

Shallow quantum shuffling via measurement-driven multi-scale rotations

Simone Faro^{1,†}, Francesco Pio Marino^{1,2,†} and Arianna Pavone^{3,‡}

¹Università di Catania, viale A. Doria n.6, 95125, Catania, Italy

²Univ Rouen Normandie, INSA Rouen Normandie, Université Le Havre Normandie, Normandie Univ, LITIS UR 4108, CNRS NormaSTIC FR 3638, IRIB, Rouen F-76000, France

³Università di Palermo, via Archirafi n.34, 90123, Palermo, Italy

Abstract

Shuffling under structural constraints highlights a fundamental tension between randomness and parallelism: while global mixing requires information to propagate across the entire system, shallow circuits inherently limit such propagation. In this work, we introduce a multi-scale randomized procedure for approximate shuffling within a measurement-driven shallow circuit model. The construction combines parametrized cyclic rotations with intermediate measurements, using locally generated randomness to drive structured data movement across a hierarchy of spatial scales. Each round applies independent random rotations on disjoint blocks, together with shifted alignments that promote interaction across block boundaries and progressively reduce dependence on the initial configuration. We show that the resulting process defines an ergodic Markov chain on the symmetric group with the uniform distribution as stationary measure, and we provide explicit convergence guarantees in terms of the spectral properties of a single round. In particular, we establish a non-trivial polynomial lower bound on the spectral gap, yielding explicit mixing guarantees for this class of multi-scale shuffling procedures. The multi-scale structure of the construction suggests substantially faster convergence than this baseline bound, potentially beyond diffusive behavior, which we leave as an open problem. The procedure can be implemented by a uniform circuit of polylogarithmic depth and near-linear size (up to polylogarithmic factors). Empirical evidence on small instances illustrates the qualitative mixing behavior of the construction. Beyond the specific application to shuffling, our results suggest that measurement-driven randomness, when combined with structured circuit primitives, provides a flexible framework for designing randomized computations under depth constraints.

Keywords

quantum circuits, shuffling, random permutations, parallel algorithms

1. Introduction

Randomness plays a central role in the design of efficient algorithms, particularly for sampling and approximate counting tasks [1]. In many computational models, however, randomness is treated as an external resource, supplied by an oracle or assumed to be freely available. This raises a natural question: how does the ability to *generate* and *reuse* randomness within a

ICTCS 2026: Italian Conference on Theoretical Computer Science, September 7–9, 2026, Udine, Italy

[†]Supported by University of Catania, progetto PIACERI 2024-2026 - Linea di Intervento 1

[‡]Supported by Progetto Eurostart 2026

✉ faro@dmf.unict.it (S. Faro); francesco.marino@phd.unict.it (F. P. Marino); ariannamaria.pavone@unipa.it (A. Pavone)

ORCID 0000-0001-5937-5796 (S. Faro); 0000-0003-4722-9542 (F. P. Marino); 0000-0002-8840-6157 (A. Pavone)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

computational model affect its algorithmic power, especially under structural constraints such as limited depth or restricted interaction patterns?

In the quantum circuit model, classical randomness can be obtained intrinsically via measurement [2, 3]. For instance, preparing a qubit in state $|0\rangle$, applying a Hadamard gate, and measuring yields an unbiased random bit in constant depth. More generally, mid-circuit measurements provide a mechanism to generate classical random variables and reuse them adaptively within a computation. This paradigm has been explored in several contexts, including measurement-based and hybrid quantum-classical computation [4, 5, 6, 7], and suggests a model of shallow computation where randomness is produced internally and interleaved with unitary operations.

In this work, we investigate the algorithmic consequences of this capability in the context of *approximate shuffling*. We consider the problem of transforming an ordered sequence of n items into a permutation whose distribution is ε -close to uniform over the symmetric group S_n , under constraints on circuit depth and data organization.

Shuffling under strict locality constraints is inherently limited. In classical models where interactions are restricted to nearest neighbors on a one-dimensional array, information propagates at a bounded speed, and any process that achieves near-uniform mixing requires $\Omega(n)$ depth. This phenomenon is well understood in the study of adjacent transposition shuffles and interchange processes on low-dimensional graphs [8, 9, 10], and reflects a structural limitation: randomness introduced at one location cannot influence distant positions without a linear number of sequential steps.

Our construction departs from these strictly local models. Rather than restricting interactions to a fixed radius, we allow operations on contiguous blocks of increasing size. This enables structured long-range information propagation within a shallow circuit, and places our approach in an intermediate regime between strictly local and fully global models. In particular, while each operation acts locally within a block, the hierarchy of block sizes allows interactions across arbitrary distances within polylogarithmic depth.

