"High-dimensional quantum Schur transforms"

Adam Burchardt, Jiani Fei, Dmitry Grinko, Martin Larocca, Maris Ozols, Sydney Timmerman, Vladyslav Visnevskyi

AND

"Quantum Fourier transform for the symmetric group"

Carli Bruinsma, Dmitry Grinko, Maris Ozols

Background. The *quantum Schur transform* is an isometry U_{Sch} that implements the basis change in Schur–Weyl duality. It maps n-qudit basis states $|x_1,\dots,x_n\rangle \in (\mathbb{C}^d)^{\otimes n}$ to *Schur basis*:

$$|x_1,\dots,x_n\rangle \xmapsto{U_{\mathrm{Sch}}} \sum_{\lambda} \sum_{S\in \mathrm{SYT}(\lambda),\, M\in \mathrm{SSYT}(\lambda)} c^{x_1,\dots,x_n}_{\lambda,S,M}\ |\lambda\rangle_{\mathrm{irrep}} \otimes |S\rangle_{\mathrm{S}_n} \otimes |M\rangle_{\mathrm{U}_d} \tag{1}$$

where $\mathrm{SYT}(\lambda)$ and $\mathrm{SSYT}(\lambda)$ denote respectively the sets of *standard* and *semi-standard Young tableaus* of shape λ, $|\lambda\rangle_{\mathrm{irrep}}$ simultaneously labels irreducible representations (irreps) of the unitary group U_d and the symmetric group S_n, and $|S\rangle_{\mathrm{S}_n}$ and $|M\rangle_{\mathrm{U}_d}$ are the basis vectors[1] of the λ-irrep of S_n and U_d, respectively. Remarkably, Schur–Weyl duality guarantees that the symmetric group acts only on the symmetric group register $|S\rangle_{\mathrm{S}_n}$, while the unitary group acts only on the unitary group register $|M\rangle_{\mathrm{U}_d}$. This separation makes Schur transform a powerful tool for quantum algorithms and protocols that exploit permutation and unitary symmetries [14].

Two fundamentally different approaches to Schur transform are known. The original 2005 approach by Bacon, Chuang, and Harrow (BCH) [2, 14] relies on the representation theory of the unitary group and uses a cascade of Clebsch–Gordan transforms. In 2019, Krovi proposed a fundamentally different approach [26], rooted in the representation theory of the symmetric group, which is better suited for the *high-dimensional* ($d \gg n$) regime. Krovi's algorithm induces from a Young subgroup of the symmetric group S_n and uses the *quantum Fourier transform* (QFT) over S_n. The S_n-QFT maps any basis state $|g\rangle \in \mathbb{C}^G$ where $g \in \mathrm{S}_n$ to *Fourier basis*:

$$|g\rangle \xmapsto{\mathrm{QFT}} \sum_{\lambda} \sum_{S,\,T\in \mathrm{SYT}(\lambda)} \sqrt{\frac{d_\lambda}{n!}} R_\lambda(g^{-1})_{T,S}\ |\lambda\rangle_{\mathrm{irrep}} \otimes |S\rangle_{\mathrm{S}_n} \otimes |T\rangle_{\mathrm{S}_n} \tag{2}$$

where R_λ is the λ-irrep of S_n and d_λ is its dimension. Fourier basis has the property that left- and right-multiplication act only on $|S\rangle_{\mathrm{S}_n}$ and $|T\rangle_{\mathrm{S}_n}$, respectively.

Two decades after its invention, a substantial amount of literature on Schur transform and its applications exists. Unfortunately, this literature is not in an ideal state due to the highly technical nature of the subject matter. It is often difficult – even for domain experts – to assess the current state of the art and to agree on what has been established and what has not. In particular, some of the existing papers on Schur transform disagree on the very definition of the Schur transform, some omit important implementation details, and some contain significant errors, both in algorithms and their analysis. The goal of this joint submission is to rectify these issues and to provide a definitive treatment of the subject, particularly in the high-dimensional ($d \gg n$) case.

