

# Measurement-Induced Non-Commutativity in Adaptive Fermionic Linear Optics

Chenfeng Cao<sup>12</sup>, Yifan Tang<sup>1</sup>, Jens Eisert<sup>13</sup>

<sup>1</sup> *Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*

<sup>2</sup> *HK Institute of Quantum Science & Technology, The University of Hong Kong, Hong Kong, China*

<sup>3</sup> *Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany*

**Abstract.** Fermionic linear optics (FLO) with Gaussian resources is efficiently classically simulable, with many-body amplitudes reducing to determinants and Pfaffians. We show this is no longer true when fermions carry internal  $d$ -level labels and undergo mid-circuit number monitoring with classical feedforward. The measurement record routes the selected blocks into a fixed Bell-fusion geometry that ties teleportation byproducts to specific measured blocks: the antisymmetrized permutation sum no longer collapses to a single determinant or Pfaffian, and each post-selected branch amplitude becomes a matrix element of an ordered non-commutative trace polynomial in outcome-dressed FLO blocks. We prove an exponential sequential-memory barrier for any fusion-order-respecting MPO contraction in the free trace-polynomial benchmark, and numerically observe Porter-Thomas statistics together with rapid bond-dimension growth. Under a standard average-case hardness conjecture, this establishes mid-circuit measurement-induced non-commutativity as a route to sampling hardness for noninteracting fermions, without coherent two-body interactions in the FLO evolution.

**Keywords:** fermionic linear optics, mid-circuit measurement, quantum advantage

Full proofs are available at arXiv:2603.24950.

## 1 Main Idea and Results

Quantum random sampling protocols aim to outperform classical computers using restricted dynamics whose output distributions encode  $\#P$ -hard matrix functions [1, 2]. Amplitudes are permanents in boson sampling [3], Hafnians in Gaussian boson sampling [4], and Ising partition functions in commuting (IQP) circuits [5]. By sharp contrast, fermionic linear optics (FLO) with Gaussian inputs is efficiently simulable: amplitudes reduce to determinants and Pfaffians, and matchgate circuits admit polynomial-time simulation [6, 7, 8, 9, 10]. Closing the gap to a fermionic sampling advantage has so far required non-Gaussian resources such as magic input states or interaction gadgets [11, 12].

This work proposes a complementary route: monitoring and classical feedforward on otherwise Gaussian fermion evolution. Fermions carry internal  $d$ -level labels and propagate through a number-conserving FLO circuit. Mid-circuit coarse-grained occupation monitoring records which spatial blocks are singly occupied, without resolving the internal orbital, and the collision-free record routes the selected blocks into a fixed Bell-fusion pairing geometry. Pinning teleportation byproducts to fixed fusion steps enforces a contraction order that prevents

the antisymmetrized permutation sum from collapsing to a determinant or Pfaffian: each branch amplitude becomes the matrix element of an ordered non-commutative trace polynomial in outcome-dressed FLO blocks. This single-occupation projection lies outside standard mode-resolved Gaussian/FLO simulation frameworks [7, 9].

The results are threefold. *First*, the FLO layer induces a determinantal operator kernel  $\det_{\otimes}(\mathbf{S})$  on the  $n$  selected blocks. *Second*, the fixed-order Bell fusion converts this kernel into a branch operator  $\mathcal{T}_{\beta}(\mathbf{S})$  admitting a unique path-cycle decomposition: each permutation contributes one open boundary path (an ordered product of dressed blocks) and disjoint closed loops (scalar trace factors). *Third*, any fusion-order-respecting MPO representation of  $\mathcal{T}_{\beta}(\mathbf{S})$  has bond dimension at least  $2^{\Omega(n)}$  in the free trace-polynomial benchmark. With numerical Porter-Thomas statistics and rapid bond-dimension growth, an average-case hardness conjecture yields conditional sampling hardness via Stockmeyer’s approximate counting [13, 3, 2].

