Workspace-Bounded Quantum Pointer Chasing

Michael Aaron Cody *

Posted Date: 10 October 2025

doi: 10.20944/preprints202510.0382.v2

Keywords: quantum communication complexity; pointer chasing; workspace bounds; space-time tradeoffs; Kadison-Schwarz inequality; quantum lower bounds; multi-round protocols; quantum memory constraints; communication-memory tradeoff; quantum information theory

 doi:10.20944/preprints202510.0382.v2

Article

Workspace-Bounded Quantum Pointer Chasing

Michael Aaron Cody

Independent Theorist, USA; Mac92Contact@gmail.com

Abstract

Pointer chasing measures how information propagates through multiple rounds of communication. Quantum versions of the problem have never incorporated bounded local memory, even though every realistic protocol operates under finite workspace. This paper defines a workspace-bounded quantum pointer-chasing model, where each party has at most S qubits of reusable memory and total communication T. A multi-round form of the Kadison–Schwarz packing lemma shows that bounded workspace limits distinguishable state evolution across k rounds, giving

$$T\sqrt{S} \;\geq\; \Omega\big(k\sqrt{n}\big).$$

The bound recovers $T \geq \Omega(k\sqrt{n})$ when $S = 1$ and becomes trivial at $T \geq \Omega(k)$ when $S \geq n$. It extends the single-round framework established in *Workspace Bound* (Cody, 2025, DOI:10.5281/zenodo.17270825) and provides the first explicit multi-round tradeoff between communication and local memory in quantum protocols.

Keywords: quantum communication complexity; pointer chasing; workspace bounds; space-time tradeoffs; Kadison-Schwarz inequality; quantum lower bounds; multi-round protocols; quantum memory constraints; communication-memory tradeoff; quantum information theory

---

Introduction

Pointer chasing is one of the core problems that measure how information moves through rounds of communication. Each round depends on the message from the previous one, which makes the task sensitive to the number of interactions between parties. Classical lower bounds for multi-round pointer chasing remain active and sharp through recent work [12]. None of those results include a local workspace limit, and no quantum paper since Klauck's early work [9,10] has treated pointer chasing under explicit memory bounds. In realistic quantum communication, local memory cannot be ignored. Qubits that remain coherent during message exchange occupy space and control resources, and that physical limit changes the information flow of a protocol. A bounded workspace means fewer orthogonal states can be stored between rounds, which in turn restricts how much new information can be carried forward. This work asks how that memory limit changes the cost of pointer chasing.

The main result shows that total communication T and local workspace S satisfy the inequality

$$T\sqrt{S} \geq \Omega(k\sqrt{n}).$$

Here k is the number of rounds and n the pointer length. The bound recovers $T \geq \Omega(k\sqrt{n})$ when $S = 1$ and reduces to $T \geq \Omega(k)$ when $S \geq n$. It extends the single-round workspace relation proved in *Workspace Bound* [5] and moves the analysis to multi-round settings where both communication and local memory act as limited resources.

The bound connects directly to prior results. Classical k-round pointer chasing without workspace limits requires $\Omega(n \log^{k-1} n)$ communication [15], while quantum protocols without workspace constraints achieve $\Omega(k\sqrt{n})$ communication [9]. The workspace-bounded inequality derived here recovers the quantum bound when $S = 1$ (minimal workspace) and becomes trivial when $S \geq n$ (unbounded

 doi:10.20944/preprints202510.0382.v2

workspace), showing that the $\sqrt{S}$ factor quantifies how limited memory compresses distinguishable evolution across rounds. The tradeoff is multiplicative: reducing workspace from n to $\sqrt{n}$ increases the required communication by a factor of $\sqrt{n}$.

Section 2 defines the workspace-bounded pointer-chasing model. Section 3 proves the multi-round packing lemma using the Kadison–Schwarz inequality. Section 4 derives the tradeoff and checks endpoint cases. Section 5 discusses implications and open directions for interactive quantum protocols.

Related Work

Quantum communication complexity lower bounds began with Yao's early work and were formalized by Klauck [9,10], who proved that quantum protocols require $\Omega(\sqrt{n})$ communication for functions like disjointness and established general adversary frameworks. Those results assume unlimited local memory. Pointer chasing appeared in Klauck's framework but without workspace constraints. Classical pointer chasing has sharp multi-round bounds: Ponzio, Radhakrishnan, and Venkatesh [15] showed that k-round protocols need $\Omega(n/k)$ communication, and recent work by Mao, Yang, and Zhang [12] tightened these bounds using gadgetless lifting. None of those papers restrict how much memory each party can reuse between rounds.