Summary of results. Our three main contributions are as follows:

(i) We identify and correct a crucial error in the final step of Krovi's Schur transform algorithm [26], provide a complete and detailed quantum circuit for the whole algorithm, and show its gate and depth complexity to be $\widetilde{O}(n^{7/2})$.

(ii) We correct errors in the algorithm and analysis of the S_n-QFT [23], a key ingredient of Krovi's Schur transform, and show its gate and depth complexity to be $\widetilde{O}(n^{7/2})$.

(iii) We provide details missing from Harrow's thesis [14] on how pre-processing can be used to adapt the BCH Schur transform to the high-dimensional ($d \gg n$) regime, and show that this updated BCH algorithm has gate and depth complexity $\widetilde{O}(\min(n^5, nd^4))$, improving upon its original $\widetilde{O}(nd^4)$ scaling.

Taken together, these results significantly advance the current state of the art of high-dimensional Schur transforms, making them practical in regimes where the local dimension d is much larger than the number of qudits n. This fills a key gap in the literature and strengthens the algorithmic foundations of a wide range of results in quantum information and quantum algorithms that rely on Schur–Weyl duality.

Revised version of Krovi's Schur transform

Problem. The Schur transform circuit proposed by Krovi [26] has two issues. The first is fundamental: in the final stage of the algorithm [26, Section 4, Theorem 3], the step preparing the correct state on the unitary register after the S_n-QFT is claimed to be classical (i.e., a unitary matrix with entries in $\{0,1\}$).[2] We show that this cannot be the case. Namely, the required operation has to be quantum (i.e., requires coherent superpositions) and cannot be implemented by classical gates alone. Second, we show that the Generalized Phase Estimation (GPE) in Krovi's algorithm is, in fact, unnecessary.

The main result of our first submission is to correct these two issues. We identify the correct (coherent) final transformation as the inverse of a certain isometry $\mathrm{V}_{\lambda,\mu}$. We describe this isometry in detail, analyze its properties, and provide a complete quantum circuit implementation. This constitutes a correct and fully consistent version of Krovi's

[1] These are so-called *Gelfand–Tsetlin bases*, meaning that the group action restricts nicely for a certain canonical tower of subgroups.

[2] This claim has to do with the fact that Krovi's definition of Schur transform assumes a measurement on the final unitary register.

high-dimensional quantum Schur transform (see Fig. 1):

Result 1 (Revised version of Krovi's Schur transform)**.** *The quantum circuit in Fig. 1 performs a (high-dimensional) quantum Schur transform and achieves Gelfand–Tsetlin bases for the symmetric and unitary registers. The algorithm has total gate and depth complexity* $\widetilde{O}(n^{7/2})$ *and total space complexity* $\widetilde{O}(n^2)$ *(the* $\widetilde{O}$ *notation hides a* $\mathrm{polylog}(n, d, 1/\epsilon)$ *factor), where* n *is the number of input qudits,* d *is their local dimension, and* ϵ *is the target precision.*

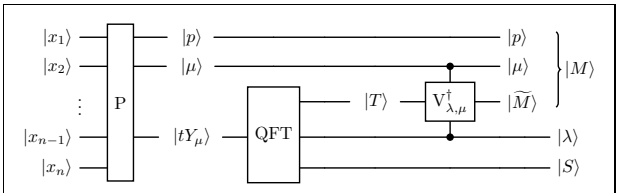

FIGURE 1. Our corrected version of Krovi's approach to the high-dimensional Schur transform.

For comparison, Krovi [26] reports the gate complexity of his algorithm only as $O(\mathrm{poly}(n, \log d, \log 1/\epsilon))$, without specifying an explicit exponent for n. Note also that his algorithm only achieves the *weaker* task of implementing a quantum channel that corresponds to measuring the irrep and unitary registers after Schur transform.

