Space-bounded quantum interactive proof systems

François Le Gall [1] **Yupan Liu** [1] Harumichi Nishimura [1] Qisheng Wang [2,1]

[1]Nagoya University

[2]University of Edinburgh

Available on arXiv soon.

CS Theory Student Seminar, Columbia University, October 2024

1. Space-bounded quantum computation meets interactive proofs
2. Definitions of space-bounded quantum interactive proof systems
3. Main results on $\mathsf{QIP_UL}$, QIPL, and $\mathsf{QSZK_UL}$
4. Proof techniques: Upper bounds and properties for $\mathsf{QIP_UL}$ and QIPL
5. Open problems

What is **space-bounded** quantum computation?

Time-bounded quantum computation (BQP):

- Uses $\mathrm{poly}(n)$ elementary quantum gates, and thus requires $\mathrm{poly}(n)$ qubits.
- The goal is to find *a small corner* of a $2^{\mathrm{poly}(n)}$-dimension Hilbert space that holds the relevant information, which can only be extracted through measurements.

Space-bounded quantum computation (BQL) is introduced in [Watrous'98, Watrous'99]:

- Limits computation to $O(\log n)$ qubits, but allows $\mathrm{poly}(n)$ quantum gates.
- A quantum logspace computation operates on a $2^{O(\log n)}$-dimension Hilbert space, making this model appear weak and contained in NC.

However, BQL has shown *notable* power and gained recent increased attention:

- ◇ INVERTING WELL-CONDITIONED MATRICES [Ta-Shma'13, Fefferman-Lin'16] is BQL-complete, fully saturating the *quadratic* space advantage over classical suggested by BQL $\subseteq$ DSPACE$[\log^2(n)]$ [Watrous'99].
- ◇ Intermediate measurements appear to make BQL stronger than $\mathrm{BQ_UL}$, but provide *no advantage* for promise problems [Fefferman-Remscrim'21, Girish-Raz-Zhan'21].
- ◇ Quantum singular value transformation, a unifying quantum algorithm framework, has a logspace version [Gilyén-Su-Low-Weibe'18, Metger-Yuen'23, Le Gall-**L.**-Wang'23].

What is (quantum) **interactive proofs**?

Classical and quantum interactive proof systems

Given a promise problem $(\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})$, there is an interactive proof system $P \rightleftharpoons V$ that involves at most $\text{poly}(n)$ messages exchanged between the prover P and the verifier V:

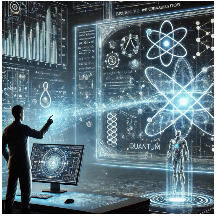

The image is generated using OpenAI's DALL·E model.

- ◇ P is typically all-powerful but untrusted;
- ◇ V is computationally bounded, possibly quantum;
- ◇ P and V may share entanglement in a quantum setting.

For any $x \in \mathcal{L}_{\text{yes}} \cup \mathcal{L}_{\text{no}}$, this proof system $P \rightleftharpoons V$ guarantees:

- ▶ For *yes* instances, $(P \rightleftharpoons V)(x)$ accepts w.p. at least $2/3$;
- ▶ For *no* instances, $(P \rightleftharpoons V)(x)$ accepts w.p. at most $1/3$.

Classical interactive proofs were introduced in [Babai'85, Goldwasser-Micali-Rackoff'85]:

1. Public-coin (AM$[k+2]$) matches the power of private-coin (IP$[k]$) [Goldwasser-Sipser'86].
2. IP $=$ PSPACE [Lund-Fortnow-Karloff-Nisan'90, Shamir'90], but IP$[O(1)] \subseteq$ IP$[2] \subseteq$ PH [B85, GS86].

Quantum interactive proofs were introduced in [Watrous'99, Kitaev-Watrous'00]:

1. "Parallelization": PSPACE $\subseteq$ QIP $\subseteq$ QIP$[3]$ [Watrous'99, Kitaev-Watrous'00].
2. QIP$[3] \subseteq$ QMAM $\subseteq$ PSPACE [Marriott-Watrous'04, Jain-Ji-Upadhyay-Watrous'09].

