefficients from random. This gives a lower bound on the number of rounds to reach 4-wise independence, which in turn is also a lower bound for convergence to any $t > 4$.

2 Preliminaries

For the entirety of this paper, our inputs and operations will be over a finite field $\mathbb{F}_n$, where n is a prime power p^b. For fields of odd characteristic, we use 2 to denote the sum of the multiplicative identity with itself. We denote by INV the inverse over the field that maps x to x^{n-2}. It holds that $\mathrm{INV}(x) \cdot x = 1$ for all non-zero x, and $\mathrm{INV}(0) = 0$. We will also define ARK_K to be the AddRoundKey operation with secret round key K that maps x to $x + K$. Moreover, we use the symbols $\gtrsim, \lesssim$ to compare two asymptotic quantities without specifying the constant factors.

2.1 The INV Key-Alternating Cipher

A Key-Alternating Cipher (KAC) is parameterized by a field size n, number of rounds r, and fixed permutation $P : \mathbb{F}_n \to \mathbb{F}_n$. A KAC is a family of functions indexed by $r+1$ sub-keys $K_0, K_1, \ldots, K_r$, and defined recursively as follows:

$$F_P^{(0)}(x) = x + K_0$$

[4] Here 1 is the multiplicative identity of the field, and 0 the additive identity.

$$F^{(i)}_{P,K_0,\ldots,K_i}(x) = P\left(F^{(i-1)}_{P,K_0,\ldots,K_{i-1}}(x)\right) + K_i.$$

The family of functions is

$$\mathcal{F}_P = \{F^{(r)}_{P,K_0,\ldots,K_r}(x) \mid K_i \in \mathbb{F}_n\}\ .$$

One can also naturally extend this to have different permutations in each round. In this paper, we consider the INV KAC, for which we use the INV map over $\mathbb{F}_n$ as the fixed permutation P.

Approximate t-wise independence. We say that a KAC $\mathcal{F}_P$ is ε-close t-wise independent if the images $(f(x_1),\ldots,f(x_t))$ of t arbitrary, distinct inputs $(x_1,\ldots,x_t)$ when f is drawn uniformly from $\mathcal{F}_P$ is at most ε away in total variation distance from a uniformly random t-tuple of distinct elements of $\mathbb{F}_n$. A KAC that is ε-close t-wise independent is also ε-close to t'-wise independent for any $t' < t$. Moreover, a KAC close to $(n-2)$-wise independent implies that the uniform distribution over the functions in the KAC family is close to a uniformly random even permutation.

2.2 Markov Chain Preliminaries

In this section, we recall some basic definitions of Markov chains, variational forms, and mixing time results. The interested reader may refer to [SC97, WLP09] for more details and proofs.

Let Π be the transition matrix of an ergodic Markov chain over a finite state space Ω, and let π denote its stationary distribution. We further use E to refer to the set of edges of the underlying graph, that is $E = \{(x,y) : \Pi(x,y) > 0\}$. We identify a Markov chain with its transition matrix, so we will often say that Π is both the transition matrix for a Markov chain and also the Markov chain itself.

Definition 2.1 (Reversible Markov chain)**.** *A Markov chain Π is* reversible *if for all $x,y \in \Omega$,*

$$\pi(x)\Pi(x,y) = \pi(y)\Pi(y,x).$$

Definition 2.2 (Mixing time)**.** *The ε-mixing time of an ergodic Markov chain Π is defined as the minimum number of steps t such that any random walk of length t is ε-close to the stationary distribution in total variation distance*

$$\tau_\varepsilon(\Pi) = \min_{t\geq 0}\left\{\max_{x\in\Omega}\|p_x^t - \pi\|_{TV} \leq \varepsilon\right\}.$$

2.3 Mixing times via functional inequalities

Functional inequalities have been very useful in studying the mixing times of Markov chains. The interested reader is encouraged to consult [LPW06, MT06] for a more thorough treatment.

On a very high and informal level, these functional inequalities capture how the "distance" from the stationary distribution decays with the number of steps in the Markov chain. The notion of "distance" is often instantiated by the χ^2 divergence (which gives rise to the variance in Definition 2.4) and the KL divergence (which gives rise to the entropy in Definition 2.3). One may think of the function f in the definitions below as a weighted probability distribution over the state space for a specific time step of our random walk.

Definition 2.3 (Entropy)**.** *The* entropy *of a function $f : \Omega \to \mathbb{R}_{\geq 0}$ with respect to Π is*

$$\mathrm{Ent}_\pi[f] = \sum_{x\in\Omega}\pi(x)f(x)\log\frac{f(x)}{\mathbb{E}_\pi[f]},$$

where $\mathbb{E}_\pi[f] = \sum_{x\in\Omega}\pi(x)f(x)$.

Definition 2.4 (Variance)**.** *The* variance *of a function* $f : \Omega \to \mathbb{R}_{\geq 0}$ *with respect to* Π *is*

$$\mathrm{Var}_\pi[f] = \sum_{x \in \Omega} \pi(x) \left(f(x) - \mathbb{E}_\pi[f]\right)^2,$$

where $\mathbb{E}_\pi[f] = \sum_{x \in \Omega} \pi(x) f(x)$*.*

It turns out that the way these two divergences decay at every time step is captured by the following Dirichlet form. Indeed, the definition of the Dirichlet form is sensitive to how the current probability distribution of our walk varies *locally*, whereas the entropy and variance above only give a *global* view of the current probability distribution of the walk.

Definition 2.5 (Dirichlet form)**.** *Let* $f : \Omega \to \mathbb{R}_{\geq 0}$ *be a function. The* Dirichlet form *of* f, g *with respect to* Π *is*

$$\mathcal{E}^\Pi(f, g) = \frac{1}{2} \sum_{x, y \in \Omega} \left(f(x) - f(y)\right)\left(g(x) - g(y)\right) \pi(x) \Pi(x, y).$$

The functional inequalities that we make use of in this paper can be written as the minimum ratio of the Dirichlet form of function f when divided by the entropy (or variance) of f. Keeping in mind that the Dirichlet form acts as the derivative of the entropy (and variance), can we view each of the constants that we introduce below as bounding the minimum amount of progress that we make per time step towards the stationary distribution in the respective distance.

Definition 2.6 (Log-Sobolev constant of Markov chain)**.** *The* log-Sobolev constant *of* Π *is defined by*

$$\alpha^\Pi = \inf_{\substack{f : \Omega \to \mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{\mathcal{E}^\Pi(f, f)}{\mathrm{Ent}_\pi[f^2]}.$$

Definition 2.7 (Modified log-Sobolev constant of Markov chain)**.** *The* modified log-Sobolev constant *of* Π *is defined by*

$$\beta^\Pi = \inf_{\substack{f : \Omega \to \mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{\mathcal{E}^\Pi(f, \log f)}{\mathrm{Ent}_\pi[f]}.$$

Definition 2.8 (Spectral gap of Markov chain)**.** *The* spectral gap *of* Π *is equal to*

$$\lambda^\Pi = \inf_{\substack{f : \Omega \to \mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{\mathcal{E}^\Pi(f, f)}{\mathrm{Var}_\pi[f]}.$$

In the Markov chain analysis literature, the spectral gap is often referred to as the *Poincaré constant*. Due to its relation to the spectral gap, and the fact that the spectral gap may be more familiar to a cryptography audience, we use the term spectral gap.

The functional inequalities satisfy the following inequalities between them.

Theorem 2.9 ([BT03])**.** *For every reversible Markov chain* Π*, the three functional inequalities satisfy*

$$4\alpha^\Pi \leq \beta^\Pi \leq 2\lambda^\Pi.$$

The above result implies that a lower bound on the log-Sobolev constant also implies the same bound for the remaining two constants. Moreover, we can transfer a modified log-Sobolev lower bound to a log-Sobolev lower bound using a recent result of [STY22].

Theorem 2.10 ([STY22], Theorem 1)**.** *For any reversible Markov chain* Π*, let* α *be its log-Sobolev constant, and* β *be its modified log-Sobolev constant. Moreover, let* p *be defined as follows*

$$p = \min_{\substack{x,y\in\Omega \\ \Pi(x,y)\neq 0}} \frac{\Pi(x,y)}{\max\left\{\sum_{y\neq x}\Pi(x,y), \Pi(y,x)\right\}}.$$

Then

$$\alpha \geq \frac{1}{20\log\frac{1}{p}} \cdot \beta.$$

Recall that the functional inequalities characterize the minimum progress towards the stationary distribution after every step of the random walk. A lower bound on these quantities can be directly translated to a mixing time upper bound, as we see in the following theorems.

Theorem 2.11 (Mixing time by log-Sobolev and spectral gap, [DSC96], Theorem 3.7)**.** *Let* Π *be the transition matrix of a reversible Markov chain whose stationary distribution is* π*, and* $\pi_{\min}$ *be the smallest stationary probability. For* $\varepsilon \leq \frac{1}{e}$*, the* ε*-mixing time is bounded by*

$$\tau_\varepsilon(\Pi) \leq \frac{1}{4\alpha^\Pi}\log\log\frac{1}{\pi_{min}} + \frac{1}{\lambda^\Pi}\cdot\log\frac{1}{2\varepsilon^2}$$

Theorem 2.12 (Mixing time by modified log-Sobolev, [BT03], Corollary 2.2)**.** *Let* Π *be the transition matrix of a reversible Markov chain whose stationary distribution is* π*, and* $\pi_{\min}$ *be the smallest stationary probability. For* $\varepsilon \leq \frac{1}{e}$*, the* ε*-mixing time is bounded by*

$$\tau_\varepsilon(\Pi) \leq \frac{1}{\beta^\Pi}\left(\log\log\frac{1}{\pi_{min}} + \log\frac{1}{2\varepsilon^2}\right)$$

Remark 2.13. *The mixing theorems above are often stated in the context of continuous-time Markov chains. That is, the random walk does not update at every integer time step, but rather it takes a step according to a random point process. A popular setting is the Poisson point process, in which the random walk takes its* k^{th} *step and then waits for* X_k *amount of time, where* X_k *is a random exponential variable.*

The connection between the discrete-time and the continuous-time settings is fairly well understood. In particular, Corollary 2.2 of [DSC96] implies that the continuous-time bounds of Theorems 2.11 and 2.12 also hold in the discrete-time setting for a lazy Markov chain. Making any Markov chain lazy, that is imposing that the random walk remains at the current state with probability $\frac{1}{2}$ *at each step will increase the functional inequalities by a factor of* 2 *(since the Dirichlet form is exactly halved) and thus worsen our mixing time bounds by a factor of* 2*.*

2.4 The INV KAC cipher as a Markov Chain

To study the t-wise independence of the INV KAC cipher, we will model its execution as a random walk over the alternating group of permutations over $\mathbb{F}_n$. Even though generating a truly random even permutation may initially seem too strong for t-wise independence, we will see (perhaps surprisingly) that the number of rounds required to reach 4-wise independence and $(n-2)$-wise independence are close (up to logarithmic factors). Thus considering convergence to the entire alternating group makes our analysis more convenient.

A first idea is to consider our block cipher to be a random walk over the alternating group A_n, where every step of the walk applies first the ARK_K operation with a random key K from $\mathbb{F}_n$, and then the INV operation. The main issue of representing the block cipher this way is that the underlying Markov chain is not reversible, and thus it is harder to apply the mixing time result of Theorem 2.11.

We will instead use the following reversible Markov chain to represent our cipher:

Definition 2.14 (INV KAC Markov chain)**.** *The chain P_n^{INV} on the alternating group A_n has the following transition matrix. Given the current even permutation σ_t, one step in this Markov chain corresponds to drawing uniformly and independently two random keys K_1, K_2 from $\mathbb{F}_n$, and setting*

$$\sigma_{t+1} = \mathrm{ARK}_{K_2} \circ \mathrm{INV} \circ \mathrm{ARK}_{K_1} \circ \sigma_t.$$

Note that the degree of the underlying graph has increased from n to n^2. It is not hard to observe that this transformation has not introduced any parallel edges to our Markov chain.

Lemma 2.15. *The Markov chain P_n^{INV} does not have any parallel edges for $n \geq 5$. That is, the following holds*

$$\mathrm{ARK}_i \circ \mathrm{INV} \circ \mathrm{ARK}_j = \mathrm{ARK}_k \circ \mathrm{INV} \circ \mathrm{ARK}_\ell,$$

only when $(i, j) = (k, \ell)$.