Our high-dimensional quantum Schur transform algorithm consists of three steps: (1) a pre-processing circuit P, (2) QFT over the regular representation of S_n, and (3) the inverse of a certain non-trivial isometry V. Our key innovation is the isometry V that corrects the final step of Krovi's algorithm. It maps *Gelfand–Tsetlin* (GT) *patterns* to superpositions of SYTs, whereas its inverse prepares a superposition of GT patterns from the SYT produced by the QFT. It is also responsible for obtaining a Gelfand–Tsetlin basis on the U_d register, while QFT obtains a Gelfand–Tsetlin basis on the S_n register. The isometry V has the following block structure:

$$V = \sum_{\mu} \sum_{\lambda} |\mu\rangle\langle\mu| \otimes V_{\lambda,\mu} \otimes |\lambda\rangle\langle\lambda| \tag{3}$$

where each block $V_{\lambda,\mu} : \mathbb{C}^{\mathrm{GT}(\lambda,\mu)} \to \mathbb{C}^{\mathrm{SYT}(\lambda)}$ is an isometry and $\mathrm{GT}(\lambda, \mu)$ denotes the set of GT patterns of shape λ and weight μ. Here μ is a composition that encodes the type of the input string $x = (x_1, x_2, \ldots, x_n) \in [d]^n$. To define it, let p_1 be the smallest entry of x, p_2 the second smallest, etc., and p_l ($l \leqslant n$) the largest. Then $p = (p_1, p_2, \ldots, p_l)$ captures the full alphabet of the entries of x. The i-th component of $\mu = (\mu_1, \mu_2, \ldots, \mu_l)$ is then equal to the number of times p_i appears in x.

For each irrep $\lambda \vdash n$ and composition $\mu = (\mu_1, \mu_2, \ldots, \mu_l)$, the rows and columns of the corresponding isometry $V_{\lambda,\mu}$ are labeled by $T \in \mathrm{SYT}(\lambda)$ and $\widetilde{M} \in \mathrm{GT}(\lambda, \mu)$, respectively, where the range of $\widetilde{M}$ depends on μ and hence the alphabet vector p used to compress an arbitrary $d \gg n$ to $d = n$. We compute the matrix entries of $V_{\lambda,\mu}$ by contracting a tensor network of Clebsch–Gordan tensors, and express them in terms of products of F-symbols. We show that after QFT the quantum state *always* lies within $\mathrm{im}(V)$ and therefore the circuit of the isometry V can be reversed. We also show that V can be implemented with $\widetilde{O}(n^3)$ gates, implying that the most expensive step is the S_n-QFT. Our next contribution is a correct implementation and complexity analysis of the S_n-QFT.

Quantum Fourier transform for the symmetric group

Krovi's algorithm as well as our corrected version of it rely on the symmetric group QFT. A high-level idea on how to implement the S_n-QFT was first provided by Beals [3] in 1997 and later generalized to other groups by Moore, Rockmore and Russell [30]. The first explicit circuit for S_n-QFT was derived only in 2013 by Kawano and Sekigawa [22], which they later claim to improve from $\widetilde{O}(n^4)$ to $\widetilde{O}(n^3)$ [23].

Problem. While Kawano and Sekigawa claim in [23] that their algorithm uses $\widetilde{O}(n^3)$ elementary gates for implementing the S_n-QFT, this is not accurate since in their analysis they assume black-box access to certain non-trivial operations that do not admit constant computational complexity in terms of standard elementary gates. Indeed, the main contribution of our second submission is to show that the true gate complexity of their algorithm is higher.