What is space-bounded (classical) interactive proofs?

Space-bounded classical interactive proofs were introduced in [Dwork-Stockmeyer'92, Condon'91], where the verifier operates in *logspace* but can run in *polynomial time*.

Public coins *weaken* the computational power of such proof systems:

- Classical interactive proofs with a logspace verifier using $O(\log n)$ private (random) coins ("IPL") exactly characterizes NP [Condon-Ladner'92].
  - Key ingredient: The fingerprinting lemma of multisets [Lipton'90].
- The model of *public-coin* space-bounded classical interactive proofs is weaker:
  - With $\mathrm{poly}(n)$ public coins, this model is contained in P [Condon'89].
  - With $O(\log n)$ public coins, it contains SAC[1] [Fortnow'89]; while with $\mathrm{poly}\log(n)$ public coins, it contains NC [Fortnow-Lund'91].
  - With $\mathrm{poly}(n)$ public coins, it contains P [Goldwasser-Kalai-Rothblum'15], connecting to *doubly-efficient interactive proofs*, where the prover is also efficient in some sense.

In this work, the verifier has *direct access* to messages during interaction, generalizing the space-bounded quantum Merlin-Arthur proofs (QMAL):

- **Direct access**: A QMAL verifier has *direct access* to an $O(\log n)$-qubit message, processing it directly in the verifier's workspace qubit, similar to QMA.
- QMAL $=$ BQL [Fefferman-Kobayashi-Lin-Morimae-Nishimura'16, Fefferman-Remscrim'21].

1st attempt: Space-bounded UNITARY quantum interactive proofs

Space-bounded *unitary* quantum interactive proofs ($\mathsf{QIP_U L}$)

Consider a $2l$-turn space-bounded unitary quantum interactive proof system $P \rightleftharpoons V$ for $(\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})$, where the verifier V operates in quantum logspace and has direct access to messages during interaction with the prover P:

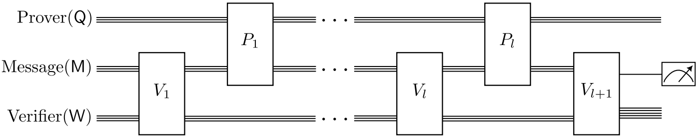

- The verifier V maps $x \in \mathcal{L}_{\text{yes}} \cup \mathcal{L}_{\text{no}}$ to $(V_1, \cdots, V_{l+1})$, where each V_j is unitary.
- Both M and W are of size $O(\log n)$, with M being accessible to both P and V.
- **Strong uniformity**: The description of $(V_1, \cdots, V_{l+1})$ can be computed by a single deterministic logspace Turing machine, intuitively implying $\{V_j\}$'s *repetitiveness*.

$\mathsf{QIP_U L}$ does not contain "IPL", particularly the model from [Condon-Ladner'92]:

- The prover P can somehow reveal private coins through shared entanglement, meaning soundness against classical messages does not extend to quantum.
- To show IP $\subseteq$ QIP, the verifier needs to *measure* the received messages at the beginning of each action, and treat the outcome as classical messages.

2^{nd} attempt: Space-bounded ISOMETRIC quantum interactive proofs

Space-bounded *isometric* quantum interactive proofs ($\mathsf{QIPL}^\diamond$)

Consider a $2l$-turn space-bounded isometric quantum interactive proof system $P \rightleftharpoons V$ for $(\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})$, where V acts on $O(\log n)$ qubits and has direct access to messages:

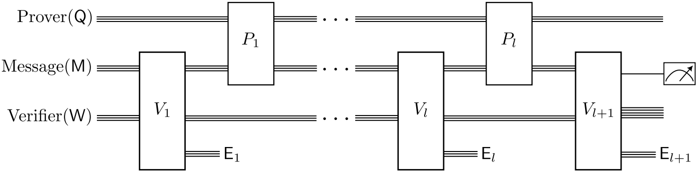

- Each V_j is an *isometric* quantum circuit, specifically allowing $O(\log n)$ ancillary gates that introduce an ancillary qubit $|0\rangle$ in the environment register E_j.
- Each environment register E_j is *only accessible* in the round of V_j belongs.
- The qubits in E_j cannot be altered after V_j, but entanglement with W can change!