Rather than attempting to overcome locality limitations within fully deterministic or reversible frameworks, we adopt a different perspective. We study a model of *measurement-driven shallow computation*, in which intermediate measurements are treated as an explicit algorithmic resource. In this model, randomness is generated internally, stored classically, and used to control subsequent operations within a shallow circuit.

Our main contribution is a structured *multi-scale randomized shuffling procedure* that exploits this paradigm. The algorithm applies, in parallel, random cyclic rotations on blocks of increasing size, with alternating alignments to promote interaction across block boundaries. By iterating this multi-scale process, the procedure progressively reduces dependence on the initial configuration and induces a stochastic evolution on S_n .

From a complexity-theoretic perspective, our results highlight differences between models of shallow computation. Deterministic classical circuits with fixed-radius local interactions require linear depth to propagate information across the system, whereas allowing internally generated randomness together with structured multi-scale operations enables polylogarithmic-depth constructions for randomized transformations. We emphasize that our focus is not on exploiting uniquely quantum phenomena such as interference, but rather on understanding how the availability of *internally generated randomness* and structured non-local primitives affects the power of shallow circuits.

Our main contributions can be summarized as follows:

- We introduce a structured *multi-scale shuffling scheme* based on parallel cyclic rotations at geometrically increasing spatial scales. The construction is uniform, parameterizable, and tailored to shallow circuit implementations.
- We formalize a model of *measurement-driven randomized computation*, in which intermediate measurements are used to generate and reuse classical randomness within shallow circuits.
- We show that the proposed procedure can be implemented by circuits of depth $O(\log^3 n \cdot r)$ and near-linear size, where r is the number of rounds. The induced process defines an ergodic Markov chain on S_n with the uniform distribution as stationary measure. We establish a non-trivial polynomial lower bound on the spectral gap of a single round, yielding explicit (polynomial) mixing guarantees, and characterize convergence in terms of this gap.
- We contrast this setting with classical models under strict locality constraints, emphasizing that our construction operates in a relaxed locality regime enabled by multi-scale structured operations.

The remainder of the paper is organized as follows. Section 2 reviews classical limitations of shuffling and related work. Section 3 introduces the computational model and basic primitives, including the parametrized cyclic rotation construction. Section 4 describes the multi-scale shuffling algorithm together with its complexity and mixing analysis. Section 5 concludes with directions for further investigation.

2. Shuffling: classical limits and related work

Shuffling and randomization on permutations have been extensively studied in probability theory, combinatorics, and theoretical computer science. A classical approach models shuffling as a random walk on the symmetric group S_n , where the objective is to converge rapidly to the uniform distribution. Foundational work by Diaconis and Shahshahani [11] on random transpositions established sharp bounds on mixing times and highlighted the importance of global operations in achieving efficient randomization. Comprehensive treatments of mixing phenomena can be found in standard references such as [12, 13].

A key factor governing the efficiency of shuffling processes is the underlying interaction structure. In highly connected models such as random transpositions, mixing occurs in $\Theta(n \log n)$ steps. In contrast, when interactions are restricted by locality, the propagation of randomness becomes significantly slower. This effect is particularly pronounced in one-dimensional settings, where only neighboring elements can interact, and where shuffling processes can be interpreted as variants of the interchange process on a path graph.

It is well known that locality imposes strong limitations on mixing. Adjacent transposition shuffles require polynomially larger mixing times than their non-local counterparts [8, 10]. More generally, results on the interchange process [9, 14] show that the spectral properties of the underlying graph fundamentally constrain convergence rates. Even when randomness is

injected at every step, locality enforces a form of sequential information propagation, leading in one dimension to diffusive mixing behavior.

A related line of work studies shuffling schemes with structured or partially deterministic updates. Mossel, Peres, and Sinclair [15] analyze semi-random transposition shuffles, showing that mixing remains on the order of $n \log n$. In the one-dimensional setting, Morris [16] studies the Card-Cyclic-to-Random shuffle and establishes that locality constraints enforce essentially sequential behavior despite the presence of randomness. Collectively, these results support the principle that, under locality constraints, efficient global mixing requires sufficient depth to propagate randomness across the entire system.

Quantum and measurement-driven approaches. The quantum setting offers additional primitives, most notably the ability to generate classical randomness via measurement and to interleave classical control with unitary evolution. This raises the question of whether such capabilities can be leveraged to design efficient randomized procedures under depth constraints.

Several works have explored quantum analogues of classical shuffling processes. Quantum walk-based approaches study mixing behavior on highly connected graphs, where quantum dynamics may exhibit different convergence properties than classical random walks. However, these results typically focus on abstract graph models and do not directly address circuit complexity or implementations under locality constraints.