Result 2 (Revised S_n-QFT)**.** *For* $n \geqslant 1$*, the Fourier transform for the symmetric group* S_n *has a quantum circuit with* $\widetilde{O}(n^{7/2})$ *one- and two-qubit gates and* $\widetilde{O}(n^{3/2})$ *qubits. Both complexities scale as* $\mathrm{polylog}(1/\epsilon)$ *in the diamond norm error* ϵ*.*

Our analysis also uncovers a mistake in [23] that has to do with how the basis vectors of a certain Hilbert space are labeled. We address this as well as simplify their algorithm by removing one of the gates. In addition, we implement the algorithm in *Qrisp* and verify on a classical simulator that it works correctly for $n = 2$ and $n = 3$ (simulating it for $n > 4$ is not feasible on a laptop). Overall, our second submission is the first work that provides a full comprehensive description and analysis of all steps of the S_n–QFT algorithm.

High-dimensional BCH Schur transform

Problem. Our first submission also provides new insights into the original variant of the Schur transform introduced by Bacon, Chuang, and Harrow [2]. The original BCH algorithm has time complexity $\widetilde{O}(nd^4)$. In his PhD thesis [14], Harrow pointed out that in the $d \gg n$ regime the d-dependence can likely be replaced by an n-dependence, suggesting the possibility of achieving $\mathrm{polylog}(d)$ complexity.[3] However, the details of such a construction were never worked out. This brings us to our third result:

[3] See the footnote at the beginning of Section 8.1.2 in Ref. [14].

Result 3 (Description of the high-dimensional BCH circuit)**.** *The quantum circuit presented in Fig. 2 performs a high-dimensional version of the BCH quantum Schur transform. This algorithm has total gate complexity $\widetilde{O}(n^5)$, and total space complexity $\widetilde{O}(n^2)$ in terms of the number of qudits n in the input and their local dimension d.*

Our high-dimensional version of the BCH Schur transform consists of two stages: a preparation isometry $\hat{\mathrm{P}}$, and a sequence of n Clebsch–Gordan transforms on systems of local dimension $n \ll d$, see Fig. 2. The crucial difference from the standard BCH algorithm [2] is that we use the preparation isometry $\hat{\mathrm{P}}$ to reduce the local dimension of the main part of the circuit from an arbitrary d down to n. This enables the use of Clebsch–Gordan transforms of local dimension n instead of d, thereby improving the gate complexity from $\widetilde{O}(nd^4)$ to $\widetilde{O}(n^5)$, which gives better performance whenever $d \gg n$.

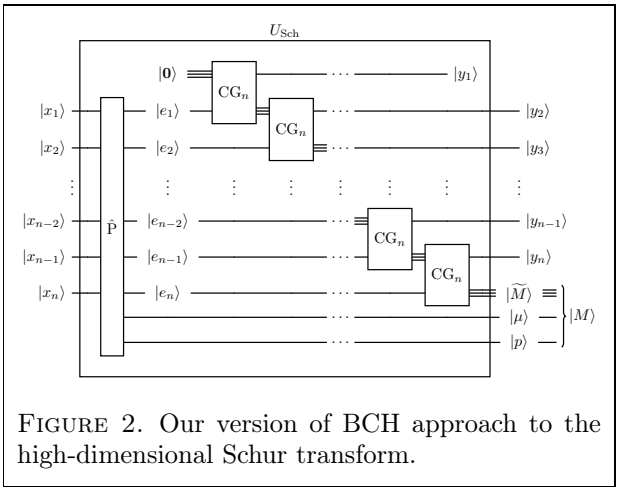

FIGURE 2. Our version of BCH approach to the high-dimensional Schur transform.

This version of the BCH Schur transform also implements a Gelfand–Tsetlin basis on both the symmetric and unitary group registers. But, unlike the standard encoding used in [2, 14], our high-dimensional version outputs the SYT on the symmetric group register in the so-called Yamanouchi encoding $|y\rangle$ instead of the tableau sequence $T = (T^0, \ldots, T^n \equiv \lambda) \in \mathrm{SYT}(\lambda)$. Here y_i specifies the row in which a box is added to obtain the Young diagram T^i from T^{i-1}, providing a more compact description of the standard Young tableau. This approach achieves space complexity $\widetilde{O}(n^2)$.