Quantum lower-bound techniques rely on distinguishability arguments through the adversary method [1] and Gram matrix rank bounds [2]. The hybrid argument used here comes from Ambainis [1] and applies to multi-round settings where progress is measured step by step. Surveys by de Wolf [6] and Buhrman, Cleve, and Wigderson [4] cover standard quantum communication methods but do not address memory limits during interaction. Bounded quantum memory has been studied for single-party computation. Beame, Saks, Sun, and Vee [3] analyzed time-space tradeoffs for quantum search and showed that workspace constraints affect query complexity. Klauck, Špalek, and de Wolf [11] proved strong direct-product theorems with space bounds for quantum query models. These results apply to computation, not communication, and the techniques do not transfer directly because communication involves two parties with separate workspaces and message transmission between them.

The verify-and-reset structure used here follows the bounded-memory model from the workspace analysis in [5] and connects to reset operations in fault-tolerant quantum computation [13]. The Kadison–Schwarz inequality in operator form appears in Paulsen's operator algebra text [14] and has been used for quantum channel capacity bounds [8], but applying it to multi-round communication with workspace projectors is new.

This paper combines workspace constraints with multi-round pointer chasing for the first time. The result shows that the $\sqrt{S}$ factor changes communication scaling even when protocols can interact over k rounds.

Model and Definitions

Preliminaries. Let $\mathcal{H}$ denote a finite-dimensional Hilbert space. For density matrices ρ, σ on $\mathcal{H}$, the trace distance is $D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$. The Hilbert–Schmidt inner product is $\langle A, B\rangle_{\mathrm{HS}} = \mathrm{Tr}(A^\dagger B)$ with induced norm $\|A\|_2 = \sqrt{\mathrm{Tr}(A^\dagger A)}$. A quantum communication protocol uses local workspaces of dimension 2^S (S qubits per party) and message registers of C_r qubits per round. The total communication cost is $T = \sum_r C_r$. The pointer-chasing function $PC_k : [n]^k \times [n] \to [n]$ is $PC_k(f_1, \ldots, f_k; i_0) = f_k(f_{k-1}(\ldots f_1(i_0)\ldots))$. The goal is to obtain lower bounds on T as a function of workspace S. The pointer-chasing task involves two parties that exchange quantum messages through multiple rounds. Alice holds an array of functions $f_1, \ldots, f_k$ and Bob starts with an initial pointer i_0. Each round uses one of the functions to update the pointer value. The goal is for Bob to output the final pointer after k rounds with error at most $1/3$.

At round r, the message register has size at most C_r qubits. Each party maintains a private workspace of at most S qubits that may be reused but not expanded. The total Hilbert space for a

 doi:10.20944/preprints202510.0382.v2

sender in round r can be written as $\mathcal{H}_{\text{msg}}^{(C_r)} \otimes \mathcal{H}_{\text{work}}^{(S)}$. The workspace limit defines how much coherent information can be stored locally between rounds.

The protocol evolves by completely positive trace-preserving maps with an intermediate verification and reset (see [13]). After round r, the global state satisfies

$$\rho_{r+1} \;=\; V_r\, \mathcal{M}_r(\rho_r)\, V_r^\dagger,$$

where $\mathcal{M}_r$ represents the message transmission and V_r is a local unitary that resets a verified operator subspace to a clean state before the next round.

Success requires that Bob outputs the correct final pointer value with probability at least $2/3$. The total communication is the sum of all message sizes,

$$T \;=\; \sum_{r=1}^{k} C_r.$$

Distinguishability among message states is expressed through the Gram matrix

$$G_{xy} \;=\; \mathrm{Tr}(\rho_x^\dagger \rho_y) \;=\; \langle \rho_x, \rho_y \rangle_{\mathrm{HS}},$$

which uses the Hilbert–Schmidt inner product and satisfies $G \succeq 0$ (standard in quantum lower-bound methods [1,8]). For an S-qubit workspace, the underlying Hilbert space has dimension 2^S, so the Gram matrix obeys the conservative rank bound

$$\mathrm{rank}(G) \;\le\; 2^S.$$

The verify–and–reset structure also induces a Hilbert–Schmidt orthogonal projector Π onto a verified operator subspace. In later sections this projector provides a finite operator-norm budget through Kadison–Schwarz, recorded as $\mathrm{Tr}(\Pi) = S$ in the proofs (cf. [5]). That budget controls sums of squared Hilbert–Schmidt norms inside the verified subspace. The packing bounds in the next section use this operator-budget together with the Gram structure above to quantify how limited workspace restricts total information flow across rounds.