## 2 Non-Commutative Branch Operators

The architecture (Fig. 1) carries each spatial block by  $d$  fermionic orbitals coupled to auxiliary registers. A number-conserving FLO circuit acts on the  $dm$  encoded orbitals with single-particle propaga-




Figure 1: Monitored adaptive FLO architecture. Occupied input blocks are Bell-entangled with an auxiliary chain. A number-conserving FLO circuit (blue) with propagator  $\mathbf{V}$  acts on the encoded orbitals. Mid-circuit monitoring (MCM) produces a coarse-grained occupation record  $\mathbf{c}$ ; collision-free records select  $n$  blocks and fix an  $n \times n$  block sub-matrix  $\mathbf{S}$ , after which feedforward routes them into a fixed nearest-neighbor Bell-fusion geometry with outcomes  $\beta$ .

tor  $\mathbf{V} \in \text{U}(dm)$ , viewed as an  $m \times m$  block matrix  $\mathbf{V} = (V_{j,k})$  with  $V_{j,k} \in \text{M}_d(\mathbb{C})$ . Post-selecting on collision-free monitoring records selects  $n$  effective blocks  $\ell_1 < \dots < \ell_n$  and fixes the post-selected block sub-matrix  $\mathbf{S} = (S_{t,k})_{t,k=1}^n$  with  $S_{t,k} := V_{\ell_t, m-n+k}$ . In the dilute regime  $m = \kappa n^2$  the Haar-averaged collision-free monitoring rate tends to  $p_\infty(\kappa, d) = \exp[-(1 - 1/d)/(2\kappa)]$  and is therefore constant in  $n$ ; per-instance rates concentrate around this value.

**Lemma 1 (FLO kernel)** *Conditioned on a collision-free record, the encoded FLO layer induces on the selected  $n$  blocks the operator*

$$\det_{\otimes}(\mathbf{S}) = \sum_{\sigma \in \mathfrak{S}_n} \text{sgn}(\sigma) \left( \bigotimes_{t=1}^n S_{t, \sigma(t)} \right) \mathcal{P}_{\sigma^{-1}}, \quad (1)$$

where  $\mathcal{P}_\sigma$  permutes  $n$  tensor factors.

The proof tracks fermionic anticommutation in the second-quantized expansion of  $U(\mathbf{V})|\Psi_{\text{in}}\rangle$ ; the signed sum is the residue of single-occupation post-selection.

Each Bell-fusion outcome  $\beta_t = (a_t, b_t) \in \mathbb{Z}_d^2$  implements an unnormalized teleportation update  $\xi \mapsto d^{-1} \xi^\top \hat{D}_{\beta_t}$  on the boundary leg, with  $\hat{D}_\beta := X^a Z^b$  the Weyl operator. Absorbing the byproducts into dressed blocks  $\tilde{S}_{t,k} := S_{t,k}^\top \hat{D}_{\beta_t}$ , the post-selected fusion layer defines a linear contraction map  $\mathcal{C}_\beta$  that

absorbs the fixed  $d^{-1}$  factors, and the branch operator

$$\mathcal{T}_\beta(\mathbf{S}) := \mathcal{C}_\beta(\det_{\otimes}(\mathbf{S})) \in \text{M}_d(\mathbb{C}), \quad (2)$$

with branch amplitude  $\mathcal{A}_d(\mathbf{c}, \beta) = \langle r | \mathcal{T}_\beta(\mathbf{S}) | \ell \rangle$  for fixed boundary states.

**Proposition 2 (Branch structure)**  *$\mathcal{T}_\beta(\mathbf{S})$  is an ordered non-commutative trace polynomial in the dressed blocks  $\{\tilde{S}_{t,k}\}_{t,k=1}^n$ , with scalar trace factors generated by closed wiring loops in the path-cycle decomposition.*

The mechanism is a path-cycle decomposition. Under the fixed pairing  $(A_{t-1}, Q_t)$ , each permutation tensor  $\mathcal{P}_{\sigma^{-1}}$  induces a directed graph on  $\{A_0, \dots, A_n\}$  with edges  $A_{t-1} \rightarrow A_{\sigma(t)}$  carrying label  $\tilde{S}_{t, \sigma(t)}$ . Since  $\sigma$  is a permutation, the graph decomposes uniquely into one open path from  $A_0$  to  $A_n$  and disjoint directed cycles. The path yields an ordered product of dressed blocks (later steps multiplying on the left), each cycle yields a scalar trace-loop factor, and already the swap at  $n = 2$  produces an irreducible loop  $\text{Tr}(\tilde{S}_{2,1})$ . The dressed blocks are generically non-commuting matrices in  $\text{M}_d(\mathbb{C})$ , so  $\mathcal{T}_\beta(\mathbf{S})$  does not collapse to a determinant or Pfaffian.