$\mathsf{QIPL}^\diamond$ contains the Condon-Ladner model ("IPL"), but it appears too powerful:

- For instance, P can send an n-qubit state using $\lceil n/\log n \rceil$ messages of $(\log n)$-qubit states, while V takes only $O(\log n)$ qubits without P detecting the choices.
- $\mathsf{QIPL}^\diamond$ can verify the local Hamiltonian problem, and thus contains QMA:
  - A similar observation appeared in [Gharibian-Rudolph'22] on a *streaming* version of QMAL.

3rd attempt: Space-bounded quantum interactive proofs (QIPL)

Consider a $2l$-turn space-bounded quantum interactive proof system $P \rightleftharpoons V$ for $(\mathcal{L}_{\rm yes}, \mathcal{L}_{\rm no})$, where V acts on $O(\log n)$ qubits and has direct access to messages:

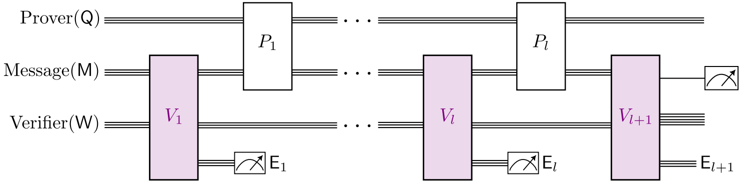

- Each V_j is an *almost-unitary* quantum circuit, meaning a unitary quantum circuit with $O(\log n)$ intermediate measurements in the computational basis.
  - QIPL also contains the Condon-Ladner model ("IPL").
- Applying the *principle of deferred measurements* to this almost-unitary quantum circuit V_j transforms it into **a special class of** isometric quantum circuits, followed by *measuring* the register E_j, with the outcome denoted by u_j.
- For *yes* instances, the distribution of intermediate measurement outcomes $u = (u_1, \cdots, u_l)$, condition on acceptance, must be *highly concentrated*.
  - This requirement leads to the NP containment for any QIPL proof system.
  - Specifically, let $\omega(V)|^u$ be the contribution of u to $\omega(V)$, where $\omega(V)$ is the maximum acceptance probability of $P \rightleftharpoons V$. There must exists a u^* such that $\omega(V)|^{u^*} \geq c(n)$.

Main results on $\mathsf{QIP_U L}$ and QIPL

Theorem 1. $\mathsf{QIPL} = \mathsf{NP}$.

- QIPL is the *weakest* model that includes space-bounded classical interactive proofs, ensuring that soundness against classical messages extends to quantum.
- The lower bound is inspired by space-bounded (private-coin) classical interactive proof systems for NP, particularly 3-SAT, in [Condon-Ladner'95].

Theorem 2. $\mathsf{SAC}^1 \cup \mathsf{BQP} \subseteq \mathsf{QIP_U L} \subseteq \cup_{c(n)-s(n)\geq 1/\mathrm{poly}(n)} \mathsf{QIPL}_{O(1)}[c,s] \subseteq \mathsf{P}$.

- Intermediate measurements enhance the model: $\mathsf{QIP_U L} \subsetneq \mathsf{QIPL}$ unless $\mathsf{P} = \mathsf{NP}$.
- $\mathsf{QIP_U L}$ proof systems, regarded as the most natural space-bounded analog to QIP, do not achieve the aforementioned soundness guarantee.
- The lower bound is inspired by space-bounded classical interactive proof systems with $O(\log n)$ public coins for evaluating SAC^1 circuits [Fortnow'89].
    - ◇ It is known that $\mathsf{NL} \subseteq \mathsf{SAC}^1 = \mathsf{LOGCFL} \subseteq \mathsf{AC}^1 \subseteq \mathsf{NC}^2$ [Venkateswaran'91].