Proof. First, observe that if $i = k$, then the statement is true. Indeed, we can apply the permutation $\mathrm{INV} \circ \mathrm{ARK}_{-i}$ to both sides and obtain

$$\mathrm{ARK}_j = \mathrm{ARK}_\ell \implies j = \ell.$$

Similarly, if $j = \ell$, then the statement also holds. Hence, let us consider the case when both $i \neq k$, and $j \neq \ell$. For all $n - 2$ values of x such that $y = \sigma(x) \notin \{-j, -\ell\}$, x is mapped under the two permutations to equal values:

$$i + \frac{1}{y+j} = k + \frac{1}{y+\ell}$$

Therefore,

$$\begin{aligned} & \frac{iy + ij + 1}{y+j} = \frac{ky + k\ell + 1}{y + \ell} \\ \implies \ & iy^2 + (ij + 1 + i\ell)y + ij\ell + \ell = ky^2 + (k\ell + 1 + kj)y + jk\ell + j \\ \implies \ & (i-k)y^2 + \left((i-k)j + (i-k)\ell\right)y + (i-k)j\ell + (\ell - j) = 0. \end{aligned}$$

For the above equality to be true for more than 2 values of y, it must hold that $i = k$. This concludes the proof for $n \geq 5$. □

2.5 Functional inequalities on some well-known Markov chains

We will bound the mixing time of the INV KAC Markov chain by comparing it to two well-studied Markov chains. In the first case, when $\mathbb{F}_n$ is of characteristic 2, we will use the random transposition Markov chain.

Definition 2.16 (Random transposition Markov chain)**.** *The chain $P_n^{2\mathrm{cyc}}$ on the symmetric group S_n has the following transition matrix. Given the current permutation σ_t, one step in this Markov chain corresponds to drawing uniformly a random transposition (i, j) from S_n, and setting*

$$\sigma_{t+1} = (i, j) \circ \sigma_t.$$

Prior work has obtained tight estimates for the functional inequalities of this chain.

Theorem 2.17 ([Sal20], Theorem 5)**.** *The log-Sobolev constant of the random transposition chain satisfies*

$$\frac{\log 2}{2(n-1)\log n} \leq \alpha_n^{2\mathrm{cyc}} \leq \frac{2}{(n-1)\log n}.$$

Theorem 2.18 ([Goe04], Corollary 3.1)**.** *The modified log-Sobolev constant of the random transposition Markov chain satisfies*

$$\frac{1}{n-1} \leq \beta_n^{\text{2cyc}} \leq \frac{4}{n-1}.$$

Using the monotonicity of the functional inequalities from Theorem 2.9 we can obtain a lower bound for the spectral gap bound from the modified log-Sobolev constant.

Corollary 2.19 (Spectral gap of the random transposition chain)**.** *The spectral gap of the random transposition Markov chain satisfies* $\lambda_n^{\text{2cyc}} \geq \frac{1}{2n}$.

In the second case, when $\mathbb{F}_n$ is of odd characteristic, we will use the random 3-cycle Markov chain.

Definition 2.20 (Random 3-cycle Markov chain)**.** *The chain* P_n^{3cyc} *on the alternating group* A_n *has the following transition matrix. Given the current permutation* σ_t*, one step in this Markov chain corresponds to drawing uniformly a random* 3*-cycle* (i, j, k) *from* A_n*, and setting*

$$\sigma_{t+1} = (i, j, k) \circ \sigma_t.$$

The underlying graph of P_n^{3cyc} is $2\binom{n}{3}$-regular and thus $P_n^{\text{3cyc,ap}}(x, y) = \frac{1}{2\binom{n}{3}}$ for $(x, y) \in E_n^{\text{3cyc}}$.

Theorem 2.21 ([Goe04], Corollary 3.2)**.** *The modified log-Sobolev constant of the* 3*-cycle Markov chain satisfies*

$$\frac{1}{n-2} \leq \beta_n^{\text{3cyc}} \leq \frac{6}{n-1}.$$

Corollary 2.22 (Log-Sobolev constant of the 3-cycle chain)**.** *The log-Sobolev constant of the* 3*-cycle Markov chain satisfies*

$$\alpha_n^{\text{3cyc}} \geq \frac{1}{60n \log n}.$$

Proof. We use Theorem 2.10 to transfer our modified log-Sobolev bound to a log-Sobolev lower bound. For the 3-cycle Markov chain, it holds that $p = \frac{1}{\binom{n}{3}} \geq \frac{1}{n^3}$. Thus

$$\alpha_n^{\text{3cyc}} \geq \frac{1}{20 \log n^3} \cdot \frac{1}{n-2} = \frac{1}{60(n-2) \log n}. \qquad \square$$

Again the monotonicity of the functional inequalities implies a spectral gap lower bound.

Corollary 2.23 (Spectral gap of the 3-cycle chain)**.** *The spectral gap of the* 3*-cycle Markov chain satisfies* $\lambda_n^{\text{3cyc}} \geq \frac{1}{2n}$.

2.6 The comparison method

Below we sketch the "comparison with multicommodity flows" method of Sinclair, Diaconis and Saloff-Coste [Sin92, DSC93].

Let P^{ref} and P^{tar} be two reversible Markov chains on the same ground set, with stationary distributions $\pi^{\text{ref}}, \pi^{\text{tar}}$ and edge sets $E^{\text{ref}}, E^{\text{tar}}$ respectively. We will think of P^{ref} as the "reference" chain for which we have somehow obtained estimates for its log-Sobolev constant. Our goal is to bound the log-Sobolev constant of the "target" chain P^{tar}, by relating it to that of P^{ref}.

Define a path γ_{xy} for $(x, y) \in E^{\text{ref}}$ to be a sequence of steps

$$(x = a_0, a_1, \ldots, a_k = y)$$

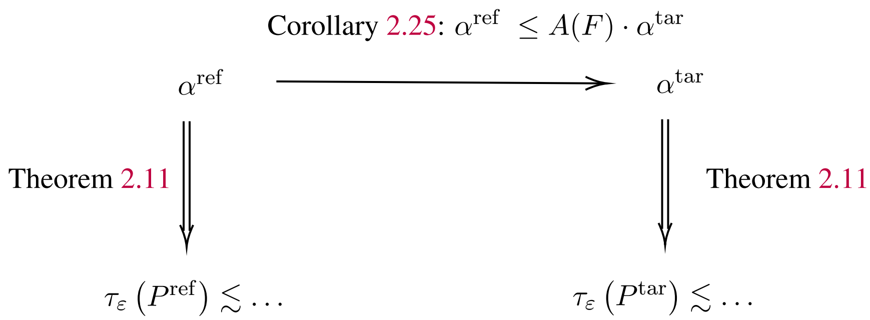

Figure 1: Schematic representation of the comparison method for the case of the log-Sobolev constant. The comparison method of Corollary 2.25 allows one to compare the log-Sobolev constant of a "target" Markov chain P^{tar} with the known constant of a "reference" chain P^{ref}, as long as they have the same stationary distribution. The resulting log-Sobolev bound α^{tar} then implies a mixing time bound $\tau_\varepsilon(P^{\text{tar}})$ using Theorem 2.11. On a high level, the comparison method allows one to transfer mixing time bounds from a well-studied chain P^{ref} to a new Markov chain P^{tar}, as long as we can construct a valid multi-commodity flow with low congestion.

in the target chain P^{tar}. For this to be a valid path, it must hold that $(a_i, a_{i+1}) \in E^{\text{tar}}$. We say that such a path has length $|\gamma_{xy}| = k$. Let $\mathcal{P}_{xy}$ be the set of all simple paths connecting x to y. Also let $\mathcal{P} = \cup_{(x,y)\in E^{\text{ref}}} \mathcal{P}_{xy}$ be the union of all such paths. For $(a,b) \in E^{\text{tar}}$, let $\mathcal{P}(a,b) = \{\gamma \in \mathcal{P} \mid (a,b) \in \gamma\}$. That is, $\mathcal{P}(a,b)$ contains all paths that use the edge (a,b) of the target graph.

A function $F : \mathcal{P} \to \mathbb{R}_{\geq 0}$ is called a $(P^{\text{tar}}, P^{\text{ref}})$-flow if

$$\sum_{\gamma\in\mathcal{P}_{xy}} F(\gamma) = P^{\text{ref}}(x,y)\pi^{\text{ref}}(x).$$

Theorem 2.24 ([DSC93], Theorem 2.3)**.** *Let P^{tar}, P^{ref} be reversible Markov chains on a finite set Ω. For any $(P^{\text{tar}}, P^{\text{ref}})$-flow F, the Dirichlet forms satisfy*

$$\mathcal{E}^{\text{ref}}(f,f) \leq A(F) \cdot \mathcal{E}^{\text{tar}}(f,f)$$

with

$$A(F) = \max_{(a,b)\in E^{\text{tar}}} \left\{ \frac{1}{\pi^{\text{tar}}(a) P^{\text{tar}}(a,b)} \sum_{\gamma\in\mathcal{P}(a,b)} |\gamma| \cdot F(\gamma) \right\}.$$

Moreover, when these two Markov chains have the same stationary distribution, the theorem above directly implies a relation between their log-Sobolev constants.

Corollary 2.25. *Let P^{tar} and P^{ref} be reversible Markov chains on a finite set Ω with the same stationary distribution π. For any $(P^{\text{tar}}, P^{\text{ref}})$-flow F, their log-Sobolev constants and spectral gaps satisfy*

$$\alpha^{\text{ref}} \leq A(F) \cdot \alpha^{\text{tar}}, \quad \lambda^{\text{ref}} \leq A(F) \cdot \lambda^{\text{tar}}.$$

Proof. Since they have the same stationary distribution, the denominator in the definitions of the log-Sobolev and spectral gap constants is equal. Then

$$\alpha^{\text{ref}} = \inf_{\substack{f:\Omega\to\mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{\mathcal{E}^{\text{ref}}(f,f)}{\text{Ent}_\pi[f^2]} \leq \inf_{\substack{f:\Omega\to\mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{A(F)\cdot\mathcal{E}^{\text{tar}}(f,f)}{\text{Ent}_\pi[f^2]} = A(F)\cdot\alpha^{\text{tar}}.$$

The same inequality holds for λ^{ref}. ☐

Comparison method for modified log-Sobolev. Due to the use of $\log f$ in Definition 2.7 for the modified log-Sobolev inequality, the comparison method does not apply directly to that functional inequality. A tool developed in a recent paper by Liu et al. [LMR+24] provides a way to deal with that logarithmic factor and almost exactly recover the same comparison theorem for the modified log-Sobolev inequality as well. We formally describe the modified comparison method below.

The crucial ingredient, in that case, is Equation (1) that allows us to write one term of the Dirichlet form in terms of its path decomposition. This parallels the use of the Cauchy-Schwarz inequality in the original proof of the comparison method.

Claim 2.26 (Proof of Lemma 3.7, [LMR+24]). *Let P be a reversible Markov chain on a finite set Ω. Further, let $x, y \in \Omega$ be two vertices of its state space, and γ be a path on Ω that connects x and y. Then it holds that*

$$(f(x) - f(y)) \cdot \log \frac{f(x)}{f(y)} \leq 4|\gamma|^2 \cdot \sum_{(a,b)\in\gamma} (f(a) - f(b)) \cdot \log \frac{f(a)}{f(b)}. \tag{1}$$

Proof. In their proof of Lemma 3.7, the authors show that the inequality holds for $|\gamma| = 2$. That is,

$$(f(x) - f(y)) \cdot \log \frac{f(x)}{f(y)} \leq 4 \left[(f(x) - f(a)) \cdot \log \frac{f(x)}{f(a)} + (f(a) - f(y)) \cdot \log \frac{f(a)}{f(y)}\right].$$

Our proof proceeds by induction on the size of $|\gamma|$. We first show that this holds when $|\gamma|$ is a power of 2. Indeed, if $|\gamma| = 2^k$, let a be the midpoint of the path from x to y. Then the $|\gamma| = 2$ case implies that:

$$(f(x) - f(y)) \cdot \log \frac{f(x)}{f(y)} \leq 4 \left[(f(x) - f(a)) \cdot \log \frac{f(x)}{f(a)} + (f(a) - f(y)) \cdot \log \frac{f(a)}{f(y)}\right].$$

Now each of x to a and a to y are paths of length 2^{k-1}, and thus

$$\begin{aligned}
&(f(x) - f(y)) \cdot \log \frac{f(x)}{f(y)} \\
&\leq 4 \left[2^{2(k-1)} \cdot \sum_{(u,v)\in\gamma(x,a)} (f(u) - f(v)) \cdot \log \frac{f(u)}{f(v)} + 2^{2(k-1)} \cdot \sum_{(u,v)\in\gamma(a,y)} (f(u) - f(v)) \cdot \log \frac{f(u)}{f(v)}\right] \\
&\leq 2^{2k} \sum_{(u,v)\in\gamma(x,y)} (f(u) - f(v)) \cdot \log \frac{f(u)}{f(v)} = |\gamma|^2 \sum_{(u,v)\in\gamma(x,y)} (f(u) - f(v)) \cdot \log \frac{f(u)}{f(v)}.
\end{aligned}$$

The proof follows for all lengths $|\gamma|$ (not necessarily powers of 2) by padding the path up to the next power of two (for example by repeating the last vertex). This incurs at most a factor of 2 increase in the length of the path and hence the factor of 4 in the claim statement. □