Multi-Round Workspace Packing Lemma

This section quantifies how bounded workspace limits total distinguishability across k rounds. The objects are defined explicitly, and every step is recorded.

Setup. For each round $r \in \{1, \dots, k\}$ and each marked input index i, let $\rho_i^{(r)}$ denote Bob's reduced state just before the r-th message is applied on the instance where the i-th location is the unique target. Let $\rho_0^{(r)}$ be the corresponding reference state for the all-zero instance. Define the round-r difference operators

$$F_i^{(r)} \;=\; \rho_i^{(r)} - \rho_0^{(r)}.$$

Fix adversary weights w_i with $\sum_i w_i^2 = 1$. These are the same weights used in standard hybrid and adversary arguments for quantum lower bounds [1].

No-accumulation bound (per round). Verify-and-reset ensures that each round begins from a fresh verified subspace. For density matrices ρ and σ,

$$(\rho - \sigma)^2 \;\preceq\; 2(\rho^2 + \sigma^2) \;\preceq\; 2(\rho + \sigma) \;\preceq\; 4I.$$

Hence, for fixed r,

$$\sum_i w_i^2\, F_i^{(r)\dagger} F_i^{(r)} \;\preceq\; 4I.$$

 doi:10.20944/preprints202510.0382.v2

This is the same no-accumulation principle used in the earlier workspace paper to control round-local operator growth [5].

Kadison–Schwarz inequality (Hilbert–Schmidt form). Let Π be the Hilbert–Schmidt orthogonal projector onto a verified operator subspace with $\mathrm{Tr}(\Pi) = S$. For any operator X,

$$\|\Pi X\|_2^2 \;\le\; \mathrm{Tr}\big(\Pi\, X^\dagger X\big).$$

This is the standard Kadison–Schwarz inequality in operator-space form [14].

Round projectors. Let $\Pi_1, \ldots, \Pi_k$ be the verified operator projectors for rounds 1 through k, each with $\mathrm{Tr}(\Pi_r) = S$. These encode the per-round workspace budget.

Packing calculation. Apply Kadison–Schwarz to the weighted differences inside each round and sum over r and i:

$$\sum_{r,i} \big\|\Pi_r\big(w_i F_i^{(r)}\big)\big\|_2^2 \;\le\; \sum_r \mathrm{Tr}\Big(\Pi_r \sum_i w_i^2 F_i^{(r)\dagger} F_i^{(r)}\Big) \;\preceq\; \sum_r \mathrm{Tr}(\Pi_r \cdot 4I) \;=\; 4kS.$$

Each inequality is explicit: Kadison–Schwarz on the left, the no-accumulation bound per round in the middle, and the trace identity $\mathrm{Tr}(\Pi_r I) = \mathrm{Tr}(\Pi_r) = S$ on the right.

Lemma (multi-round packing). In a k-round workspace-bounded protocol with per-round verified operator projectors Π_r satisfying $\mathrm{Tr}(\Pi_r) = S$,

$$\sum_{r,i} \big\|\Pi_r\big(w_i F_i^{(r)}\big)\big\|_2^2 \;\le\; 4kS.$$

Interpretation: the total Hilbert–Schmidt disturbance available to distinguish inputs across all rounds is bounded by a linear budget $O(kS)$. This is the only place where the k factor enters; everything else will be an averaging step over inputs and rounds.

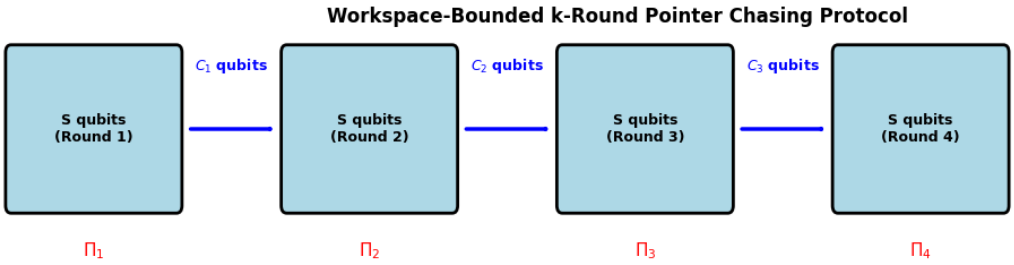

Figure 1. Workspace-bounded k-round pointer-chasing protocol. Each round r uses S qubits of local workspace and sends C_r qubits. The verify-and-reset operation applies projector Π_r with $\mathrm{Tr}(\Pi_r) = S$, enforcing the workspace bound before the next round.

Main Theorem

This section derives the communication–workspace tradeoff from the multi-round packing lemma using per-step indexing. Each transmitted qubit is treated as one step.