### 3 Sequential Memory Barrier

The fusion geometry enforces a one-dimensional contraction order. If  $\mathcal{T}_\beta(\mathbf{S})$  admitted an efficient classical contraction in this order, its minimal fusion-order-respecting MPO would have small bond dimension. The opposite is true:

**Theorem 3 (Memory barrier)** *In the free trace-polynomial benchmark, where dressed blocks are treated as algebraically independent non-commuting generators and trace-loop factors as formal central indeterminates, the minimal fusion-order-respecting bond dimension obeys  $\chi_t = \Omega_d(\binom{n}{t})$ ; in particular  $\chi_{\lfloor n/2 \rfloor} = 2^{\Omega(n)}$ .*

The proof uses Nisan’s rank method [14] adapted to the trace-polynomial extension. Across cut  $t$ , any sequential contraction must distinguish  $\binom{n}{t}$  “used-leg” sectors corresponding to which subset of  $n$  input legs has been routed through the first  $t$  fusion steps; an explicit prefix/suffix matrix witness gives the diagonal rank lower bound, while the bond-product factorization gives  $\text{rank} \leq d^2 \chi_t$ .

Numerics in the physical monitored-FLO ensemble mirror this symbolic barrier (Fig. 2(a, b)): the




Figure 2: (a, b)  $\varepsilon$ -truncated minimal fusion-order MPO bond dimension  $\chi_{\max}^{(\varepsilon)}$  grows rapidly with  $n$  in the monitored-FLO ensemble (red), while a commuting control with  $S_{t,k} \propto \mathbf{1}_d$  and  $Z$ -only byproducts (blue) is strongly suppressed. (c, d) Normalized commutator diagnostic  $\nu_{\text{nc}}$  of the dressed blocks tracks its i.i.d. Ginibre [15] reference mean (dotted). 200 instances per  $n$ ; (a, c)  $d = 2$ , (b, d)  $d = 3$ .

truncated minimal bond dimension grows rapidly in  $n$ , while the commuting control with  $Z$ -only byproducts stays bounded by a small constant. A normalized commutator diagnostic of the dressed blocks (Fig. 2(c, d)) tracks the i.i.d. Ginibre [15] reference mean, confirming that Haar correlations in the dilute regime do not effectively commute.

A complementary signature comes from the conditional output distribution. The second moment  $\Gamma(\mathbf{c}) = \sum_{\beta} p(\beta|\mathbf{c})^2$  tracks the Haar-random reference  $2/(d^{2n} + 1)$  [16, 17], and the rescaled branch weights  $x := d^{2n}p(\beta|\mathbf{c})$  follow the Porter-Thomas law  $P(x) = e^{-x}$  [18] (Fig. 3), consistent with anti-concentration: a constant fraction of outcomes satisfy  $p(\beta|\mathbf{c}) = \Omega(d^{-2n})$  [19]. The commuting control, by contrast, departs from both signatures, confirming that the observed anti-concentration tracks the same non-commutative resource that drives the bond-dimension growth.

## 4 Hardness Conjecture and Outlook

Motivated by the non-commutative branch structure (Prop. 2), the exponential memory barrier (Thm. 3), and the numerical anti-concentration, we formulate an average-case hardness conjecture in the standard Stockmeyer framework.

**Conjecture 1 (Average-case hardness)** *For Haar-random number-conserving FLO instances at  $m = \kappa n^2$  conditioned on a collision-free record  $\mathbf{c}$ , estimating  $p(\beta|\mathbf{c})$  within relative error  $1/\text{poly}(n)$  is  $\#P$ -hard on average and for a constant fraction of outcomes  $\beta \in \mathbb{Z}_d^{2n}$ .*




Figure 3: Anti-concentration diagnostics. (a, b) Haar-normalized second moment  $(d^{2n} + 1)\Gamma/2$  versus  $n$ ; the monitored-FLO ensemble (red) tracks the Haar-random benchmark, while the commuting control (blue) deviates systematically. (c, d) Rescaled branch probabilities  $x := d^{2n}p(\beta|\mathbf{c})$  for representative  $n = 8$  instances against the Porter-Thomas law (grey dashed). (a, c)  $d = 2$ , (b, d)  $d = 3$ .