Theorem 2.27. *Let P^{tar}, P^{ref} be reversible Markov chains on a finite set Ω. For any $(P^{\mathrm{tar}}, P^{\mathrm{ref}})$-flow F, the Dirichlet forms satisfy*

$$\mathcal{E}^{\mathrm{ref}}(f, \log f) \leq A'(F) \cdot \mathcal{E}^{\mathrm{tar}}(f, \log f)$$

with

$$A'(F) = \max_{(a,b)\in E^{\mathrm{tar}}} \left\{ \frac{4}{\pi^{\mathrm{tar}}(a) P^{\mathrm{tar}}(a,b)} \sum_{\gamma\in\mathcal{P}(a,b)} |\gamma|^2 \cdot F(\gamma) \right\}$$

Proof. The proof follows the original proof of the comparison theorem, but substitutes Equation (1) in place of the Cauchy-Schwarz inequality.

$$
\begin{aligned}
&\mathcal{E}^{\text{ref}}(f, \log f) \\
&= \frac{1}{2} \sum_{x,y \in \Omega} (f(x) - f(y)) \cdot \log \frac{f(x)}{f(y)} \cdot \pi^{\text{ref}}(x) P^{\text{ref}}(x,y) \\
&= \frac{1}{2} \sum_{x,y \in \Omega} \left(\sum_{\gamma \in \mathcal{P}_{xy}} (f(x) - f(y)) \cdot \log \frac{f(x)}{f(y)} \cdot F(\gamma) \right) \\
&\leq 4 \cdot \frac{1}{2} \sum_{x,y \in \Omega} \left[\sum_{\gamma \in \mathcal{P}_{xy}} \sum_{(a,b) \in \gamma} (f(a) - f(b)) \log \frac{f(a)}{f(b)} \cdot |\gamma|^2 \cdot F(\gamma) \right] && \text{(Equation (1))} \\
&= \frac{1}{2} \sum_{(a,b) \in E^{\text{tar}}} (f(a) - f(b)) \log \frac{f(a)}{f(b)} \pi^{\text{tar}}(a) P^{\text{tar}}(a,b) \left[\frac{4}{\pi^{\text{tar}}(a) P^{\text{tar}}(a,b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma|^2 \cdot F(\gamma) \right] \\
&\leq A'(F) \cdot \mathcal{E}^{\text{tar}}(f, \log f).
\end{aligned}
$$

□

Corollary 2.28. *Let P^{tar} and P^{ref} be reversible Markov chains on a finite set Ω with the same stationary distribution π. For any $(P^{\text{tar}}, P^{\text{ref}})$-flow F supported on paths of length at most $\gamma_{\max}$, their modified log-Sobolev constants satisfy*

$$\beta^{\text{ref}} \leq A'(F) \cdot \beta^{\text{tar}} \leq \gamma_{\max} \cdot A(F) \cdot \beta^{\text{tar}}.$$

Proof. We first show that $A'(F) \leq \gamma_{\max} \cdot A(F)$. This follows directly from the definition of congestion

$$
\begin{aligned}
A'(F) &= \max_{(a,b) \in E^{\text{tar}}} \left\{ \frac{4}{\pi^{\text{tar}}(a) P^{\text{tar}}(a,b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma|^2 \cdot F(\gamma) \right\} \\
&\leq \gamma_{\max} \max_{(a,b) \in E^{\text{tar}}} \left\{ \frac{4}{\pi^{\text{tar}}(a) P^{\text{tar}}(a,b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma| \cdot F(\gamma) \right\} \\
&= \gamma_{\max} \cdot A(F).
\end{aligned}
$$

Now, since the two Markov chains have the same stationary distribution, the denominator in the definition of the modified log-Sobolev chain is equal. Then

$$\beta^{\text{ref}} = \inf_{\substack{f:\Omega \to \mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{\mathcal{E}^{\text{ref}}(f, \log f)}{\text{Ent}_\pi[f]} \leq \inf_{\substack{f:\Omega \to \mathbb{R}_{\geq 0} \\ f \text{ non-constant}}} \frac{A'(F) \cdot \mathcal{E}^{\text{tar}}(f, \log f)}{\text{Ent}_\pi[f]} = A'(F) \cdot \beta^{\text{tar}}.$$

□

Representing edges and paths. To employ the comparison method, we will need a way to specify paths in the P_n^{INV} Markov chain. Since each edge of P_n^{INV} is determined by two keys r_1, r_2, we will use the notation $[[r_1, r_2]]$ to denote the edge $\text{ARK}_{r_2} \circ \text{INV} \circ \text{ARK}_{r_1}$. The starting vertex will be specified separately.

Whenever we need to describe a longer path of P_n^{INV}, we will write it as a tuple of double square brackets, by specifying the first edge first and so on, e.g. $([[r_1, r_2]], [[q_1, q_2]])$[5]. Since ARK operations form a subgroup, this path is also equal to

$$\begin{aligned}(\mathrm{ARK}_{q_2} \circ \mathrm{INV} \circ \mathrm{ARK}_{q_1}) \circ (\mathrm{ARK}_{r_2} \circ \mathrm{INV} \circ \mathrm{ARK}_{r_1})\\ = \mathrm{ARK}_{q_2} \circ \mathrm{INV} \circ \mathrm{ARK}_{q_1+r_2} \circ \mathrm{INV} \circ \mathrm{ARK}_{r_1} .\end{aligned}$$

We will also use the notation $[[r_1, r_2 + q_1, q_2]]$ to describe the above path in P_n^{INV}. In general, we will extend the double square bracket notation to mean:

$$[[r_1, r_2, \dots, r_k]] = \mathrm{ARK}_{r_k} \circ \mathrm{INV} \circ \cdots \circ \mathrm{ARK}_{r_2} \circ \mathrm{INV} \circ \mathrm{ARK}_{r_1} .$$

3 Upper Bound

In this section, we formally prove our claim that a random S-box over $\mathbb{F}_n$ can be approximated via the sequential composition of alternating AddRoundKey and INV S-box operations. Our proof is different for fields of characteristic 2 and fields of odd characteristic.

We show that $\alpha_n^{\mathrm{INV}} \geq \Omega\left(\frac{1}{n \log n}\right)$ and $\beta_n^{\mathrm{INV}}, \lambda_n^{\mathrm{INV}} \geq \Omega\left(\frac{1}{n}\right)$ for all fields $\mathbb{F}_n$ with $n \geq 5$. This log-Sobolev constant bound implies a mixing time bound from Theorem 2.11, and thus the following theorem.

Theorem 3.1. *Let $n \geq 5$. The* INV *KAC over the field of size n with $r \leq M_n \cdot n(2 \log n + \log \frac{1}{2\varepsilon^2})$ rounds generates a permutation that is ε-close to a uniformly random even permutation (equivalently a uniformly random permutation from the alternating group A_n). The constant M_n satisfies*

$$M_n = \begin{cases} 9130 \cdot 54 & n \textit{ is power of } 2 \\ 96 \cdot 729 \cdot 2 & n \textit{ is power of odd prime}. \end{cases}$$

Proof. In Lemmas 3.3 and 3.11 we show that the modified log-Sobolev constant of P_n^{INV} satisfies

$$\beta_n^{\mathrm{INV}} \geq \begin{cases} \frac{1}{9130 \cdot 54 \cdot n} & n \text{ is power of } 2 \\ \frac{1}{96 \cdot 729 \cdot 2n} & n \text{ is power of odd prime} \end{cases} = \frac{1}{M_n \cdot n}.$$

Thus Theorem 2.12 implies:

$$\tau_\varepsilon(P_n^{\mathrm{INV}}) \leq M_n \cdot n \left(\log \log n! + \log \frac{1}{2\varepsilon^2}\right) \leq M_n \cdot n \left(2 \log n + \log \frac{1}{2\varepsilon^2}\right).$$

Thus after this many rounds, the distribution of permutations computed by the INV KAC is ε-close to a uniformly random permutation from A_n. □

The log-Sobolev constant and spectral gap also imply a different mixing time result that is asymptotically weaker but with slightly better constants. It turns out that for fields of characteristic 2 of size smaller than $2^{2 \cdot M_n / L_n} \approx 2^{47}$, Theorem 3.2 provides a stronger bound than Theorem 3.1.

Theorem 3.2. *Let $n \geq 5$. The* INV *KAC over the field of size n with $r \leq n(L_n/2 \cdot \log^2 n + S_n \log \frac{1}{2\varepsilon^2})$ rounds generates a permutation that is ε-close to a uniformly random even permutation (equivalently a uniformly random permutation from the alternating group A_n). The constants L_n and S_n satisfy*

$$L_n = \begin{cases} \frac{830 \cdot 18 \cdot 2}{\log 2} & n \textit{ is power of } 2 \\ 24 \cdot 81 \cdot 60 & n \textit{ is power of odd prime} \end{cases}$$

[5]If we are writing the edges as permutations, we write the permutations corresponding to the edges from right-to-left.

$$S_n = \begin{cases} 830 \cdot 18 \cdot 2 & n \text{ is power of } 2 \\ 24 \cdot 81 \cdot 2 & n \text{ is power of odd prime.} \end{cases}$$

Proof. In Lemmas 3.3 and 3.11 we show that the log-Sobolev constant and spectral gap of P_n^{INV} satisfy

$$\alpha_n^{\mathrm{INV}} \geq \begin{cases} \frac{\log 2}{830 \cdot 18 \cdot 2n \log n} & n \text{ is power of } 2 \\ \frac{1}{24 \cdot 81 \cdot 60 \cdot n \log n} & n \text{ is power of odd prime} \end{cases} = \frac{1}{L_n \cdot n \log n},$$

$$\lambda_n^{\mathrm{INV}} \geq \begin{cases} \frac{1}{830 \cdot 18 \cdot 2n} & n \text{ is power of } 2 \\ \frac{1}{24 \cdot 81 \cdot 2n} & n \text{ is power of odd prime} \end{cases} = \frac{1}{S_n \cdot n}.$$

Thus Theorem 2.11 implies:

$$\begin{aligned} \tau_\varepsilon(P_n^{\mathrm{INV}}) &\leq \frac{L_n}{4} \cdot n \log n \cdot \log \log n! + S_n \cdot n \log \frac{1}{2\varepsilon^2} \\ &\leq n \left(\frac{L_n}{2} \cdot \log^2 n + S_n \log \frac{1}{2\varepsilon^2} \right). \end{aligned}$$

Thus after this many rounds, the distribution of permutations computed by the INV KAC is ε-close to a uniformly random permutation from A_n. ☐

3.1 Fields of characteristic 2

The main result of this section is a bound on the log-Sobolev constant of P_n^{INV} when n is a power of 2.

Lemma 3.3. *Let $n = 2^b$. The log-Sobolev constant, modified log-Sobolev constant, and spectral gap of P_n^{INV} satisfies*

- *Log-Sobolev*
$$\alpha_n^{\mathrm{INV}} \geq \frac{\log 2}{830 \cdot 18 \cdot 2n \log n}$$
- *Modified log-Sobolev*
$$\beta^{\mathrm{INV}} \geq \frac{1}{9130 \cdot 54 \cdot n}$$
- *Spectral gap*
$$\lambda_n^{\mathrm{INV}} \geq \frac{1}{830 \cdot 18 \cdot 2n}$$

Proof. Our comparison method proceeds in two steps. We introduce an intermediate chain $P_n^{\mathrm{2cyc},0}$, a slight variant of the random transposition chain, and compare the functional inequalities of P_n^{INV} with it. Then we bound the functional inequalities of $P_n^{\mathrm{2cyc},0}$ by comparing it to the random transposition chain P_n^{2cyc}. The proof follows directly from the comparison method in Lemmas 3.5 and 3.6 and known bounds on the functional inequalities of the random transpositions chain.