Theorem 3. For any $c(n) - s(n) \geq \Omega(1)$, $\mathsf{QIPL}_{O(1)}[c,s] \subseteq \mathsf{NC}$.

- For constant-turn space-bounded quantum proofs, all three models are equivalent!
- An exponentially down-scaling version of $\mathsf{QIP} = \mathsf{PSPACE}$ [Jain-Ji-Upadhyay-Watrous'09].

Main results on space-bounded *unitary* quantum statistical zero-knowledge

Zero-knowledge property. A $\mathsf{QIP_UL}$ proof system has *the zero-knowledge property* if there is a *space-bounded* simulator that well approximates the snapshot states ("the verifier's view") in (M,W) after each turn, with respect to the trace distance.

- $\mathsf{QSZK_UL_{HV}}$ and $\mathsf{QSZK_UL}$ are space-bounded variants of quantum statistical zero-knowledge against an honest and arbitrary verifier, $\mathsf{QSZK_{HV}}$ and QSZK, respectively, introduced in [Watrous'02] and [Watrous'09].

Theorem 4. $\mathsf{QSZK_UL} = \mathsf{QSZK_UL_{HV}} = \mathsf{BQL}$.

The IndivProdQSD$[k,\alpha,\delta]$ **problem** (Individual Product State Distinguishability) involves two k-tuples of $O(\log n)$-qubit states, $\sigma_1,\cdots,\sigma_k$ and $\sigma'_1,\cdots,\sigma'_k$, whose purifications can be prepared by unitary quantum logspace circuits, satisfying $\alpha(n) - k(n)\cdot\delta(n) \geq 1/\mathrm{poly}(n)$ and $1 \leq k(n) \leq \mathrm{poly}(n)$, with the following conditions:

- For *yes* instances, the k-tuples are *globally far*, i.e., $\mathrm{T}(\sigma_1\otimes\cdots\otimes\sigma_k, \sigma'_1\otimes\cdots\otimes\sigma'_k) \geq \alpha$.
- For *no* instances, the k-tuples are *pairwise close*, i.e., $\forall j \in [k]$, $\mathrm{T}(\sigma_j,\sigma'_j) \leq \delta$.

$\mathsf{QSZK_UL_{HV}} \subseteq \mathsf{BQL}$ follows since IndivProdQSD is $\mathsf{QSZK_UL_{HV}}$-complete:

- The complement of IndivProdQSD is $\mathsf{QSZK_UL_{HV}}$-hard, similar to [Watrous'02].
- Since IndivProdQSD implies an "existential" version of GapQSD$_{\log}$, which is BQL-complete [Le Gall-**L.**-Wang'23], it follows that IndivProdQSD $\in$ QMAL $\subseteq$ BQL.

1. Space-bounded quantum computation meets interactive proofs
2. Definitions of space-bounded quantum interactive proof systems
3. Main results on $\mathsf{QIP_UL}$, QIPL, and $\mathsf{QSZK_UL}$
4. Proof techniques: Upper bounds and properties for $\mathsf{QIP_UL}$ and QIPL
5. Open problems

Proof techniques: Upper bounds for QIPL and $\mathsf{QIP_U L}$

Our approach is inspired by the SDP formulation for QIP [Vidick-Watrous'16].

$\mathsf{QIPL}_{\mathsf{O(1)}} = \mathsf{QIP_U L}_{\mathsf{O(1)}} \subseteq \mathsf{P}$. Consider a $2l$-turn $\mathsf{QIPL}_{\mathsf{O(1)}}$ proof system $P \rightleftharpoons V$, with $l \leq O(1)$. Let $\rho_{\mathtt{M}_j \mathtt{W}_j}$ and $\rho_{\mathtt{M}'_j \mathtt{W}_j}$, for $j \in [l]$, be snapshot states in the registers (M, W) after the verifier's and prover's action in the j-th round in $P \rightleftharpoons V$, respectively.