Another line of research investigates shallow quantum circuits with intermediate measurements as a computational resource. Measurement-based constructions, including hierarchical architectures such as quantum convolutional circuits [5], demonstrate that classical randomness can be generated and reused within low-depth quantum processes. These works suggest that combining shallow circuit structure with measurement-driven randomness may enable new forms of efficient randomized computation.

Circuit complexity of permutations. From the perspective of circuit complexity, the problem of implementing permutations has been studied in both classical and quantum settings. It is known that any fixed permutation on n items can be realized by a logspace-uniform quantum circuit of constant depth [17], placing arbitrary permutations in QNC^0 , while classical constant-depth circuit classes cannot realize nontrivial routing operations. Such constructions are inherently non-parametric: each circuit implements a specific permutation and cannot be efficiently reused to sample from a distribution over permutations. For randomized tasks such as shuffling, one instead requires parameterized and reusable primitives that can be controlled by random inputs.

In this work, we build on the parametrized cyclic rotation gate introduced in [18], which implements rotations by arbitrary offsets in depth $O(\log n)$ using a uniform and parallelizable structure. These properties make it a natural building block for the multi-scale randomized constructions studied in this paper.

Overall, classical locality constraints lead to diffusive mixing, while measurement-driven quantum models enable randomness within shallow circuits. We combine structured primitives and internally generated randomness to achieve approximate shuffling in polylogarithmic depth, proving a non-trivial polynomial lower bound on the spectral gap. Whether the multi-scale structure yields substantially faster mixing remains open.

3. Computational model and basic primitives

In this section we formalize the computational model underlying our construction and introduce the basic primitives used in the shuffling procedure. Our focus is on shallow circuits augmented with the ability to generate and reuse classical randomness via intermediate measurements.

We work in the standard quantum circuit model with bounded fan-in gates and uniform circuit families. A circuit is a sequence of layers of operations acting on disjoint qubits in parallel. The depth is the number of layers and measures parallel time, while the size is the total number of gates. We assume a conventional universal gate set and unrestricted connectivity, allowing gates to act on arbitrary pairs of qubits within each layer.

A key feature of our model is the use of intermediate measurements. In addition to unitary gates, circuits may include measurements in the computational basis, whose outcomes are stored as classical bits and used to control subsequent operations. We assume that classical control can be implemented without asymptotic overhead in depth.

Measurement enables the internal generation of classical randomness. For example, preparing $|0\rangle$, applying a Hadamard gate, and measuring yields an unbiased random bit in constant depth. More generally, the computation defines a randomized process whose output distribution depends on measurement outcomes, making this model naturally suited for sampling tasks.

For comparison, in classical nearest-neighbor circuits on a one-dimensional array, information propagates locally and global dependencies require linear depth. In contrast, the present model allows randomness to be generated locally and reused globally via classical control.

We represent a deck of n items as a quantum data register $D = (D[0], D[1], \dots, D[n-1])$, where each entry $D[i]$ encodes a symbol from an alphabet Σ of size σ . Each symbol is stored using $\lceil \log \sigma \rceil$ qubits, so the total number of qubits in the data register is $n \lceil \log \sigma \rceil$. A classical configuration corresponds to a computational basis state $|\pi\rangle$, where $\pi \in S_n$ specifies which symbol occupies each position. A (classical) shuffle samples a permutation σ from a distribution $\mathcal{D}$ over S_n and applies the transformation $|\pi\rangle \mapsto |\sigma \circ \pi\rangle$.

Our objective is to implement, using a measurement-driven shallow circuit, a randomized transformation whose output distribution is close to uniform over S_n . More precisely, let μ denote the distribution over permutations induced by the circuit, and let π be the uniform distribution on S_n . We measure the quality of the shuffle using total variation distance:

$$\|\mu - \pi\|_{\text{TV}} = \frac{1}{2} \sum_{\sigma \in S_n} |\mu(\sigma) - \pi(\sigma)|.$$

Given an accuracy parameter $\varepsilon \in (0, 1)$, the goal is to construct a circuit such that $\|\mu - \pi\|_{\text{TV}} \leq \varepsilon$, while minimizing circuit depth.

Efficient manipulation of permutations within shallow circuits requires structured primitives that are both *uniform* and *parameterizable*. While it is known that any fixed permutation on n items can be implemented by a logspace-uniform quantum circuit of constant depth [17], such constructions are inherently non-parametric: each circuit is tailored to a specific permutation and cannot be efficiently reused to sample from a family of permutations.