**Theorem 4 (Sampling hardness)** *Assuming Conjecture 1 and anti-concentration, no classical probabilistic polynomial-time algorithm can approximately sample from  $p(\beta|\mathbf{c})$  within total variation  $1/\text{poly}(n)$  on a non-negligible fraction of monitored-FLO instances, unless the polynomial hierarchy collapses.*

The proof is the standard Stockmeyer chain: anti-concentration converts total-variation closeness into multiplicative-error sampling estimates on a constant fraction of outcomes, and approximate counting then yields a  $\text{BPP}^{\text{NP}}$  procedure approximating  $p(\beta|\mathbf{c})$  to  $1/\text{poly}$  relative error, contradicting Conjecture 1.

This route is complementary to fermionic-advantage proposals via state injection [12]: there the non-Gaussian resource is supplied by magic inputs, whereas here the non-commutative branch structure is induced by monitoring, feedforward, and fixed-order fusion of an otherwise Gaussian FLO evolution. On locally connected architectures the feedforward routing compiles into a swap network; platforms with long-range gates [20] or dynamic qubit rearrangement [21] implement the required connectivity with minimal overhead. More broadly, this work identifies mid-circuit measurement and feedforward as a mechanism that can fundamentally alter the algebraic structure of free-fermion interference and induce computational power into an otherwise classically simulable quantum circuit.

## References

- [1] A. W. Harrow and A. Montanaro. Quantum computational supremacy. *Nature*, 549:203–209, 2017.
- [2] D. Hangleiter and J. Eisert. Computational advantage of quantum random sampling. *Rev. Mod. Phys.*, 95:035001, 2023.
- [3] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9:143–252, 2013.
- [4] C. S. Hamilton et al. Gaussian boson sampling. *Phys. Rev. Lett.*, 119:170501, 2017.
- [5] M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. R. Soc. A*, 467:459–472, 2011.
- [6] E. Knill. Fermionic linear optics and matchgates. arXiv:quant-ph/0108033, 2001.
- [7] B. M. Terhal and D. P. DiVincenzo. Classical simulation of noninteracting-fermion quantum circuits. *Phys. Rev. A*, 65:032325, 2002.
- [8] L. G. Valiant. Quantum circuits that can be simulated classically in polynomial time. *SIAM J. Comput.*, 31:1229–1254, 2002.
- [9] R. Jozsa and A. Miyake. Matchgates and classical simulation of quantum circuits. *Proc. R. Soc. A*, 464:3089–3106, 2008.
- [10] S. Bravyi. Lagrangian representation for fermionic linear optics. *Quantum Inf. Comput.*, 5:216–238, 2005.
- [11] M. Hebenstreit et al. All pure fermionic non-Gaussian states are magic states for matchgate computations. *Phys. Rev. Lett.*, 123:080503, 2019.
- [12] M. Oszmaniec, N. Dangniam, M. E. S. Morales, and Z. Zimborás. Fermion sampling: a robust quantum computational advantage scheme using fermionic linear optics and magic input states. *PRX Quantum*, 3:020328, 2022.
- [13] L. Stockmeyer. The complexity of approximate counting. *SIAM J. Comput.*, 14:849–861, 1985.
- [14] N. Nisan. Lower bounds for non-commutative computation. In *STOC '91*, pages 410–418, 1991.
- [15] J. Ginibre. Statistical ensembles of complex, quaternion, and real matrices. *J. Math. Phys.*, 6:440–449, 1965.
- [16] B. Collins and P. Śniady. Integration with respect to the Haar measure on unitary, orthogonal and symplectic groups. *Commun. Math. Phys.*, 264:773–795, 2006.
- [17] A. W. Harrow and R. A. Low. Random quantum circuits are approximate 2-designs. *Commun. Math. Phys.*, 291:257–302, 2009.
- [18] C. E. Porter and R. G. Thomas. Fluctuations of nuclear reaction widths. *Phys. Rev.*, 104:483–491, 1956.
- [19] A. M. Dalzell, N. Hunter-Jones, and F. G. S. L. Brandão. Random quantum circuits anticoncentrate in log depth. *PRX Quantum*, 3:010333, 2022.
- [20] L. Postler et al. Demonstration of fault-tolerant universal quantum gate operations. *Nature*, 605:675–680, 2022.
- [21] D. Bluvstein et al. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626:58–65, 2024.