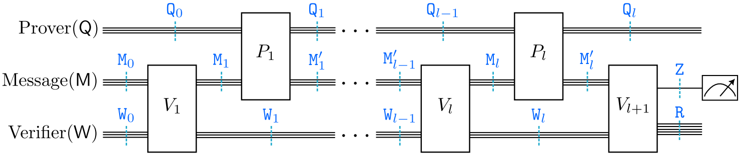

In this SDP formulation, we consider the following:

- ◇ Variables $\Leftrightarrow$ The snapshot states $\rho_{\mathtt{M}'_j \mathtt{W}_j}$ for $j \in [l]$ after each prover action;
- ◇ Objection function $\Leftrightarrow$ Maximum acceptance probability $\omega(V)$.

These variables are **independent** due to the *unitary* verifier. The SDP constraints are:

1. Verifier is always honest:
   $\rho_{\mathtt{M}_j \mathtt{W}_j} = V_j \rho_{\mathtt{M}'_{j-1} \mathtt{W}_{j-1}} V_j^\dagger$ for $j \in \{2, \cdots, l\}$, and $\rho_{\mathtt{M}_1 \mathtt{W}_1} = V_1 |\bar{0}\rangle\langle\bar{0}|_{\mathsf{MW}} V_1^\dagger$.
2. Prover's actions do not change the verifier's private register:
   $\mathrm{Tr}_{\mathtt{M}_j}(\rho_{\mathtt{M}_j \mathtt{W}_j}) = \mathrm{Tr}_{\mathtt{M}'_j}(\rho_{\mathtt{M}'_j \mathtt{W}_j})$ for $j \in [l]$.

As any SDP solution holds $O(\log n)$ qubits, standard SDP solvers ensure the efficiency.

Proof techniques: Upper bounds for QIPL and $\mathrm{QIP_U L}$ (Cont.)

QIPL $\subseteq$ NP. Now the verifier's actions are *almost-unitary* quantum circuits.

There is a family of SDP programs depending on the measurement outcome $\{u\}$:

- Variables $\Leftrightarrow$ The *unnormalized* snapshot states $\rho_{\mathtt{M}_j\mathtt{W}_j} \otimes |u_j\rangle\langle u_j|_{\mathsf{E}_j}$ for $j \in [l]$.
- Objection function $\Leftrightarrow$ $\omega(V)|^u$, namely the contribution of u to $\omega(V)$.

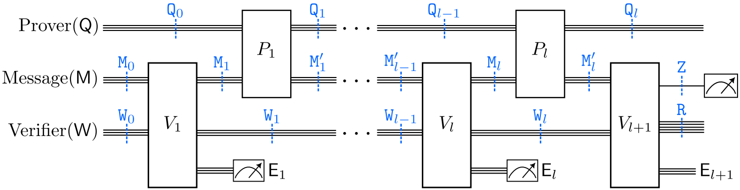

For a given $u = (u_1, \cdots, u_l)$, the SDP program includes two types of constraints:

1. Verifier is always honest: Let $\rho_{\mathtt{M}'_0\mathtt{W}_0} := |\bar{0}\rangle\langle\bar{0}|_{\mathsf{MW}}$, then
$\rho_{\mathtt{M}_j\mathtt{W}_j} \otimes |u_j\rangle\langle u_j|_{\mathsf{E}_j} = \left(I_{\mathtt{M}_j\mathtt{W}_j} \otimes |u_j\rangle\langle u_j|_{\mathsf{E}_j}\right) V_j \rho_{\mathtt{M}'_{j-1}\mathtt{W}_{j-1}} V_j^\dagger$ for $j \in [1]$.
2. Prover's actions do not change the verifier's private register:
$\mathrm{Tr}_{\mathtt{M}_j}\left(\rho_{\mathtt{M}_j\mathtt{W}_j} \otimes |u_j\rangle\langle u_j|_{\mathsf{E}_j}\right) = \mathrm{Tr}_{\mathtt{M}'_j}\left(\rho_{\mathtt{M}'_j\mathtt{W}_j} \otimes |u_j\rangle\langle u_j|_{\mathsf{E}_j}\right)$ for $j \in [l]$.