For randomized tasks such as shuffling, one instead requires primitives that can be controlled by external parameters and composed in a regular manner. In this work, we adopt *cyclic rotations* as a basic building block.

A (rightward) cyclic rotation by offset k on a register of size n is the permutation R_k defined by $R_k(x) = (x + k) \bmod n$, which preserves the relative order of elements modulo a global shift. As a transformation on the data register, R_k induces a permutation of computational basis states, extended linearly to a unitary operator. Cyclic rotations are a convenient primitive due to their simple algebraic structure and their ability to generate nontrivial global rearrangements.

We rely on the construction of Pavone and Viola [18], which provides a logspace-uniform circuit family implementing cyclic rotations on a register of size n with depth $O(\log n)$. The construction proceeds iteratively: at each layer, a carefully chosen set of disjoint swap operations is applied in parallel, placing a constant fraction of the elements into their final positions. As a consequence, the number of misplaced elements decreases geometrically across layers, yielding logarithmic depth overall. This construction is uniform and admits efficient parallel execution across disjoint subregisters.

A key requirement for our application is that cyclic rotations admit an efficient *parameterized* implementation. We consider a circuit acting on a data register $q = (q_0, \dots, q_{n-1})$ of size n , together with a classical parameter $j \in \{0, \dots, n-1\}$ specifying the rotation offset, and denote by ROT_j the operation that applies the cyclic rotation R_j .

To implement ROT_j , we decompose the offset j into its binary representation. For each bit position i , with $0 \leq i < \lceil \log n \rceil$, a rotation by 2^i positions is applied conditionally on the value of the i -th bit of j . Each such subrotation can be realized using the logarithmic-depth construction described above. Composing these subrotations yields a uniform implementation of ROT_j with depth $O(\log^2 n)$.

4. Multi-scale shuffling procedure

Shuffling processes are constrained by the rate at which information can propagate across the system. In classical nearest-neighbor models, interactions are local and influence spreads at most linearly with circuit depth. As a consequence, achieving global mixing requires depth at least linear in n in such settings.

In this section we present a randomized shuffling procedure implementable within the measurement-driven shallow circuit model introduced in Section 3. The procedure produces a permutation whose distribution approaches uniformity over S_n , and is designed to operate in shallow depth by combining structured data-movement primitives with generated randomness.

A central difficulty in designing shallow shuffling procedures is that purely local randomization does not eliminate global structure within a small number of layers. To address this, our construction adopts a *multi-scale strategy*: randomized operations are applied across a hierarchy of spatial scales, progressively propagating local perturbations to larger regions.

For simplicity, we assume that n is a power of two. The general case can be handled by padding with dummy items up to the next power of two and discarding them after the shuffle.

At each scale, the data register is partitioned into disjoint blocks, and independent random cyclic rotations are applied in parallel within each block. To promote interaction across block boundaries, each scale is processed using two sub-layers with shifted block alignments. This organization ensures that elements that are separated at one scale may interact at another, enabling global mixing over multiple scales.

Input: Data register $D[0..n-1]$	Input: Parameter register K of $n/2$ qubits
Input: Number of rounds $r \in \mathbb{N}$	Output: Permutation of D
<pre> 1 for $t = 1$ to r do 2 for $k = 2, 4, 8, \dots, n$ do 3 $m \leftarrow n/k$ 4 partition K into $K_0, \dots, K_{m-1}$ of $\log k$ qubits 5 for $j = 0$ to $m-1$ in parallel 6 sample $s_j \in \{0, \dots, k-1\}$ via measurement 7 apply R_{s_j} to block B_j 8 for $j = 0$ to $m-1$ in parallel 9 sample $s'_j \in \{0, \dots, k-1\}$ 10 apply $R_{s'_j}$ to shifted block B'_j 11 return D </pre>	

Figure 1: Multi-scale shuffling procedure. The algorithm performs r rounds of structured random rotations across multiple spatial scales. The value of r is chosen as a function of ε (see Section 4.2).

The algorithm (described in Figure 1) operates on a data register $D = (D[0], D[1], \dots, D[n-1])$ and uses an auxiliary parameter register K to generate random offsets. Each round traverses a sequence of scales $k = 2, 4, 8, \dots, n$. At scale k , the register is partitioned into $m = n/k$ contiguous blocks of size k , and each block undergoes a cyclic rotation by a randomly sampled offset. All rotations at a given scale are performed in parallel.

The number of rounds r will be chosen as a function of the desired accuracy ε , based on the mixing bounds established in Section 4.2.

To enable interaction between adjacent blocks, a second sub-layer applies the same procedure to blocks shifted by $k/2$. The parameter register K is reused across all scales and rounds: after each measurement, qubits collapse to classical states and can be reused without overhead.