Log-Sobolev.

$$\alpha_n^{\mathrm{INV}} \underset{\text{Lemma 3.6}}{\geq} \frac{\alpha_n^{\mathrm{2cyc},0}}{830} \underset{\text{Lemma 3.5}}{\geq} \frac{\alpha_n^{\mathrm{2cyc}}}{830 \cdot 18} \underset{\text{Theorem 2.17}}{\geq} \frac{\log 2}{830 \cdot 18 \cdot 2 \cdot n \log n}.$$

Modified log-Sobolev.

$$\beta_n^{\text{INV}} \underset{\text{Lemma 3.6}}{\geq} \frac{\beta_n^{\text{2cyc},0}}{9130} \underset{\text{Lemma 3.5}}{\geq} \frac{\beta_n^{\text{2cyc}}}{9130 \cdot 54} \underset{\text{Theorem 2.18}}{\geq} \frac{1}{9130 \cdot 54 \cdot n}.$$

Spectral gap.

$$\lambda_n^{\text{INV}} \underset{\text{Lemma 3.6}}{\geq} \frac{\lambda_n^{\text{2cyc},0}}{830} \underset{\text{Lemma 3.5}}{\geq} \frac{\lambda_n^{\text{2cyc}}}{830 \cdot 18} \underset{\text{Corollary 2.19}}{\geq} \frac{1}{830 \cdot 18 \cdot 2n}.$$

□

Definition 3.4 (Transposition with fixed element Markov chain)**.** *The chain $P_n^{\text{2cyc},0}$ on the alternating group A_n has the following transition matrix. Given the current permutation σ_t, one step in this Markov chain corresponds to drawing a uniformly random non-zero element i from $[n]$. Then set*

$$\sigma_{t+1} = (0, i) \circ \sigma_t.$$

Lemma 3.5. *The functional inequalities of $P_n^{\text{2cyc},0}$ satisfy*

- *Log-Sobolev*
$$\alpha_n^{\text{2cyc},0} \geq \frac{\alpha_n^{\text{2cyc}}}{18}.$$
- *Modified log-Sobolev*
$$\beta_n^{\text{2cyc},0} \geq \frac{\beta_n^{\text{2cyc}}}{54}.$$
- *Spectral gap*
$$\lambda_n^{\text{2cyc},0} \geq \frac{\lambda_n^{\text{2cyc}}}{18}.$$

Proof. We will assign exactly one path to every edge $(x, y) \in E_n^{\text{2cyc}}$. Let $y = (i, j) \circ x$ for some $i \neq j$. Then we will set $\mathcal{P}_{xy} = \{\gamma_{xy}\}$, where γ_{xy} is chosen according to the following two cases:

1. One of i, j is equal to zero. Then (i, j) is also in $E_n^{\text{2cyc},0}$ and we set
$$\gamma_{xy} = (x, (i, j) \circ x = y).$$
2. Both i, j are non-zero. Then we set
$$\gamma_{xy} = (x, \underbrace{(0, i) \circ x}_{a_1}, \underbrace{(0, j) \circ a_1}_{a_2}, (0, i) \circ a_2 = y).$$

We will assign to each path the same flow $F(\gamma_{xy}) = P_n^{\text{2cyc}}(x, y) \cdot \pi_n^{\text{2cyc}}(x) = \frac{2}{n!\binom{n}{2}} = \frac{4}{n(n-1) \cdot n!}$.

The comparison constant we get is

$$\begin{aligned}
A(F) &= \max_{(a,b) \in E^{\text{2cyc},0}} \left\{ \frac{1}{\pi^{\text{2cyc},0}(a) \cdot P_n^{\text{2cyc},0}(a, b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma| \cdot F(\gamma) \right\} \\
&= \max_{(a,b) \in E^{\text{2cyc},0}} \left\{ \frac{n!(n-1)}{2} \cdot 3 \cdot |\mathcal{P}(a, b)| \cdot \frac{4}{n(n-1) \cdot n!} \right\} \\
&\leq \frac{n!(n-1)}{2} \cdot 9(n-1) \cdot \frac{4}{n(n-1) \cdot n!} \\
&\leq 18.
\end{aligned}$$

We used the fact that the number of paths γ in $\mathcal{P}(a,b)$ is at most $3(n-1)$. This is because $\mathcal{P}(a,b)$ contains paths of length 1 and 3. The set $\mathcal{P}(a,b)$ contains at most 1 path of length 1.

We bound the number of length-3 paths in $|\mathcal{P}(a,b)|$ by $3(n-2)$ in the following way. The edge (a,b) specifies a unique ℓ such that $b=(0,\ell)\circ a$. Consider a length-3 path γ_{xy} that uses edge (a,b), where $y=(i,j)\circ x$. This edge can be one of 3 edges of γ_{xy}, and the position of (a,b) in the path specifies one of i,j to be equal to ℓ. Thus the remaining variable has $n-2$ possible values (except 0 and ℓ). The total number of paths is at most $1+3(n-2)\leq 3(n-1)$.

The modified log-Sobolev bound follows since we use paths of length at most 3, i.e. $\gamma_{\max}=3$. □

Lemma 3.6. *Let $n=2^b\geq 8$. The functional inequalities of P_n^{INV} satisfy*

- *Log-Sobolev*

$$\alpha_n^{\mathrm{INV}}\geq\frac{\alpha_n^{\mathrm{2cyc},0}}{830}.$$

- *Modified log-Sobolev*

$$\beta_n^{\mathrm{INV}}\geq\frac{\beta_n^{\mathrm{2cyc},0}}{9130}.$$

- *Spectral gap*

$$\lambda_n^{\mathrm{INV}}\geq\frac{\lambda_n^{\mathrm{2cyc},0}}{830}.$$

The proof uses the following way to generate the transposition $(0,\frac{v}{uv+1}+\frac{v+1}{uv+u+1})$ by combining INV and ARK operations.

Lemma 3.7. *For any $u,v,w\in\mathbb{F}_n$ such that $uv\neq 1$, $u(v+1)\neq 1$, and $v\notin\{0,1\}$, we can generate the transposition $\left(0,\frac{v}{uv+1}+\frac{v+1}{uv+u+1}\right)$ using the following sequence $\psi_{\mathrm{even}}(u,v,w)$ of* ARK *and* INV *operations*

$$\psi_{\mathrm{even}}(u,v,w)=\left[\left[\frac{v}{uv+1},\ u,\ v,\ \mathrm{INV}(w)+1,\ w,\right.\right.$$
$$\left.\left.\mathrm{INV}(w),\ w,\ \mathrm{INV}(w),\ w+1,\ v+1,\ u,\ \frac{v}{uv+1}\right]\right].$$

Our proof of Lemma 3.7 will follow from Lemmas 3.8 and 3.9.

Lemma 3.8. *For any $u,v\in\mathbb{F}_n$ such that $uv\neq 1$, $u(v+1)\neq 1$, and $v\notin\{0,1\}$, we can generate the transposition $\left(u+\frac{v}{uv+1},u+\frac{v+1}{uv+u+1}\right)$ using the following sequence $\gamma(u,v)$ of* AddRoundKey *and* INV *S-box operations*

$$\gamma(u,v)=[[u,\ u,\ v,\ 1,\ 1,\ v+1,\ u,\ u]]\,.$$

Proof. We will consider the following cases.

Case 1. $u=0$**:** We will show that the sequence $\gamma(0,v)=[[0,\ 0,\ v,\ 1,\ 1,\ v+1,\ 0,\ 0]]$ generates the transposition $(v,v+1)$.

We consider first the application of π on some x that is not equal to v or $v+1$.

$$\begin{aligned}
x &\xrightarrow{\mathrm{ARK}_0} x \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(x) \\
&\xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} x+v \xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} \frac{1}{x+v}+1 = \frac{x+v+1}{x+v} \\
&\xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} \frac{x+v}{x+v+1}+1 = \frac{1}{x+v+1} \xrightarrow{\mathrm{ARK}_{v+1} \circ \mathrm{INV}} x \\
&\xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(x) \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} x.
\end{aligned}$$

Now consider what happens when $x = v$:

$$\begin{aligned}
v &\xrightarrow{\mathrm{ARK}_0} v \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(v) \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} 0 \\
&\xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} 1 \xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} 0 \xrightarrow{\mathrm{ARK}_{v+1} \circ \mathrm{INV}} v+1 \\
&\xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(v+1) \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} v+1.
\end{aligned}$$

And when $x = v+1$:

$$\begin{aligned}
v+1 &\xrightarrow{\mathrm{ARK}_0} v+1 \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(v+1) \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} 1 \xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} 0 \\
&\xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} 1 \xrightarrow{\mathrm{ARK}_{v+1} \circ \mathrm{INV}} v \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(v) \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} v.
\end{aligned}$$

Case 2. We will now prove the result for all valid parameters $u \neq 0, v$ and inputs x that do not cause any input to $\mathrm{INV}(\cdot)$ to vanish. Thus, for these calculations, we will use the fact that $y \cdot \mathrm{INV}(y) = 1$ for all non-zero y. Also for brevity, we will use $\mathrm{INV}(y)$ and $\frac{1}{y}$ interchangeably.

$$
\begin{aligned}
& x \\
\xrightarrow{\mathrm{ARK}_u}\ & x+u \\
\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}}\ & \frac{1}{x+u}+u=\frac{ux+u^2+1}{x+u} \\
\xrightarrow{\mathrm{ARK}_v\circ\mathrm{INV}}\ & \frac{x+u}{ux+u^2+1}+v=\frac{(uv+1)x+u^2v+v+u}{ux+u^2+1} \\
\xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}}\ & \frac{ux+u^2+1}{(uv+1)x+u^2v+v+u}+1 \\
& =\frac{(uv+u+1)x+u^2v+v+u+u^2+1}{(uv+1)x+u^2v+v+u} \\
\xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}}\ & \frac{(uv+1)x+u^2v+v+u}{(uv+u+1)x+u^2v+v+u+u^2+1}+1 \\
& =\frac{ux+u^2+1}{(uv+u+1)x+u^2v+v+u+u^2+1} \\
\xrightarrow{\mathrm{ARK}_{v+1}\circ\mathrm{INV}}\ & \frac{(uv+u+1)x+u^2v+v+u+u^2+1}{ux+u^2+1}+v+1=\frac{x+u}{ux+u^2+1} \\
\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}}\ & \frac{ux+u^2+1}{x+u}+u=\frac{1}{x+u} \\
\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}}\ & x+u+u=x
\end{aligned}
$$

So we have seen that for u, v, x such that no input to $\mathrm{INV}(\cdot)$ is zero, $\gamma(u, v)$ acts like the identity and maps x to itself. To complete our proof, we now consider what happens if some input to $\mathrm{INV}(\cdot)$ equals 0. This happens when one of the following equalities hold:

(a) $x+u=0 \implies x=u$.

(b) $ux+u^2+1=0 \implies ux=u^2+1 \implies x=u+\frac{1}{u}$, since the equality doesn't hold if $u=0$.

(c) $(uv+1)x+u^2v+v+u=0 \implies x=u+\frac{v}{uv+1}$, since we have imposed that $uv\neq 1$.

(d) $(uv+u+1)x+u^2v+v+u+u^2+1=0 \implies x=u+\frac{v+1}{uv+u+1}$, since we have imposed that $u(v+1)\neq 1$.

Note that the third and fourth cases are the claimed non-fixed points of $\gamma(u, v)$. Looking forward, we will verify that $\gamma(u, v)$ transposes these two inputs.

Case 2(a). $x=u\neq 0$: The permutation $\gamma(u, v)$ maps u to itself as we show below:

$$
\begin{aligned}
u &\xrightarrow{\mathrm{ARK}_u} u+u=0 \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} 0+u=u \xrightarrow{\mathrm{ARK}_v\circ\mathrm{INV}} \frac{1}{u}+v=\frac{uv+1}{u} \\
&\xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} \frac{u}{uv+1}+1=\frac{uv+u+1}{uv+1} \xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} \frac{u}{uv+u+1} \\
&\xrightarrow{\mathrm{ARK}_{v+1}\circ\mathrm{INV}} \frac{uv+u+1}{u}+v+1=\frac{1}{u} \\
&\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} u+u=0 \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} 0+u=u.
\end{aligned}
$$

Note that in all the above computations, we have only evaluated $\mathrm{INV}(\cdot)$ at the values $u, uv+1$, and $uv+u+1$, which are non-zero.

Case 2(b). $x = u + \frac{1}{u}$**:** The permutation $\gamma(u,v)$ maps $u + \frac{1}{u}$ to itself as we show below:

$$\begin{aligned} u+\frac{1}{u} &\xrightarrow{\mathrm{ARK}_u} u+\frac{1}{u}+u=\frac{1}{u} \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} u+u=0 \\ &\xrightarrow{\mathrm{ARK}_v\circ\mathrm{INV}} 0+v=v \xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} \frac{1}{v}+1=\frac{v+1}{v} \\ &\xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} \frac{v}{v+1}+1=\frac{1}{v+1} \xrightarrow{\mathrm{ARK}_{v+1}\circ\mathrm{INV}} 0 \\ &\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} u \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{1}{u}+u. \end{aligned}$$

Note that in all the above computations, we have only evaluated $\mathrm{INV}(\cdot)$ at the values u, v, and $v+1$, which are non-zero.

Case 2(c). $x = u + \frac{v}{uv+1}$**:** The permutation $\gamma(u,v)$ maps $u + \frac{v}{uv+1}$ to $u + \frac{v+1}{uv+u+1}$ as we show below:

$$\begin{aligned} u+\frac{v}{uv+1} &\xrightarrow{\mathrm{ARK}_u} u+\frac{v}{uv+1}+u=\frac{v}{uv+1} \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{uv+1}{v}+u=\frac{1}{v} \\ &\xrightarrow{\mathrm{ARK}_v\circ\mathrm{INV}} v+v=0 \xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} 0+1=1 \xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} 1+1=0 \\ &\xrightarrow{\mathrm{ARK}_{v+1}\circ\mathrm{INV}} 0+v+1=v+1 \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{1}{v+1}+u=\frac{uv+u+1}{v+1} \\ &\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{v+1}{uv+u+1}+u. \end{aligned}$$

Note that in all the above computations, we have only evaluated $\mathrm{INV}(\cdot)$ at the values $uv+1, uv+u+1, v$, and $v+1$, which are non-zero.