Next, we explain the NP containment:

- The classical witness w includes an l-tuple u and a feasible poly-size solution.
- The verification procedure involves checking whether (1) the solution encoded in w satisfies the SDP constraints based on u; and (2) $\omega(V)|^u \geq c(n)$.

Proof techniques: Basic properties for QIPL and $\mathsf{QIP_UL}$

Theorem 5 (Properties for QIPL and $\mathsf{QIP_UL}$). Let $c(n)$, $s(n)$, and $m(n)$ be functions such that $0 \leq s(n) < c(n) \leq 1$, $c(n) - s(n) \geq 1/\mathrm{poly}(n)$, and $1 \leq m(n) \leq \mathrm{poly}(n)$. Then, we have:

1. ***Closure under perfect completeness***.
$\mathsf{QIPL}_m[c,s] \subseteq \mathsf{QIPL}_{m+2}\left[1, 1 - \frac{1}{2}(c-s)^2\right]$ and $\mathsf{QIP_UL}_m[c,s] \subseteq \mathsf{QIP_UL}_{m+2}\left[1, 1 - \frac{1}{2}(c-s)^2\right]$.
2. ***Error reduction***. For any polynomial $k(n)$, there is a polynomial $m'(n)$ such that:
$$\mathsf{QIPL}_m[c,s] \subseteq \mathsf{QIPL}_{m'}\left[1, 2^{-k}\right] \text{ and } \mathsf{QIPL}_m[c,s] \subseteq \mathsf{QIPL}_{m'}\left[1, 2^{-k}\right].$$
3. ***Parallelization***. $\mathsf{QIP_UL}_{4m+1}[1,s] \subseteq \mathsf{QIP_UL}_{2m+1}[1, (1+\sqrt{s})/2]$.

Proof Sketch.

- Theorem 5 ❶ is directly adapted from [Vidick-Watrous'16].
- Theorem 5 ❷ uses *sequential repetition* due to the space constraint, with the key being to force the prover to "clean" the workspace.
- For establishing Theorem 5 ❸:
  - The original approach in [Kitaev-Watrous'00] fails, since it requires sending all snapshot states in a single message, which *exceeds* logarithmic size.
  - The turn-halving approach in [Kempe-Kobayashi-Matsumoto-Vidick'07] works, a "dequantized" version of the above approach, which leverages the *reversibility* of the verifier's actions.

Conclusions and open problems

Take-home messages on our work

1. Intermediate measurements play a *distinct* role in space-bounded quantum interactive proofs compared to space-bounded quantum computation: $\mathsf{QIP_UL} \subsetneq \mathsf{QIPL}$ unless $\mathsf{P} = \mathsf{NP}$ (this work), while $\mathsf{BQ_UL} = \mathsf{BQL}$ [FR21, GRZ21].
2. We define three models of space-bounded quantum interactive proofs:

	$\mathsf{QIP_UL}$	QIPL	$\mathsf{QIPL}^{\diamond}$
Verifier's actions	unitary	almost-unitary	isometry
Lower bounds	$\mathsf{SAC}^1 \cup \mathsf{BQL}$ "IPL" with $O(\log n)$ *public* coins	NP "IPL" with $O(\log n)$ *private* coins	QMA
Upper bounds	P	NP	PSPACE

3. Introducing the *zero-knowledge* property for $\mathsf{QIP_UL}$ proof systems, i.e., $\mathsf{QSZK_UL}$, eliminates the usual advantage gained from interaction ($\mathsf{QSZK_UL} = \mathsf{BQL}$).

Open problems

1. Is it possible to obtain a tighter characterization of $\mathsf{QIP_UL}$? For example, does $\mathsf{QIP_UL}$ contain "IPL" with $\omega(\log n)$ public coins?
2. What is the computational power of space-bounded quantum interactive proofs with a general quantum logspace verifier that allows "erasure" gates?

Thanks!