Adversary potential. For step index $\tau \in \{0, 1, \ldots, T\}$ define

$$\Phi_\tau \;=\; \frac{1}{n} \sum_{i=1}^{n} D\big(\rho_\tau^{(i)}, \rho_\tau^{(0)}\big),$$

with $D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ the trace distance. This measures the average distinguishability after τ message steps. Initially $\Phi_0 = 0$. Success with bounded error requires $\Phi_T \ge c$ for some constant $c > 0$. This follows from the Helstrom bound relating distinguishability to trace distance [7].

 doi:10.20944/preprints202510.0382.v2

Per-step progress bound. Let $r(\tau)$ be the round active at step τ, let $F_i^{(r)} = \rho_i^{(r)} - \rho_0^{(r)}$, fix weights w_i with $\sum_i w_i^2 = 1$, and let Π_r be the verified operator projector for round r (Section 3). The hybrid argument [1] implies

$$\Delta_\tau \;:=\; \Phi_{\tau+1} - \Phi_\tau \;\le\; \frac{1}{\sqrt{n}} \sqrt{\sum_i \big\|\Pi_{r(\tau)}\big(w_i F_i^{(r(\tau))}\big)\big\|_2^2}.$$

Packing budget across all steps. From the multi-round packing lemma,

$$\sum_{r=1}^{k} \sum_i \big\|\Pi_r\big(w_i F_i^{(r)}\big)\big\|_2^2 \;\le\; 4kS.$$

Each step τ belongs to some round $r(\tau)$, so reindexing gives

$$\sum_{\tau=1}^{T} \sum_i \big\|\Pi_{r(\tau)}\big(w_i F_i^{(r(\tau))}\big)\big\|_2^2 \;\le\; 4kS.$$

Accumulating progress (Cauchy–Schwarz). Summing Δ_τ and applying Cauchy–Schwarz,

$$\Phi_T \;=\; \sum_{\tau=0}^{T-1} \Delta_\tau \;\le\; \frac{1}{\sqrt{n}} \sum_{\tau=0}^{T-1} \sqrt{\sum_i \big\|\Pi_{r(\tau)}\big(w_i F_i^{(r(\tau))}\big)\big\|_2^2} \;\le\; \frac{1}{\sqrt{n}} \sqrt{T} \sqrt{\sum_{\tau=0}^{T-1} \sum_i \big\|\Pi_{r(\tau)}\big(w_i F_i^{(r(\tau))}\big)\big\|_2^2}.$$

By the packing budget, the last factor is at most $\sqrt{4kS}$, hence

$$\Phi_T \;\le\; O\left(\sqrt{\frac{TkS}{n}}\right).$$

Conclusion. Since $\Phi_T \ge c$, rearranging yields

$$T\sqrt{S} \;\ge\; \Omega\big(k\sqrt{n}\big).$$

This is the communication–workspace tradeoff for k-round pointer chasing.

Theorem (main result). In any k-round quantum pointer-chasing protocol with local workspace at most S qubits per party and total communication T qubits,

$$T\sqrt{S} \;\ge\; \Omega\big(k\sqrt{n}\big).$$

Endpoints: for $S = 1$, $T \ge \Omega(k\sqrt{n})$. For $S \ge n$, the bound reduces to $T \ge \Omega(k)$. This extends the single-round workspace relation of [5] to the multi-round setting.

Discussion

The bound connects directly to classical pointer chasing. Classical lower bounds such as Mao, Yang, and Zhang (ITCS 2025) remain active, showing that interaction depth controls information cost when local memory is unrestricted. Those results assume parties can reuse unlimited workspace between rounds. This paper introduces the first quantum version with an explicit workspace parameter since the early studies of Klauck (2000–2002), establishing how finite quantum memory alters the scaling of communication.

Workspace acts as an information resource. The Hilbert–Schmidt budget per round shows that limited workspace compresses distinguishable state evolution. The $\sqrt{S}$ factor proves the tradeoff between memory and communication is multiplicative, not additive. Each round gains only $\sqrt{n/S}$ in communication advantage rather than the full factor of n that appears when workspace is unbounded. Several questions remain. A matching upper bound is not known: it is open whether a protocol

 doi:10.20944/preprints202510.0382.v2

achieving $T = O(k\sqrt{n/S})$ exists. The role of entanglement assistance is also unclear; pre-shared EPR pairs may change the effective workspace scaling. Another direction is to remove the verify-and-reset assumption and test whether the Kadison–Schwarz argument survives in a fully general setting. Finally, a query-model version using multi-oracle access could determine whether the same Hilbert–Schmidt control extends beyond communication tasks.