Case 2(d). $x = u + \frac{v+1}{uv+u+1}$**:** The permutation $\gamma(u,v)$ maps $u + \frac{v+1}{uv+u+1}$ to $u + \frac{v}{uv+1}$ as we show below:

$$\begin{aligned} u+\frac{v+1}{uv+u+1} &\xrightarrow{\mathrm{ARK}_u} u+\frac{v+1}{uv+u+1}+u=\frac{v+1}{uv+u+1} \\ &\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{uv+u+1}{v+1}+u=\frac{1}{v+1} \xrightarrow{\mathrm{ARK}_v\circ\mathrm{INV}} v+1+v=1 \\ &\xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} 1+1=0 \xrightarrow{\mathrm{ARK}_1\circ\mathrm{INV}} 0+1=1 \xrightarrow{\mathrm{ARK}_{v+1}\circ\mathrm{INV}} 1+v+1=v \\ &\xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{1}{v}+u=\frac{uv+1}{v} \xrightarrow{\mathrm{ARK}_u\circ\mathrm{INV}} \frac{v}{uv+1}+u. \end{aligned}$$

Note that in all the above computations, we have only evaluated $\mathrm{INV}(\cdot)$ at the values $uv+1, uv+u+1, v$, and $v+1$, which are non-zero. □

Note that Lemma 3.8 is already enough to give us a bound on the number of operations required to simulate a random S-box. We use Lemma 3.9 to construct a flow with lower congestion, and thus a better comparison constant.

Lemma 3.9. *The following two sequences of* ARK *and* INV *operations implement the same permutations*

$$[[1,1]] = [[\mathrm{INV}(w)+1, w, \mathrm{INV}(w), w, \mathrm{INV}(w), w+1]]$$

Proof. Denote by $\sigma_{LHS}, \sigma_{RHS}$ as the permutations of the LHS and RHS respectively. Then σ_{LHS} maps $x \to \mathrm{INV}(x+1)+1$. We will show that this is the case of σ_{RHS}. For simplicity, we will first compute the image of x under σ_{RHS}, assuming that no input to $\mathrm{INV}(\cdot)$ is equal to zero. Thus, we will use the fact that $y \cdot \mathrm{INV}(y) = 1$ for all non-zero y. Also for brevity, we will use $\mathrm{INV}(y)$ and $\frac{1}{y}$ interchangeably.

$$\begin{aligned}
& x \\
\xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)+1}} \; & x + \frac{1}{w} + 1 = \frac{xw+w+1}{w} \\
\xrightarrow{\mathrm{ARK}_w \circ \mathrm{INV}} \; & \frac{w}{xw+w+1} + w = \frac{xw^2+w^2}{xw+w+1} \\
\xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)} \circ \mathrm{INV}} \; & \frac{xw+w+1}{xw^2+w^2} + \frac{1}{w} = \frac{1}{xw^2+w^2} \\
\xrightarrow{\mathrm{ARK}_w \circ \mathrm{INV}} \; & xw^2+w^2+w \\
\xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)} \circ \mathrm{INV}} \; & \frac{1}{xw^2+w^2+w} + \frac{1}{w} = \frac{xw+w}{xw^2+w^2+w} = \frac{x+1}{xw+w+1} \\
\xrightarrow{\mathrm{ARK}_{w+1} \circ \mathrm{INV}} \; & \frac{xw+w+1}{x+1} + w + 1 = \frac{xw+w+1+(xw+x)+(w+1)}{x+1} \\
& = \frac{x}{x+1} = \mathrm{INV}(x+1)+1
\end{aligned}$$

To complete our proof, we now consider what happens if some input to $\mathrm{INV}(\cdot)$ equals 0. Thus, we will consider the following cases separately:

1. $w = 0$,
2. $xw + w + 1 = 0 \implies x = \frac{w+1}{w}$
3. $xw^2 + w^2 = 0 \implies x = 1$
4. $x + 1 = 0 \implies x = 1$.

Case 1. $w = 0$: σ_{RHS} becomes the permutation denoted by $[[1, 0, 0, 0, 0, 1]]$. This permutation maps

$$\begin{aligned}
x \quad & \xrightarrow{\mathrm{ARK}_1} x+1 \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(x+1) \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} x+1 \\
& \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} \mathrm{INV}(x+1) \xrightarrow{\mathrm{ARK}_0 \circ \mathrm{INV}} x+1 \xrightarrow{\mathrm{ARK}_1 \circ \mathrm{INV}} \mathrm{INV}(x+1)+1.
\end{aligned}$$

Note that in the above expression, we only used the fact that $\mathrm{INV}(\mathrm{INV}(y)) = y$, which holds for all y. Hence the above mapping holds for all x.

Case 2. $x = \frac{w+1}{w}$: For simplicity we will assume that $w \neq 0$, as this case was already covered above. This allows us to replace $\mathrm{INV}(w)+1$ with $\frac{1}{w} + 1 = \frac{w+1}{w}$. The permutation σ_{RHS} maps $\frac{w+1}{w}$ to

$$\begin{aligned}
\frac{w+1}{w} & \xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)+1}} \frac{w+1}{w} + \frac{w+1}{w} = 0 \xrightarrow{\mathrm{ARK}_w \circ \mathrm{INV}} 0 + w = w \\
& \xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)} \circ \mathrm{INV}} \mathrm{INV}(w) + \mathrm{INV}(w) = 0 \xrightarrow{\mathrm{ARK}_w \circ \mathrm{INV}} 0 + w = w \\
& \xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)} \circ \mathrm{INV}} \mathrm{INV}(w) + \mathrm{INV}(w) = 0 \xrightarrow{\mathrm{ARK}_{w+1} \circ \mathrm{INV}} 0 + w + 1 = w + 1.
\end{aligned}$$

Note that

$$\mathrm{INV}(x+1)+1 = \mathrm{INV}\left(\frac{w+1}{w} + 1\right) + 1 = \mathrm{INV}\left(\frac{1}{w}\right) + 1 = w + 1.$$

Thus σ_{RHS} maps this value of x to the same image as the $[[1, 1]]$ permutation.

Case 3. $x = 1$**:** For simplicity we will assume that $w \neq 0$, as this case was already covered above. This allows us to replace $\mathrm{INV}(w) + 1$ with $\frac{1}{w} + 1 = \frac{w+1}{w}$. The permutation σ_{RHS} maps 1 to

$$1 \xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)+1}} 1 + \mathrm{INV}(w) + 1 = \mathrm{INV}(w) \xrightarrow{\mathrm{ARK}_w \circ \mathrm{INV}} w + w = 0$$
$$\xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)} \circ \mathrm{INV}} 0 + \mathrm{INV}(w) = \mathrm{INV}(w) \xrightarrow{\mathrm{ARK}_w \circ \mathrm{INV}} w + w = 0$$
$$\xrightarrow{\mathrm{ARK}_{\mathrm{INV}(w)} \circ \mathrm{INV}} 0 + \mathrm{INV}(w) = \mathrm{INV}(w) \xrightarrow{\mathrm{ARK}_{w+1} \circ \mathrm{INV}} w + w + 1 = 1.$$

Again, $\mathrm{INV}(1+1) + 1 = 1$, thus σ_{RHS} maps $x = 1$ to the same image as the $[[1,1]]$ permutation.

□

Proof of Lemma 3.7. From Lemmas 3.8 and 3.9 we know that the sequence

$$\psi'_{\text{even}}(u,v,w) = [[u,\ u,\ v,\ \mathrm{INV}(w)+1,\ w,\ \mathrm{INV}(w),$$
$$w,\ \mathrm{INV}(w),\ w+1,\ v+1,\ u,\ u]].$$

generates the transposition $(u + \frac{v}{uv+1}, u + \frac{v+1}{uv+u+1})$. This sequence is obtained by substituting the statement of Lemma 3.9 into the middle part of $\gamma(u,v)$ from Lemma 3.8.

Now we modify ψ'_{even} to obtain transpositions with the fixed element 0, i.e. $(0, i)$ for non-zero $i \in \mathbb{F}_n$. We do this by "relabelling" the left endpoint of the resulting transposition to be a 0. It suffices to conjugate ψ'_{even} with a permutation that maps $u + \frac{v}{uv+1}$ to 0, e.g. $\mathrm{ARK}_{u+\frac{v}{uv+1}}$.

Composing two ARK operations gives another ARK operation with a key equal to the sum of the two original round keys. Thus

$$\psi_{\text{even}}(u,v,w) = \mathrm{ARK}_{u+\frac{v}{uv+1}} \circ \psi'_{\text{even}}(u,v,w) \circ \mathrm{ARK}_{u+\frac{v}{uv+1}}\ .$$

□

Using Lemma 3.7, we can construct paths that connect adjacent vertices of $P_n^{\text{2cyc},0}$ using edges of P_n^{INV}.

Corollary 3.10 (Corollary of Lemma 3.7)**.** *Let $n = 2^b$. For any $u, v, w \in \mathbb{F}_n$ such that $uv \neq 1$, $u(v+1) \neq 1$, and $v \notin \{0,1\}$, and any $r_1, \ldots, r_{10} \in \mathbb{F}_n$ we can generate the transposition $\left(0, \frac{v}{uv+1} + \frac{v+1}{uv+u+1}\right)$ using the following path $\phi_{\text{even}}(u,v,w,r_1,\ldots,r_{10})$ of 11 edges of P_n^{INV}:*

$$\phi_{\text{even}}(u,v,w,r_1,\ldots,r_{10}) = \left(\left[\left[\frac{v}{uv+1},\ r_1\right]\right],\ \ [[u+r_1,\ r_2]]\,,\right.$$
$$[[v+r_2,\ r_3]]\,,\ \ [[\mathrm{INV}(w)+1+r_3,\ r_4]]\,,\ \ [[w+r_4,\ r_5]]\,,$$
$$[[\mathrm{INV}(w)+r_5,\ r_6]]\,,\ \ [[w+r_6,\ r_7]]\,,\ \ [[\mathrm{INV}(w)+r_7,\ r_8]]\,,$$
$$\left. [[w+1+r_8,\ r_9]]\,,\ \ [[v+1+r_9,\ r_{10}]]\,,\ \ \left[\left[u+r_{10},\ \frac{v}{uv+1}\right]\right] \right).$$

Proof of Lemma 3.6. Towards applying the Comparison Theorem (Theorem 2.24), we will only assign a non-zero amount of flow to the ϕ_{even} paths defined in Corollary 3.10. Formally, let $y = (0, \frac{v}{uv+1} + \frac{v+1}{uv+u+1}) \circ x$, then:

$$\mathcal{P}_{xy} = \big\{\phi_{\text{even}}(u,v,w,r_1,r_2,\ldots,r_{10})$$
$$: u,v,w,r_1,r_2,\ldots,r_{10} \in \mathbb{F}_n, uv \neq 1,\ u(v+1) \neq 1,\ v \notin \{0,1\}\big\}.$$

We further denote by $\mathcal{P}_x$ be the set of ϕ_{even} paths that start from x. It holds that $|\mathcal{P}_x| = n^{11}(n-2)^2 = \Theta(n^{13})$.

Lemma A.1 implies that if we push the same amount of flow through all paths in $\mathcal{P}_x$, we will get an almost uniform flow routed through all edges $(x, y) \in E_n^{\text{2cyc},0}$. Formally, for $n \geq 8$ it holds that for any $(x, y), (x, y') \in E_n^{\text{2cyc},0}$: $\frac{1}{3} \leq \frac{|\mathcal{P}_{xy}|}{|\mathcal{P}_{xy'}|} \leq 3$. In other words, the maximum total flow along any $(x, y) \in P_n^{\text{2cyc},0}0$ is at most 3 times the total flow along any other edge.

Since the stationary distribution for $P_n^{\text{2cyc},0}$ is uniform over the alternating group and $P_n^{\text{2cyc},0}(x, y) = \frac{1}{n-1}$ for all $(x, y) \in E_n^{\text{2cyc},0}$, it should hold that that total amount of flow through the simple paths that connect the vertices x and y is

$$\sum_{\gamma \in \mathcal{P}_{xy}} F(\gamma) = \frac{2}{n!(n-1)} =: F.$$

Since the number of edges incident to x in $P_n^{\text{2cyc},0}0$ is $(n-1)$, pushing one unit of flow through each of the $n^{11}(n-2)^2$ paths will result in each edge (x, y) getting $c_{xy} \cdot \frac{n^{11}(n-2)^2}{n-1} = \Theta(n^{12})$ (where $1/3 \leq c_{xy} \leq 3$) units of flow. Since our goal is to push F units through each edge, the flow through each path will be

$$F(\gamma) = \frac{F}{c_{xy} \cdot \frac{n^{11}(n-2)^2}{n-1}} = d_{xy} \cdot \frac{F}{n^{12}} = \Theta\left(\frac{F}{n^{12}}\right),$$

where $d_{xy} = \frac{n(n-1)}{(n-2)^2 c_{xy}} \in [1/3, 6]$ for $n \geq 8$.