The effectiveness of the procedure relies on the interaction between scales. At small scales, randomness is injected locally within blocks, while at larger scales these local perturbations are combined across progressively larger regions. After $\log n$ scales, each element is affected—directly or indirectly—by operations spanning the entire register. The alternation between aligned and shifted partitions ensures that dependencies are not confined within fixed block boundaries. Figure 2 illustrates a single round of the procedure.

A single round of the procedure does not produce a uniform permutation. However, the induced process defines a Markov chain on S_n that admits the uniform distribution as a stationary distribution. Repeating the procedure for r rounds progressively reduces the distance from uniform; quantitative convergence guarantees will be established in Section 4.2.

The construction combines structured data movement via cyclic rotations, parallel execution across disjoint blocks, and internally generated randomness via measurement. These ingredients enable randomized transformations that operate within shallow circuit depth. We emphasize that the procedure targets approximate rather than exact uniformity, as exact uniform sampling under locality constraints would require maintaining global dependencies incompatible with shallow depth.




Figure 2: One round of the multi-scale procedure for $n = 8$. Each scale applies independent random rotations on disjoint blocks, followed by a shifted layer.

4.1. Circuit size and depth analysis

We analyze the circuit complexity of the algorithm in Figure 1, focusing on depth and size.

Depth. Fix a scale k . Each scale consists of two sub-layers (aligned and shifted), where independent parametrized cyclic rotations are applied in parallel over disjoint blocks.

Generating $\log k$ random bits and applying classical control both require constant depth. The dominant cost is the parametrized rotation: a rotation on a block of size k is implemented via $\Theta(\log k)$ controlled subrotations, each of depth $O(\log k)$, yielding depth $O(\log^2 k)$ per sub-layer. Since blocks are disjoint, all rotations at a given scale are executed in parallel. Thus, each scale contributes depth $O(\log^2 k)$. Summing over scales $k = 2, 4, \dots, n$, we obtain

$$\sum_{i=1}^{\log n} O(i^2) = O(\log^3 n)$$

per round, and hence total depth $O(\log^3 n \cdot r)$ over r rounds.

Size. At scale k , there are n/k blocks, each undergoing two parametrized rotations. A rotation on a block of size k has size $O(k \log^2 k)$, yielding total size $O(n \log^2 k)$ at scale k .

Summing over all scales gives

$$\sum_{i=1}^{\log n} O(ni^2) = O(n \log^3 n)$$

per round, and total size $O(n \log^3 n \cdot r)$ over r rounds.

4.2. Mixing analysis and ε -close guarantees

The randomized procedure induces a stochastic process on the symmetric group S_n . Each round maps an input permutation to a random output permutation, and repeated rounds correspond to iterating this transition. Accordingly, the evolution of the output distribution can be modeled as a Markov chain on S_n .

Let μ_r denote the distribution after r rounds, and let π denote the uniform distribution on S_n . Let P be the Markov kernel corresponding to a single round, so that $\mu_r = \mu_0 P^r$.

Lemma 1 (Stationarity). *The uniform distribution π is stationary for P .*

Proof. Each sub-layer applies independent random cyclic rotations on disjoint blocks, where offsets are chosen uniformly from $\{0, \dots, k-1\}$. For each block, this induces a distribution that is uniform over a cyclic group, and hence preserves the uniform distribution locally.

Since the operations on disjoint blocks act independently, each sub-layer defines a stochastic operator that preserves the uniform distribution over S_n . As a composition of such operators, the transition kernel P also preserves π , implying stationarity. $\square$

Lemma 2 (Irreducibility and aperiodicity). *The Markov chain defined by P is irreducible and aperiodic.*

Proof. Aperiodicity follows since the identity permutation occurs with positive probability (all offsets equal to zero), implying $P(\sigma, \sigma) > 0$ for all σ . For irreducibility, observe that at scale $k=2$, the aligned and shifted sub-layers allow any adjacent transposition $(i \ i+1)$ to occur with positive probability. Since adjacent transpositions generate S_n , the chain is irreducible. $\square$

Lemmas 1 and 2 imply that P defines an ergodic Markov chain with unique stationary distribution π .

Convergence via spectral properties. Let λ_* denote the second largest singular value of P in $L^2(\pi)$, and define $\gamma = 1 - \lambda_*$.