The verify-and-reset rule simplifies the analysis but may reflect real hardware constraints. Systems with fixed qubit registers that must be re-initialized after each communication naturally follow this pattern. Testing these bounds experimentally on such architectures would show whether physical coherence limits match the theoretical scaling.

Table 1. Pointer-chasin g communication-complexity bounds.

Setting	Bound	Reference
Classical (unlimited workspace)	$\Omega(n \log^{k-1} n)$	[15]
Quantum (unlimited workspace)	$\Omega(k\sqrt{n})$	[9,10]
Quantum ($S = 1$)	$\Omega(k\sqrt{n})$	Theorem 1
Quantum ($S = \sqrt{n}$)	$\Omega(k\, n^{1/4})$	Theorem 1
Quantum ($S \geq n$)	$\Omega(k)$	Theorem 1

Conclusions

This work establishes $T\sqrt{S} \geq \Omega(k\sqrt{n})$ for quantum pointer chasing with S-qubit workspace, the first workspace–communication tradeoff for multi-round quantum protocols. The bound recovers $\Omega(k\sqrt{n})$ when $S = 1$ and reduces to $\Omega(k)$ when $S \geq n$, showing that the $\sqrt{S}$ factor quantifies how limited memory constrains communication. Several questions remain open. A matching upper bound achieving $T = O(k\sqrt{n/S})$ is unknown. The role of pre-shared entanglement in altering workspace scaling requires study. Removing the verify-and-reset assumption to test whether Kadison–Schwarz arguments survive in general settings is a natural next step. Extending these techniques beyond pointer chasing to other multi-round communication problems would determine the scope of workspace-packing methods.

The verify-and-reset structure may reflect real hardware limits. Systems with fixed qubit registers that must be re-initialized after each exchange naturally follow this model. Testing these bounds on such architectures would show whether physical coherence limits match theoretical predictions.

Author Contributions: This article is the sole work of the author.

Funding: No external funding was received for this work.

Data Availability Statement: No datasets were generated or analyzed in this study.

Declaration of Generative AI Use: Language and formatting assistance were provided by generative AI tools. All mathematical reasoning, results, and conclusions are the author's own.

Ethical Approval: Not applicable.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Ambainis, A. (2002). Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4), 750–767.
2. Barnum, H., Saks, M., & Szegedy, M. (2003). Quantum query complexity and semi-definite programming. In *Proceedings of the 18th IEEE Conference on Computational Complexity* (pp. 179–193).
3. Beame, P., Saks, M., Sun, X., & Vee, E. (2003). Time-space tradeoff lower bounds for randomized computation of decision problems. *Journal of the ACM*, 50(2), 154–195.
4. Buhrman, H., Cleve, R., & Wigderson, A. (1998). Quantum vs. classical communication and computation. In *Proceedings of the 30th ACM Symposium on Theory of Computing* (pp. 63–68).

 doi:10.20944/preprints202510.0382.v2

5. Cody, M. A. (2025). Workspace bound. Zenodo Preprint. https://doi.org/10.5281/zenodo.17270825
6. de Wolf, R. (2002). Quantum communication and complexity. *Theoretical Computer Science*, 287(1), 337–353.
7. Helstrom, C. W. (1976). *Quantum detection and estimation theory*. Academic Press.
8. Holevo, A. S. (1973). Bounds for the quantity of information transmitted by a quantum communication channel. *Problems of Information Transmission*, 9(3), 177–183.
9. Klauck, H. (2000). Quantum communication complexity. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science* (pp. 241–252).
10. Klauck, H. (2002). Lower bounds for quantum communication complexity. *SIAM Journal on Computing*, 30(6), 1863–1879.
11. Klauck, H., Špalek, R., & de Wolf, R. (2007). Quantum and classical strong direct product theorems and optimal time-space tradeoffs. *SIAM Journal on Computing*, 36(5), 1472–1493.
12. Mao, X., Yang, G., & Zhang, J. (2025). Gadgetless lifting beats round elimination: Improved lower bounds for pointer chasing. In *16th Innovations in Theoretical Computer Science Conference (ITCS)*, LIPIcs (Vol. 325, pp. 75:1–75:14).
13. Nielsen, M. A., & Chuang, I. L. (2000). *Quantum computation and quantum information*. Cambridge University Press.
14. Paulsen, V. I. (2002). *Completely bounded maps and operator algebras*. Cambridge University Press.
15. Ponzio, S. J., Radhakrishnan, J., & Venkatesh, S. (2001). The communication complexity of pointer chasing. *Journal of Computer and System Sciences*, 62(2), 323–355.