We are now ready to compute the comparison constant for this flow.

$$\begin{aligned}
A(F) &= \max_{(a,b) \in E_n^{\text{INV}}} \left\{ \frac{1}{\pi^{\text{INV}}(a) P_n^{\text{INV}}(a, b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma| \cdot F(\gamma) \right\} \\
&\leq \max_{(a,b) \in E_n^{\text{INV}}} \left\{ \frac{1}{\frac{2}{n!} \cdot \frac{1}{n^2}} |\mathcal{P}(a, b)| \cdot 11 \cdot \frac{2}{n!(n-1)} \cdot \frac{d_{xy}}{n^{12}} \right\} \\
&\leq \max_{(a,b) \in E_n^{\text{INV}}} \left\{ \frac{n! n^2}{2} \cdot 11 n^{11} \cdot 11 \cdot \frac{2}{n!(n-1)} \cdot \frac{d_{xy}}{n^{12}} \right\} \\
&\leq \frac{121 n^{13} \cdot d_{xy}}{n^{12}(n-1)} \leq 830.
\end{aligned}$$

The last inequality holds for $n \geq 8$. We used the fact that $|\mathcal{P}(a, b)| \leq 11 n^{11}$. This is because P_n^{INV} has no parallel edges (Lemma 2.15), and thus the edge (a, b) fully specifies a unique permutation of the form $[[r_1, r_2]]$.

Now consider a path $\gamma = \phi_{\text{even}}(u', v', w' r'_1, r'_2, \ldots, r'_{10})$ that uses edge (a, b). This edge can be one of 11 edges of γ; let's say that it is the i^{th} edge, for $i \in \{1, 2, \ldots, 11\}$. Every value of i implies two equations that the set of variables $\{u', v', w', r'_1, r'_2, \ldots, r'_{10}\}$ must satisfy. This restricts 2 of the 13 degrees of freedom; thus, we can have at most n^{11} such paths, since all paths are linearly dependent on the variables $(u', v', w', r'_1, r'_2, \ldots, r'_{10})$, or their inverses.

The modified log-Sobolev bound follows since we use paths of length at most 11, i.e. $\gamma_{\max} = 11$. □

3.2 Fields of odd characteristic

The main result of this section is a bound on the log-Sobolev constant of P_n^{INV} when $n = p^b$ is a power of an odd prime p.

Lemma 3.11. *Let $n = p^b$, where p is an odd prime. The log-Sobolev constant, modified log-Sobolev constant, and spectral gap of P_n^{INV} satisfy*

- *Log-Sobolev*

$$\alpha_n^{\mathrm{INV}} \geq \frac{1}{24 \cdot 81 \cdot 60 \cdot n \log n}$$

- *Modified log-Sobolev*

$$\beta_n^{\mathrm{INV}} \geq \frac{1}{96 \cdot 729 \cdot 2n}$$

- *Spectral gap*

$$\lambda_n^{\mathrm{INV}} \geq \frac{1}{24 \cdot 81 \cdot 2n}$$

Proof. For all three functional inequalities, we employ the same comparison method in two steps. We introduce an intermediate chain $P_n^{\mathrm{3cyc,ap}}$, which is a slight variant of the 3-cycle chain, and compare the log-Sobolev constant of P_n^{INV} with it. Then we bound $\alpha_n^{\mathrm{3cyc,ap}}$ by comparing it to the 3-cycle chain P_n^{3cyc}. The proof follows directly from the comparison method in Lemmas 3.13 and 3.14 and known bounds on the functional inequalities of the 3-cycle chain.

Log-Sobolev.

$$\alpha_n^{\mathrm{INV}} \underset{\text{Lemma 3.14}}{\geq} \frac{\alpha_n^{\mathrm{3cyc,ap}}}{24} \underset{\text{Lemma 3.13}}{\geq} \frac{\alpha_n^{\mathrm{3cyc}}}{24 \cdot 81} \underset{\text{Lemma 2.22}}{\geq} \frac{1}{24 \cdot 81 \cdot 60 \cdot n \log n}.$$

Modified log-Sobolev.

$$\beta_n^{\mathrm{INV}} \underset{\text{Lemma 3.14}}{\geq} \frac{\beta_n^{\mathrm{3cyc,ap}}}{96} \underset{\text{Lemma 3.13}}{\geq} \frac{\beta_n^{\mathrm{3cyc}}}{96 \cdot 729} \underset{\text{Theorem 2.21}}{\geq} \frac{1}{96 \cdot 729 \cdot n}.$$

Spectral gap.

$$\lambda_n^{\mathrm{INV}} \underset{\text{Lemma 3.14}}{\geq} \frac{\lambda_n^{\mathrm{3cyc,ap}}}{24} \underset{\text{Lemma 3.13}}{\geq} \frac{\lambda_n^{\mathrm{3cyc}}}{24 \cdot 81} \underset{\text{Corollary 2.23}}{\geq} \frac{1}{24 \cdot 81 \cdot 2n}.$$

□

Definition 3.12 (Arithmetic progression 3-cycle Markov chain)**.** *The chain $P_n^{\mathrm{3cyc,ap}}$ on the alternating group A_n has the following transition matrix. Given the current permutation σ_t, one step in this Markov chain corresponds to drawing a uniformly random 3-cycle (i,j,k) from A_n, conditioned on the fact that i,k,j form an arithmetic progression, that is, $k-i=j-k$. Then set*

$$\sigma_{t+1} = (i,j,k) \circ \sigma_t.$$

Our paths construction shows that the underlying graph is connected, and thus the stationary distribution of this Markov chain is uniform over A_n, i.e. $\pi_n^{\mathrm{3cyc,ap}}(x) = \frac{1}{|A_n|} = \frac{2}{n!}$. Moreover, the underlying graph is $n(n-1)$-regular. This is because there are n options for the first term of the arithmetic progression u, and $n-1$ options for the difference $-\frac{1}{v}$. Thus if $(x,y) \in E_n^{\mathrm{3cyc,ap}}$, then $P_n^{\mathrm{3cyc,ap}}(x,y) = \frac{1}{n(n-1)}$.

Lemma 3.13. *The functional inequalities of $P_n^{\mathrm{3cyc,ap}}$ satisfy*

- *Log-Sobolev*

$$\alpha_n^{\mathrm{3cyc,ap}} \geq \frac{\alpha_n^{\mathrm{3cyc}}}{81}.$$

- *Modified log-Sobolev*
$$\beta_n^{\text{3cyc,ap}} \geq \frac{\beta_n^{\text{3cyc}}}{729}.$$
- *Spectral gap*
$$\lambda_n^{\text{3cyc,ap}} \geq \frac{\lambda_n^{\text{3cyc}}}{81}.$$

Proof. To apply the Comparison Theorems 2.24, and 2.27, we will define a set of paths $\mathcal{P}_{xy}$ for each edge $(x, y) \in E^{\text{3cyc}}$ and assign flow $F(\cdot)$ to each path.

Our construction of paths in this case is quite simple. We will assign exactly one path to each such edge (x, y). In particular, let $y = (i, k, j) \circ x$ for some triple of pairwise distinct elements i, j, k. Then $\mathcal{P}_{xy} = \{\gamma_{xy}\}$, and this path is chosen according to the following two cases:

1. The elements i, j, k form an arithmetic progression. This means that either $k - i = j - k$, or $i - j = k - i$, or $j - k = i - j$. Then (i, k, j) is also in $E_n^{\text{3cyc,ap}}$ and we set
$$\gamma_{xy} = (x, (i, k, j) \circ x = y).$$

2. The elements i, j, k do not form an arithmetic progression. Let $u = \frac{i+j}{2}, v = \frac{j+k}{2}, w = \frac{k+i}{2}$. Then $E_n^{\text{3cyc,ap}}$ contains the distinct 3-cycles
$$C_1 = (i, j, u), \quad C_2 = (j, k, v), \quad C_3 = (k, i, w).$$
We assign to (x, y) the length-9 path
$$\begin{aligned}\gamma_{xy} = (x, \underbrace{C_1 \circ x}_{a_1}, \underbrace{C_3 \circ a_1}_{a_2}, \underbrace{C_3 \circ a_2}_{a_3}, \underbrace{C_2 \circ a_3}_{a_4}, \\ \underbrace{C_2 \circ a_4}_{a_5}, \underbrace{C_3 \circ a_5}_{a_6}, \underbrace{C_1 \circ a_6}_{a_7}, \underbrace{C_2 \circ a_7}_{a_8}, C_2 \circ a_8 = y).\end{aligned}$$
The proof that this path connects x to y is deferred to Lemma A.2.

Since each edge $(x, y) \in E_n^{\text{3cyc}}$ has exactly one path to it, all of the flow must go through this path:
$$F(\gamma_{xy}) = P_n^{\text{3cyc}}(x, y) \cdot \pi_n^{\text{3cyc}}(x, y) = \frac{1}{2\binom{n}{3}} \cdot \frac{2}{n!} = \frac{6}{n(n-1)(n-2) \cdot n!}.$$

The comparison constant we get is thus
$$\begin{aligned} A(F) &= \max_{(a,b) \in E^{\text{3cyc,ap}}} \left\{ \frac{1}{\pi^{\text{3cyc,ap}}(a) \cdot P^{\text{3cyc,ap}}(a, b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma| \cdot F(\gamma) \right\} \\ &\leq \max_{(a,b) \in E^{\text{3cyc,ap}}} \left\{ \frac{n!}{2} \cdot \binom{n}{2} \cdot 9 \cdot |\mathcal{P}(a, b)| \cdot \frac{6}{n(n-1)(n-2) \cdot n!} \right\} \\ &\leq \frac{n(n-1) \cdot n!}{4} \cdot 54(n-2) \cdot \frac{6}{n(n-1)(n-2) \cdot n!} = 81. \end{aligned}$$

We used the fact that the number of paths γ in $\mathcal{P}(a, b)$ is at most $6(n-2)$. This is because $\mathcal{P}(a, b)$ contains paths of length 1 and 9. The set $\mathcal{P}(a, b)$ contains at most 1 path of length 1 for any $(a, b) \in E^{\text{3cyc,ap}}$.

Moreover, length-9 paths γ in $\mathcal{P}(a,b)$ must be using the edge (a,b) as their cycle C_1, C_2, or C_3. Without loss of generality, assume (a,b) is used as cycle C_1 in γ_{xy}, and we will multiply the number of paths by 3 to capture the other two cases. Then the edge (a,b) specifies the set of elements $\{i,j,u\}$. There are two ways to choose i,j from this set, and for each such setting of i,j there are $n-3$ remaining elements that could be k. Thus the total number of length-9 paths is at most $3 \cdot 2 \cdot (n-3) = 6(n-3)$. The total number of paths is at most $1 + 6(n-3) \leq 6(n-2)$.

The modified log-Sobolev bound follows since we use paths of length at most 9, i.e. $\gamma_{\max} = 9$. □

Lemma 3.14. *Let $n = p^b$, where p is an odd prime. The functional inequalities of P_n^{INV} satisfy*

- *Log-Sobolev*
$$\alpha_n^{\mathrm{INV}} \geq \frac{\alpha_n^{\mathrm{3cyc,ap}}}{24}.$$
- *Modified log-Sobolev*
$$\beta_n^{\mathrm{INV}} \geq \frac{\beta_n^{\mathrm{3cyc,ap}}}{96}.$$
- *Spectral gap*
$$\lambda_n^{\mathrm{INV}} \geq \frac{\lambda_n^{\mathrm{3cyc,ap}}}{24}.$$

The proof uses the following way to generate the arithmetic progression 3-cycle $\left(-u, \quad -u-\frac{2}{v}, \quad -u-\frac{1}{v}\right)$ by combining INV and ARK operations.