Theorem 4.1 (Spectral mixing bound). *For every initial distribution μ_0 and every $r \geq 1$,*

$$\|\mu_0 P^r - \pi\|_{\text{TV}} \leq \frac{1}{2} \sqrt{\frac{1}{\pi_{\min}}} (1 - \gamma)^r = \frac{1}{2} \sqrt{n!} (1 - \gamma)^r.$$

Proof. Let $f_0 = \mu_0/\pi - 1$, which satisfies $\mathbb{E}_\pi[f_0] = 0$, so $f_0 \in L_0^2(\pi)$. Let P^* denote the adjoint of P in $L^2(\pi)$. The density evolves under the adjoint operator:

$$\frac{\mu_0 P^r}{\pi} - 1 = (P^*)^r f_0.$$

By definition of λ_* as the operator norm of P on $L_0^2(\pi)$, the same bound holds for P^* , hence we obtain $\|(P^*)^r f_0\|_{2,\pi} \leq \lambda_*^r \|f_0\|_{2,\pi}$. Since $\lambda_* = 1 - \gamma$, this yields contraction by $(1 - \gamma)^r$.

Using the standard inequality $\|\nu - \pi\|_{\text{TV}} \leq \frac{1}{2} \|\nu/\pi - 1\|_{2,\pi}$, we obtain the following relation: $\|\mu_0 P^r - \pi\|_{\text{TV}} \leq \frac{1}{2} \|(P^*)^r f_0\|_{2,\pi} \leq \frac{1}{2} (1 - \gamma)^r \|f_0\|_{2,\pi}$.

Finally, since μ_0 is a probability distribution, one has $\|f_0\|_{2,\pi} \leq \sqrt{1/\pi_{\min}}$, which yields the claimed bound. $\square$

Lemma 3. *Let Q_2 be the transition kernel induced by the scale-2 sub-layers. Then we have that $1 - \lambda_*(Q_2) = \Omega(1/n^2)$.*

Proof. At scale $k = 2$, each block undergoes either the identity or an adjacent transposition with constant probability. The aligned and shifted sub-layers correspond to applying random swaps on the two alternating perfect matchings of the path (or cycle), yielding a parallel adjacent-transposition dynamics.

Over two consecutive sub-layers, each edge of the path is updated with constant probability, so this dynamics performs a constant-rate update of all adjacent transpositions. Its Dirichlet form is therefore comparable, up to constant factors, to that of the interchange process on the path (or cycle). It is well known that the spectral gap of the interchange process on a path or cycle is $\Theta(1/n^2)$ [9, 14]. By comparison, the same lower bound holds for the parallel alternating-matching dynamics, yielding $1 - \lambda_*(Q_2) = \Omega(1/n^2)$. $\square$

A formal comparison can be obtained via standard comparison techniques for Markov chains (see, e.g., [12]), which imply that the spectral gap of the parallel dynamics is within constant factors of that of the interchange process.

Proposition 1. *Let P be the transition kernel corresponding to one round of the multi-scale procedure, and let $\gamma = 1 - \lambda_*$. Then*

$$\gamma \geq \Omega\left(\frac{1}{n^2}\right).$$

Proof. Write the one-round kernel as $P = Q_2 R$, where Q_2 is the scale-2 component and R is the composition of all remaining layers. Each remaining layer is a convex combination of permutation operators preserving π , and therefore has operator norm at most 1 on $L_0^2(\pi)$. Therefore, $\lambda_*(P) = \|P\|_{L_0^2(\pi) \rightarrow L_0^2(\pi)} \leq \|Q_2\|_{L_0^2(\pi) \rightarrow L_0^2(\pi)} \|R\|_{L_0^2(\pi) \rightarrow L_0^2(\pi)} \leq \lambda_*(Q_2)$.

By Lemma 3 we have $1 - \lambda_*(Q_2) = \Omega(1/n^2)$.

Hence $\gamma(P) = 1 - \lambda_*(P) \geq 1 - \lambda_*(Q_2) = \Omega(1/n^2)$. $\square$

Corollary 1. *Under the bound $\gamma = \Omega(1/n^2)$, it suffices to take $r = O(n^2(n \log n + \log(1/\varepsilon)))$ to guarantee $\|\mu_r - \pi\|_{\text{TV}} \leq \varepsilon$.*

Proof. By Theorem 4.1, we have $\|\mu_r - \pi\|_{\text{TV}} \leq \frac{1}{2} \sqrt{1/\pi_{\min}} (1 - \gamma)^r$. Using $\pi_{\min} = 1/n!$ and $\log(n!) = \Theta(n \log n)$, it suffices to take $(1 - \gamma)^r \leq \varepsilon/\sqrt{n!}$. Taking logarithms and using $\log(1/(1 - \gamma)) = \Theta(\gamma)$ for small γ , we obtain $r = O\left(\frac{1}{\gamma}(n \log n + \log(1/\varepsilon))\right)$. Substituting $\gamma = \Omega(1/n^2)$ yields the claim. $\square$

Proposition 1 provides a nontrivial, unconditional lower bound on the spectral gap of a single round, implying polynomial mixing time. However, this bound only captures the contribution of the smallest scale and does not reflect the full multi-scale structure of the algorithm.