Lemma 3.15. *Let $n = p^b$, where p is an odd prime. For any $u, v \in \mathbb{F}_n$, $v \neq 0$, the following sequence $\psi_{\mathrm{odd}}(u,v)$ of* ARK *and* INV *operations*

$$\psi_{\mathrm{odd}}(u,v) = \left[\left[u, \quad v, \quad -\frac{2}{v}, \quad v, \quad -u-\frac{2}{v}\right]\right]$$

maps x to

$$\begin{cases} -u-\frac{2}{v}, & x = -u \\ -u, & x = -u-\frac{1}{v} \\ -u-\frac{1}{v}, & x = -u-\frac{2}{v} \\ x, & \textit{otherwise} \end{cases}.$$

Proof. We consider first the application of $\psi_{\mathrm{even}}(u,v)$ on some x that is not equal to any of $\{-u, -u-\frac{1}{v}, -u-\frac{2}{v}\}$.

$$\begin{aligned} x &\xrightarrow{\mathrm{ARK}_u} x+u \\ &\xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} \frac{1}{x+u} + v = \frac{xv+uv+1}{x+u} \\ &\xrightarrow{\mathrm{ARK}_{-2/v} \circ \mathrm{INV}} \frac{x+u}{xv+uv+1} - \frac{2}{v} = -\frac{xv+uv+2}{v(xv+uv+1)} \\ &\xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} -\frac{v(xv+uv+1)}{xv+uv+2} + v = \frac{v}{xv+uv+2} \\ &\xrightarrow{\mathrm{ARK}_{-u-2/v} \circ \mathrm{INV}} -\frac{xv+uv+2}{v} - u - \frac{2}{v} = x. \end{aligned}$$

Now consider what happens when $x = -u$:

$$-u \xrightarrow{\mathrm{ARK}_u} 0 \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} v \xrightarrow{\mathrm{ARK}_{-2/v} \circ \mathrm{INV}} -\frac{1}{v} \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} 0 \xrightarrow{\mathrm{ARK}_{-u-2/v} \circ \mathrm{INV}} -u - \frac{2}{v}.$$

Now consider what happens when $x = -u - \frac{1}{v}$:

$$-u - \frac{1}{v} \xrightarrow{\mathrm{ARK}_u} -\frac{1}{v} \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} 0 \xrightarrow{\mathrm{ARK}_{-2/v} \circ \mathrm{INV}} -\frac{2}{v} \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} \frac{v}{2} \xrightarrow{\mathrm{ARK}_{-u-2/v} \circ \mathrm{INV}} -u.$$

Now consider what happens when $x = -u - \frac{2}{v}$:

$$-u - \frac{2}{v} \xrightarrow{\mathrm{ARK}_u} -\frac{2}{v} \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} \frac{v}{2} \xrightarrow{\mathrm{ARK}_{-2/v} \circ \mathrm{INV}} 0 \xrightarrow{\mathrm{ARK}_v \circ \mathrm{INV}} v \xrightarrow{\mathrm{ARK}_{-u-2/v} \circ \mathrm{INV}} -u - \frac{1}{v}.$$

□

Using the above Lemma, we can construct paths that connect adjacent vertices of $P_n^{\mathrm{3cyc,ap}}$ using edges of P_n^{INV}.

Corollary 3.16 (Corollary of Lemma 3.15)**.** *Let $n = p^b$, where p is an odd prime. For any $u, v \in \mathbb{F}_n$ such that $v \neq 0$, and any $r_1, r_2, r_3 \in \mathbb{F}_n$ we can generate the 3-cycle $\left(-u, -u - \frac{2}{v}, -u - \frac{1}{v}\right)$ using the following length-4 path $\phi_{\mathrm{odd}}(u, v, r_1, r_2, r_3)$ in P_n^{INV}:*

$$\phi_{\mathrm{odd}}(u, v, r_1, r_2, r_3) = \left([[u,\ r_1]], [[v - r_1,\ r_2]],\ \left[\left[-\frac{2}{v} - r_2,\ r_3\right]\right],\ \left[\left[v - r_3,\ -u - \frac{2}{v}\right]\right] \right).$$

Proof of Lemma 3.14. Towards applying the Comparison Theorems 2.24 and 2.27, we will assign to the edge $(x, y) \in E^{\mathrm{3cyc,ap}}$ the set of ϕ_{odd} paths defined in Corollary 3.16. Formally, let $y = (-u, -u - \frac{2}{v}, -u, -\frac{1}{v}) \circ x$, then:

$$\mathcal{P}_{xy} = \{\phi_{\mathrm{odd}}(u, v, r_1, r_2, r_3) : r_1, r_2, r_3 \in \mathbb{F}_n\}.$$

It holds that $|\mathcal{P}_{xy}| = n^3$. We will assign the same amount of flow through all paths in $\mathcal{P}(x, y)$. This means that

$$F(\gamma_{xy}) = \frac{P_n^{\mathrm{3cyc,ap}}(x, y) \cdot \pi_n^{\mathrm{3cyc,ap}}(x, y)}{n^3} = \frac{1}{n^3} \cdot \frac{1}{n(n-1)} \cdot \frac{2}{n!} = \frac{2}{n^4(n-1) \cdot n!}.$$

The comparison constant we get is

$$\begin{aligned} A(F) &= \max_{(a,b) \in E^{\mathrm{INV}}} \left\{ \frac{1}{\pi^{\mathrm{INV}}(a) \cdot P^{\mathrm{INV}}(a, b)} \sum_{\gamma \in \mathcal{P}(a,b)} |\gamma| \cdot F(\gamma) \right\} \\ &= \max_{(a,b) \in E^{\mathrm{INV}}} \left\{ \frac{n^2 \cdot n!}{2} \cdot \frac{4 \cdot 2}{n^4(n-1) \cdot n!} \cdot |\mathcal{P}(a, b)| \right\} \\ &\leq \frac{n^2 \cdot n!}{2} \cdot \frac{4 \cdot 2}{n^4(n-1) \cdot n!} \cdot 4n^3 \\ &= \frac{16n}{n-1} \leq 24. \end{aligned}$$

We used the fact that $|\mathcal{P}(a,b)| \le 4n^3$. Lemma 2.15 implies that P_n^{INV} has no parallel edges, and thus the edge (a,b) fully specifies a unique permutation of the form $[[r_1, r_2]]$.

Now consider a path $\gamma = \phi_{\mathrm{odd}}(u', v', r_1', r_2', r_3')$ that uses edge (a,b). This edge can be one of 4 edges of γ; let's say that it is the i^{th} edge, for $i \in \{1,2,3,4\}$. Every value of i implies two equations that the set of variables $\{u', v', r_1', r_2', r_3'\}$ must satisfy. This restricts 2 of the 5 degrees of freedom; thus, we can have at most n^3 such paths, since all paths are linearly dependent on the variables $(u', v', r_1', r_2', r_3')$, or their inverses. The last inequality holds because $n \ge 3$.

The modified log-Sobolev bound follows since we use paths of length at most 4, i.e. $\gamma_{\max} = 4$.

□

4 Lower Bound

In this section, we demonstrate a lower bound on the number of rounds required for the INV KAC to be close to a 4-wise independent permutation, first shown by [LTV21]. This directly implies the same lower bound for the block cipher to converge to A_n.

The crucial observation is the following lemma from [LTV21], whose statement and proof we include below almost verbatim for completeness.

Lemma 4.1 (Lemma B.1 of [LTV21])**.** *For every r, with probability $1 - \frac{r}{n}$ over a random choice of $K_0, \dots, K_r \sim \mathbb{F}_n$, there are $L_1, L_2, L_3 \in \mathbb{F}_n$ such that*

$$F^{(r)}_{\mathrm{INV}, K_0, \dots, K_r}(x) = (x + L_1) \cdot \mathrm{INV}(L_2 x + L_3).$$

Proof. The proof is by induction. For $r = 0$, $L_1 = K_0$, $L_2 = 0$ and $L_3 = 1$. Let us now assume that the statement is true for i. Then:

$$\begin{aligned} F^{(i+1)}_{\mathrm{INV}, K_0, \dots, K_r}(x) &= \mathrm{INV}\left(F^{(i)}_{\mathrm{INV}, K_0, \dots, K_r}(x) + K_{i+1}\right) \\ &= \mathrm{INV}\left(\frac{x + L_1}{L_2 x + L_3} + K_{i+1}\right) = \frac{L_2 x + L_3}{(K_{i+1} L_2 + 1)x + (K_i L_3 + L_1)} \end{aligned}$$

which is of the same form as well. The only way this fails is if one of the numerators in the expression turns out to be 0. The probability of this happening for any one of the r rounds is at most $\frac{r}{n}$. □

Theorem 4.2. *An r-round* INV *KAC requires at least $r \ge \frac{(1-\varepsilon)n}{4} - \frac{1}{2}$ rounds to reach ε-close to a 4-wise independent permutation.*

Proof. We will construct the following algorithm $\mathcal{A}$ that distinguishes between an r-round INV KAC and a truly random 4-wise independent permutation. The algorithm first chooses 4 inputs x_1, x_2, x_3, x_4 and computes their images y_1, y_2, y_3, y_4. Then if there are L_1, L_2, L_3 such that $y_i = (x_i + L_1) \cdot \mathrm{INV}(L_2 x_i + L_3)$, the distinguisher will guess "INV KAC". Otherwise, it will guess a random permutation.

From the lemma above, the probability that the distinguisher correctly detects the INV KAC is at least $1 - \frac{4r}{n}$, since we can union bound over all 4 inputs.

On the other hand, the distinguisher will be fooled by a 4-wise independent permutation with probability at most $\frac{1}{n-3}$. This is because the first three inputs and outputs will give linear equations that determine the constants L_1, L_2, L_3. Thus the last input and output must satisfy $y_4 = (x_4 + L_1) \cdot \mathrm{INV}(L_2 x_4 + L_3)$. Since y_4 is uniformly random from $n - 3$ values this can only happen with probability at most $\frac{1}{n-3}$.

The total variation distance implies an upper bound in the distinguishing advantage of any adversary. Thus for the r-round INV KAC to be ε-close to a uniformly random 4-wise independent permutation, the advantage of $\mathcal{A}$ must be at most ε. Thus

$$\begin{aligned}
1 - \frac{4r}{n} - \varepsilon &\leq \frac{1}{n-3} \\
\implies \frac{4r}{n} &\geq 1 - \varepsilon - \frac{1}{n-3} \\
\implies r &\geq \frac{(1-\varepsilon)n}{4} - \frac{n}{4(n-3)} \\
&\geq \frac{(1-\varepsilon)n}{4} - \frac{1}{2}.
\end{aligned}$$

Where the last inequality holds for $n \geq 8$. □

Acknowledgments

The authors would like to thank Alistair Sinclair and Lucas Gretta for very helpful and insightful discussions regarding the functional analysis of Markov chains. We would also like to thank Amit Rajaraman and Sidhanth Mohanty for pointing out how to obtain a comparison theorem for the modified log-Sobolev constant.

References

[AcMT09] Esen Aksoy, Ayça Çeşmelioglu, Wilfried Meidl, and Alev Topuzoglu. On the carlitz rank of permutation polynomials. *Finite Fields and Their Applications*, 15(4):428–440, 2009. 1

[AGR+16] Martin Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. MiMC: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. Cryptology ePrint Archive, Paper 2016/492, 2016. https://eprint.iacr.org/2016/492. 1

[BKR11] Andrey Bogdanov, Dmitry Khovratovich, and Christian Rechberger. Biclique cryptanalysis of the full AES. pages 344–371, 2011. 1

[BS91] Eli Biham and Adi Shamir. Differential cryptanalysis of des-like cryptosystems. *J. Cryptol.*, 4(1):3–72, 1991. 1

[BT03] Sergey Bobkov and Prasad Tetali. Modified log-sobolev inequalities, mixing and hypercontractivity. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '03, page 287–296, New York, NY, USA, 2003. Association for Computing Machinery. 2.9, 2.12

[BV05] Thomas Baignères and Serge Vaudenay. Proving the security of AES substitution-permutation network. In Bart Preneel and Stafford E. Tavares, editors, *Selected Areas in Cryptography, 12th International Workshop, SAC 2005, Kingston, ON, Canada, August 11-12, 2005, Revised Selected Papers*, volume 3897 of *Lecture Notes in Computer Science*, pages 65–81. Springer, 2005. 1

[Car53] L. Carlitz. Permutations in a finite field. *Proceedings of American Mathematical Society*, page 538, 1953. 1

[Car63] L Carlitz. A note on permutations in an arbitrary field. In *Proc. Amer. Math. Soc.*, volume 14, page 101, 1963. 1, 1

[CP02] Nicolas Courtois and Josef Pieprzyk. Cryptanalysis of block ciphers with overdefined systems of equations. pages 267–287, 2002. 1

[DR02] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Information Security and Cryptography. Springer, 2002. 1, 1

[DS81] Persi Diaconis and Mehrdad Shahshahani. Generating a random permutation with random transpositions. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 57(2):159–179, 1981. 1

[DSC93] Persi Diaconis and Laurent Saloff-Coste. Comparison Theorems for Reversible Markov Chains. *The Annals of Applied Probability*, 3(3):696 – 730, 1993. 1, 2.6, 2.24

[DSC96] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *The Annals of Applied Probability*, 6(3):695 – 750, 1996. 2.11, 2.13

[FOW18] Yuval Filmus, Ryan O'Donnell, and Xinyu Wu. A log-sobolev inequality for the multislice, with applications, 2018. 1