Each round combines $\Theta(\log n)$ layers of structured randomization across scales ranging from local to global. This suggests that the true spectral gap may be significantly larger, potentially independent of n .

Open direction. Establishing stronger lower bounds on γ remains an important open problem. Possible approaches include coupling arguments tailored to the multi-scale dynamics, comparison techniques with well-understood Markov chains, or log-Sobolev inequalities adapted to hierarchical update schemes.

We observe that, if γ is bounded below by a constant, then the number of rounds required to achieve ε -closeness reduces to $r = O(\log(1/\varepsilon))$, yielding exponentially fast convergence.

4.3. Empirical mixing behavior: a small-scale simulation

To complement the theoretical analysis, we report a small-scale simulation of the proposed shuffling procedure, aimed at illustrating its qualitative behavior.

Due to the exponential growth of the permutation space, simulations were performed on a small instance with $n = 4$ elements. While limited in size, this setting already admits $4! = 24$ distinct permutations and suffices to observe nontrivial distributional effects.

We consider a data register $v = [0, 1, 2, 3]$, where each element is encoded using 2 qubits, for a total of 8 data qubits, in addition to the auxiliary parameter register. The simulation was implemented in Qiskit and executed in a noiseless setting. The multi-scale shuffling procedure was applied for $r \in \{1, 2, 3, 4, 5, 6\}$ rounds. After the final round, the data register was measured in the computational basis, and outcomes were interpreted as permutations of the input.

The experiment was repeated for 10,000 shots. All 24 permutations were observed. Under the uniform distribution, each permutation is expected to appear $10,000/24 \approx 416.7$ times.

Figure 3 shows the empirical distribution of permutations for selected values of r . For $r = 1$, the distribution exhibits significant bias, with some permutations occurring substantially more frequently than others. At $r = 2$, the distribution is already more balanced, though deviations from uniformity remain visible. For $r \geq 3$, the empirical frequencies appear close to uniform, with deviations consistent with finite-sample variability. Increasing the number of rounds beyond this point does not produce qualitative changes.

These observations provide qualitative evidence that a small number of rounds suffices to reduce visible structure in the distribution. Given the small system size, the results should be interpreted as illustrative rather than quantitative, and do not yield rigorous bounds on mixing time or spectral properties.

5. conclusion and open problems

We studied approximate shuffling under depth constraints and introduced a structured multi-scale randomized procedure based on cyclic rotations and measurement-driven randomness. The construction is uniform and parameterizable, and each round can be implemented within polylogarithmic circuit depth.

The induced process defines an ergodic Markov chain on S_n with the uniform distribution as stationary measure. We established a nontrivial polynomial lower bound on the spectral gap, $\gamma = \Omega(1/n^2)$, by relating the first scale of the construction to the interchange process. This provides a baseline guarantee for mixing, although it does not capture the full effect of the multi-scale structure.




Figure 3: Empirical evolution of the shuffling distribution on S_4 . Each panel shows the histogram of permutation counts after r rounds of the multi-scale quantum shuffling procedure (10,000 shots). Permutations are indexed in a fixed order; the dashed line marks the uniform expectation $10,000/24 \approx 416.7$. Strong bias at $r = 1$ rapidly disappears, and by $r = 3$ the distribution appears already close to uniform; $r = 6$ shows only finite-sampling fluctuations.

A central open problem is to obtain sharper bounds on the spectral gap of the full transition operator. While our analysis establishes a polynomial lower bound, the hierarchical organization of the algorithm suggests substantially faster mixing, potentially corresponding to a constant spectral gap.

More generally, it would be interesting to better understand how the multi-scale structure interacts with the underlying interaction graph, for instance by extending the construction to higher-dimensional grids or more general topologies. Another important direction is to clarify the role of measurement-driven randomness in shallow circuit models, and how it compares to classical randomized computation under similar depth constraints. Finally, exploring alternative primitives or scheduling strategies within the multi-scale framework may lead to improved mixing behavior or more efficient implementations.

More broadly, our results suggest that combining structured data movement with internally generated randomness provides a promising approach to designing efficient randomized procedures under depth constraints.

Declaration on Generative AI

Generative AI was used exclusively to assist with the linguistic refinement of the manuscript, including the rephrasing of selected sentences and paragraphs to improve clarity, conciseness, and readability.