[Goe04] Sharad Goel. Modified logarithmic sobolev inequalities for some models of random walk. *Stochastic Processes and their Applications*, 114(1):51–79, 2004. 1, 2.18, 2.21

[IW18] Leyla Işik and Arne Winterhof. Carlitz rank and index of permutation polynomials. *Finite Fields and Their Applications*, 49:156–165, 2018. 1

[JK97] Thomas Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In *FSE*, volume 1267 of *Lecture Notes in Computer Science*, pages 28–40. Springer, 1997. 1, 1

[Knu94] Lars R. Knudsen. Truncated and higher order differentials. In Bart Preneel, editor, *Fast Software Encryption: Second International Workshop. Leuven, Belgium, 14-16 December 1994, Proceedings*, volume 1008 of *Lecture Notes in Computer Science*, pages 196–211. Springer, 1994. 1

[Knu98] Lars Knudsen. Deal - a 128-bit block cipher. In *NIST AES Proposal*, 1998. 1

[KW02] Lars R. Knudsen and David A. Wagner. Integral cryptanalysis. In *FSE*, volume 2365 of *Lecture Notes in Computer Science*, pages 112–127. Springer, 2002. 1

[Lai94] Xuejia Lai. *Higher Order Derivatives and Differential Cryptanalysis*, pages 227–233. Springer US, Boston, MA, 1994. 1

[LMR+24] Kuikui Liu, Sidhanth Mohanty, Prasad Raghavendra, Amit Rajaraman, and David X. Wu. Locally stationary distributions: A framework for analyzing slow-mixing markov chains. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 203–215, 2024. 1, 2.6, 2.26

[LPTV23] Tianren Liu, Angelos Pelecanos, Stefano Tessaro, and Vinod Vaikuntanathan. Layout graphs, random walks and the t-wise independence of SPN block ciphers. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023,* *Proceedings, Part III*, volume 14083 of *Lecture Notes in Computer Science*, pages 694–726. Springer, 2023. 1, 1, 1.1, 1

[LPTV24] Tianren Liu, Angelos Pelecanos, Stefano Tessaro, and Vinod Vaikuntanathan. Layout graphs, random walks and the t-wise independence of SPN block ciphers. Cryptology ePrint Archive, Paper 2024/083, 2024. 1

[LPW06] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006. 2.3

[LTV21] Tianren Liu, Stefano Tessaro, and Vinod Vaikuntanathan. The t-wise independence of substitution-permutation networks. *CRYPTO*, 2021. 1, 1, 1, 1, 4, 4.1

[LY98] Tzong-Yow Lee and Horng-Tzer Yau. Logarithmic Sobolev inequality for some models of random walks. *The Annals of Probability*, 26(4):1855 – 1873, 1998. 1

[MT06] Ravi Montenegro and Prasad Tetali. Mathematical aspects of mixing times in markov chains. *Found. Trends Theor. Comput. Sci.*, 1(3):237–354, July 2006. 2.3

[MY92] Mitsuru Matsui and Atsuhiro Yamagishi. A new method for known plaintext attack of FEAL cipher. In *EUROCRYPT*, volume 658 of *Lecture Notes in Computer Science*, pages 81–91. Springer, 1992. 1

[Nyb93] Kaisa Nyberg. Differentially uniform mappings for cryptography. In Tor Helleseth, editor, *Advances in Cryptology - EUROCRYPT '93, Workshop on the Theory and Application of of Cryptographic Techniques, Lofthus, Norway, May 23-27, 1993, Proceedings*, volume 765 of *Lecture Notes in Computer Science*, pages 55–64. Springer, 1993. 1

[PW11] Yuval Peres and Peter Winkler. Can extra updates delay mixing?, 2011. 1

[Sal20] Justin Salez. A sharp log-sobolev inequality for the multislice, 2020. https://arxiv.org/abs/2004.05833. 1, 2.17

[SC97] Laurent Saloff-Coste. *Lectures on finite Markov chains*, pages 301–413. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997. 2.2

[Sin92] Alistair Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992. 1, 2.6

[Sta98] Richard Stafford. Groups of permutation polynomials over finite fields. *Finite Fields and Their Applications*, 4(4):450–452, 1998. 1

[STY22] Justin Salez, Konstantin Tikhomirov, and Pierre Youssef. Upgrading mlsi to lsi for reversible markov chains, 2022. 1, 2.3, 2.10

[Top14] Alev Topuzoglu. The carlitz rank of permutations of finite fields: A survey. *Journal of Symbolic Computation*, 64:53–66, 2014. Mathematical and computer algebra techniques in cryptology. 1

[WLP09] EL Wilmer, David A Levin, and Yuval Peres. Markov chains and mixing times. *American Mathematical Soc., Providence*, 2009. 2.2

[Zie13] Michael E. Zieve. On a theorem of Carlitz. *Journal of Group Theory*, 17(4):667–669, November 2013. 1

A Missing Details on the Paths Constructions

A.1 Fields of even characteristic

Lemma A.1. *Let $n = 2^b$, for $b \geq 3$, and N_i be the number of valid ψ_{even} paths that generate the transposition $(0, i)$ for $i \in \mathbb{F}_n$. That is,*

$$N_i = \left|\{(u, v, w) \mid (0, i) = \psi_{\text{even}}(u, v, w), uv \neq 1, u(v+1) \neq 1, v \notin \{0, 1\}\}\right|.$$

Then

$$N_i = \begin{cases} 0 & i = 0 \\ n(2n-4) & i = 1 \\ n(n-4) & i \notin \{0, 1\}. \end{cases}$$

And in particular

$$\frac{1}{3} \leq \frac{N_i}{N_j} \leq 3$$

for any non-zero i, j.

Proof. Recall that a valid sequence $\psi_{\text{even}}(u, v, w)$ generates the transposition $\left(0, \frac{v}{uv+1} + \frac{v+1}{uv+u+1}\right)$. Then to generate $(0, i)$ it must hold that $\frac{v}{uv+1} + \frac{v+1}{uv+u+1} = i$. We rewrite this sum below

$$\begin{aligned} s = \frac{v}{uv+1} + \frac{v+1}{uv+u+1} &= \frac{v(uv+u+1) + (v+1)(uv+1)}{(uv+1)(uv+u+1)} \\ &= \frac{uv^2 + uv + v + uv^2 + v + uv + 1}{(uv+1)(uv+u+1)} = \frac{1}{(uv+1)(uv+u+1)}. \end{aligned}$$

From our constraints on u, v, the value of s is always non-zero. Now, if we set $u = 0$, all $n - 2$ valid values of v ($v \notin \{0, 1\}$) will satisfy $s = 1$.

Finally, when $u \neq 0$, we write

$$\begin{aligned} s &= \frac{1}{(uv+1)(uv+u+1)} \\ \implies \frac{1}{s} &= u^2 v(v+1) + u + 1 \\ \implies v(v+1) &= \frac{1/s + u + 1}{u^2}. \end{aligned}$$

It is well known that in characteristic 2, the quadratic equation above has 2 solutions if the trace of the right-hand side is equal to 0, and no solutions otherwise.

$$\operatorname{Tr}\left(\frac{1/s + u + 1}{u^2}\right) = \operatorname{Tr}\left(\frac{1/s}{u^2}\right) + \operatorname{Tr}\left(\frac{1}{u}\right) + \operatorname{Tr}\left(\frac{1}{u^2}\right) = \operatorname{Tr}\left(\frac{1/s}{u^2}\right)$$

Here the first equality follows from the linearity of trace and the second equality from the fact that $\operatorname{Tr}(x^2) = \operatorname{Tr}(x)$. The square is an injective map over $\mathbb{F}_n$, and thus $\frac{1/s}{u^2}$ obtains every non-zero value in the field. Since $\operatorname{Tr}(\cdot) = 0$ defines a subspace, the number of u's that make $\frac{1/s}{u^2}$ have 0 trace is exactly $n/2 - 1$ (we exclude zero, since $\frac{\kappa}{u^2}$ is never zero and $\operatorname{Tr}(0) = 0$).

All of these values are valid, except $u = 1/s + 1$, for $s \neq 1$. This is because even though

$$\operatorname{Tr}\left(\frac{1/s}{(1/s+1)^2}\right) = \operatorname{Tr}\left(\frac{1}{1/s+1} + \frac{1}{(1/s+1)^2}\right) = 0,$$

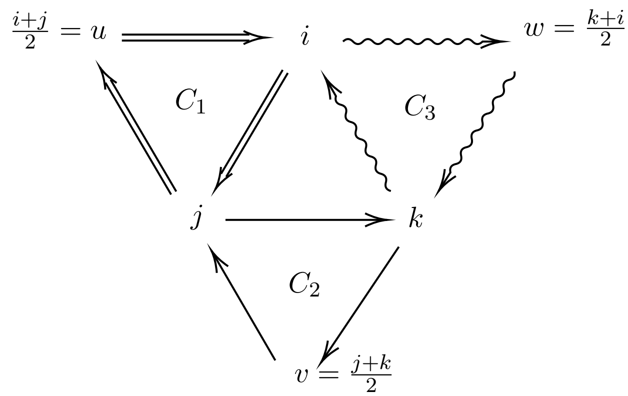

Figure 2: A schematic representation of the setting in Lemma A.2. The Claim implies that using the 3-cycles C_1, C_2, C_3 a finite number of times, we can generate the 3-cycle (i, j, k).

this implies that $v(v+1) = 0$ and thus the two solutions to this equation are $v = 0, 1$, which are not valid. The remaining $n/2 - 2$ values of u give 2 valid solutions for v, which means that all $s \notin \{0, 1\}$ have $n - 4$ solutions for non-zero u.

The case of $s = 1$ has the $n - 2$ solutions with $u = 0$ and 2 solutions for the $n/2 - 1$ non-zero valid values of u, for a total of $2n - 4$ solutions. We conclude the proof of this lemma by noting that the last parameter w can be chosen arbitrarily from the field of size n without changing the value of s. Thus the number of paths increases by a multiplicative factor of n.

□

A.2 Fields of odd characteristic

Lemma A.2. *Let $n = p^b$, where p is an odd prime. Let $i, j, k \in \mathbb{F}_n$ be distinct numbers that do not form an arithmetic progression. Moreover, let $u = \frac{i+j}{2}, v = \frac{j+k}{2}, w = \frac{k+i}{2}$. Consider the following 3-cycles acting on S_n.*

$$C_1 = (i, j, u), \quad C_2 = (j, k, v), \quad C_3 = (k, i, w).$$

Then

$$C_2^2 \circ C_1 \circ C_3 \circ C_2^2 \circ C_3^2 \circ C_1 = (i, k, j).$$

Proof of Lemma A.2. Since the C_i's are 3-cycles, applying C_i twice is equal to C_i^{-1}. We will thus compute the permutation $C_2^{-1} \circ C_1 \circ C_3 \circ C_2^{-1} \circ C_3^{-1} \circ C_1$ and show that it is equal to (i, k, j). Since the 3-cycles only touch the elements i, j, k, u, v, w, it suffices to restrict our attention to these 6 elements. Furthermore, it is convenient to arrange these 6 elements in a triangle as in Figure 2. In the following proof, we also use boldface to indicate the elements that are involved in the 3-cycle that will get applied. The statement follows by a straightforward calculation:

Applying $C_2^2 \circ C_1 \circ C_3 \circ C_2^2 \circ C_3^2 \circ C_1$ gives

$$\begin{array}{ccccc} \mathbf{u} & & \mathbf{i} & & w \\ & \mathbf{j} & & k & \\ & & v & & \end{array} \xrightarrow{C_1} \begin{array}{ccccc} j & & \mathbf{u} & & \mathbf{w} \\ & i & & \mathbf{k} & \\ & & v & & \end{array} \xrightarrow{C_3^{-1}} \begin{array}{ccccc} j & & w & & k \\ & \mathbf{i} & & \mathbf{u} & \\ & & \mathbf{v} & & \end{array}$$

$$\xrightarrow{C_2^{-1}} \begin{array}{ccccc} j & & \mathbf{w} & & \mathbf{k} \\ & u & & \mathbf{v} & \\ & & i & & \end{array} \xrightarrow{C_3} \begin{array}{ccccc} \mathbf{j} & & \mathbf{v} & & w \\ & \mathbf{u} & & k & \\ & & i & & \end{array}$$

$$\xrightarrow{C_1} \begin{array}{ccccc} u & & j & & w \\ & \mathbf{v} & & \mathbf{k} & \\ & & \mathbf{i} & & \end{array} \xrightarrow{C_2^{-1}} \begin{array}{ccccc} u & & j & & w \\ & k & & i & \\ & & v & & \end{array} .$$

Applying (i, j, k) gives

$$\begin{array}{ccccc} u & & \mathbf{i} & & w \\ & \mathbf{j} & & \mathbf{k} & \\ & & v & & \end{array} \xrightarrow{(i,k,j)} \begin{array}{ccccc} u & & j & & w \\ & k & & i & \\ & & v & & \end{array} .$$

□