References

- [1] R. Motwani, P. Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995.
- [2] S. Faro, F. P. Marino, G. Messina, Extending qutes: a practical high-level language for quantum computing, *Comput. J.* 69 (2026) 597–612. URL: <https://doi.org/10.1093/comjnl/bxaf133>. doi:10.1093/COMJNL/BXAF133.
- [3] S. Faro, F. P. Marino, G. Messina, Qutes: A high-level quantum programming language for simplified quantum computing, in: R. W. Wisniewski, I. Brandic (Eds.), *Proceedings of the 34th International Symposium on High-Performance Parallel and Distributed Computing, HPDC 2025, University of Notre Dame Conference Facilities, Notre Dame, IN, USA, July 20–23, 2025*, ACM, 2025, pp. 47:1–47:9. URL: <https://doi.org/10.1145/3731545.3744153>. doi:10.1145/3731545.3744153.
- [4] M. A. Nielsen, I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2000.
- [5] I. Cong, S. Choi, M. D. Lukin, Quantum convolutional neural networks, *Nature Physics* 15 (2019) 1273–1278. URL: <https://arxiv.org/abs/1810.03787>. doi:10.1038/s41567-019-0648-8.
- [6] Álvaro Bermejo, et al., Quantum convolutional neural networks are effectively classically simulable, *arXiv preprint* (2024). URL: <https://arxiv.org/abs/2408.12739>. arXiv:2408.12739.
- [7] X. Cao, J. Eisert, Measurement-driven quantum advantages in shallow circuits, *arXiv preprint* (2025). URL: <https://arxiv.org/abs/2505.04705>. arXiv:2505.04705.
- [8] D. B. Wilson, Mixing times of lozenge tiling and card shuffling markov chains, *The Annals of Applied Probability* 14 (2004) 274–325. URL: <https://projecteuclid.org/journals/annals-of-applied-probability/volume-14/issue-1/Mixing-times-of-lozenge-tiling-and-card-shuffling-Markov-chains/10.1214/aoap/1075828054.pdf>. doi:10.1214/aoap/1075828054.
- [9] P. Caputo, T. M. Liggett, T. Richthammer, Proof of aldous’ spectral gap conjecture, *Journal of the American Mathematical Society* 23 (2010) 831–851. URL: <https://arxiv.org/abs/0906.1238>. doi:10.1090/S0894-0347-10-00667-6.
- [10] H. Lacoïn, Mixing time and cutoff for the adjacent transposition shuffle and the simple exclusion, *The Annals of Probability* 44 (2016) 1426–1487. URL: <https://arxiv.org/abs/1309.3873>. doi:10.1214/15-AOP1002.
- [11] P. Diaconis, M. Shahshahani, Generating a random permutation with random transpositions, *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 57 (1981) 159–179. URL: <https://doi.org/10.1007/BF00535487>. doi:10.1007/BF00535487.
- [12] D. A. Levin, Y. Peres, E. L. Wilmer, *Markov Chains and Mixing Times*, American Math-

- emational Society, American Mathematical Society, Providence, RI, 2009. URL: <https://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf>.
- [13] L. Saloff-Coste, Random walks on finite groups, in: H. Kesten (Ed.), *Probability on Discrete Structures*, volume 110 of *Encyclopaedia of Mathematical Sciences*, Springer, 2004, pp. 263–346. doi:10.1007/978-3-662-09444-0_5.
 - [14] J. Jonasson, Mixing times for the interchange process, *ALEA Latin American Journal of Probability and Mathematical Statistics* 2 (2002) 1–15. URL: <https://www.math.chalmers.se/~jonasson/mip.pdf>.
 - [15] E. Mossel, Y. Peres, A. Sinclair, Shuffling by semi-random transpositions, *SIAM Journal on Computing* 35 (2004) 668–682. URL: <https://arxiv.org/abs/math/0404438>. doi:10.1137/S0097539704446131.
 - [16] B. Morris, The mixing time for the card-cyclic-to-random shuffle, *Random Structures & Algorithms* 40 (2012) 356–367.
 - [17] S. Faro, A. Pavone, C. Viola, Families of constant-depth quantum circuits for rotations and permutations, in: U. de’Liguoro, M. Palazzo, L. Roversi (Eds.), *Proceedings of the 25th Italian Conference on Theoretical Computer Science*, Torino, Italy, September 11–13, 2024, volume 3811 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024, pp. 29–41. URL: <https://ceur-ws.org/Vol-3811/paper270.pdf>.
 - [18] A. Pavone, C. Viola, The quantum cyclic rotation gate, *SN Comput. Sci.* 5 (2024) 830. URL: <https://doi.org/10.1007/s42979-024-03141-4>. doi:10.1007/s42979-024-03141-4.