Applications

Since Schur transform can be performed efficiently on a quantum computer, it has become a foundational primitive in quantum algorithms. In particular, it is used for problems with permutational or unitarily invariant structure and equivariant quantum channels, with applications including quantum state tomography [13, 20, 32, 33], purification [6, 9, 25, 27], spectrum estimation [8, 24], universal compression [15, 19, 21, 35, 37], entanglement distillation [4, 28], and quantum hypothesis testing [1, 5, 16–18, 31] (see Section 1 of the first technical manuscript).

New applications specifically of the high-dimensional Schur transform include recent random purification protocols [10, 11, 36, 38], mixed state and channel tomography [29, 34], and the upcoming work on the efficient computation of plethysm coefficients [7].

In addition to the above, our high-dimensional Schur transform has other potential applications. First, our circuit enables efficient projection onto highly symmetric Schur–Weyl sectors, including the totally antisymmetric subspace of n qudits, yielding an $\widetilde{O}(n^{7/2})$ algorithm for first-quantized fermionic state preparation that is independent of the local dimension d. Second, by compressing large alphabets prior to measurement, our transform should remove the d^4 overhead in optimal port-based teleportation [12], reducing the runtime to $\widetilde{O}(n^{7/2})$, dominated by Schur transform.

Proof ideas and intuition

We give brief intuition for the two main isometries V and $\widehat{\mathrm{P}}$ mentioned above; full constructions appear in the first technical manuscript.

V isometry. Working in Krovi's framework, we use representations of S_n induced from trivial irreps of Young subgroups $Y_\mu = \mathrm{S}_{\mu_1} \times \cdots \times \mathrm{S}_{\mu_l}$. The inverse isometry $V^\dagger$ maps a standard Young tableau $T \in \mathrm{SYT}(\lambda)$ to a split-basis element described by a semistandard Young tableau $\widetilde{M} \in \mathrm{SSYT}(\lambda, \mu)$ of shape λ and weight μ.

Both SYT and split-basis vectors admit tensor-network descriptions where each tensor consists of Clebsch–Gordan coefficients; the matrix elements of V arise from contracting these networks, see Fig. 3. The contraction is implemented by a sequence of F-moves (changes of fusion order), so V decomposes into a sequence of local F-move gates, yielding an explicit circuit (see Sec. 4.3 of the first submission).

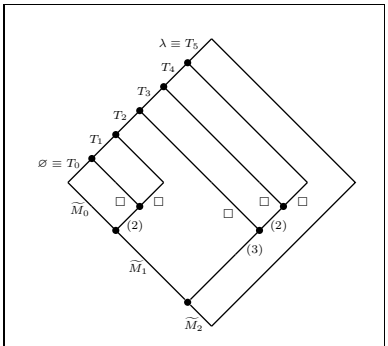

FIGURE 3. Contraction between a Young–Yamanouchi vector $T \in \mathrm{SYT}(\lambda)$ and a split basis vector $\widetilde{M} \in \mathrm{SSYT}(\lambda, \mu)$ (example: $n = 5$, $\mu = (2, 3)$).

$\widehat{\mathrm{P}}$ isometry. The preparation isometry $\widehat{\mathrm{P}}$ compresses the qudit alphabet $[d]$ to an effective alphabet of size $n \ll d$. Any string $x \in [d]^n$ can be encoded as $(p, \mu, \widetilde{x})$ with $\widetilde{x} \in [n]^n$, an ordered alphabet p, and a composition μ satisfying $x_i = p_{\widetilde{x}_i}$.

Quantumly we implement $\widehat{\mathrm{P}} : |x\rangle \mapsto |p\rangle|\mu\rangle|\widetilde{x}\rangle$ by a simple recursion that inserts symbols one-by-one. Concretely, assuming a map $\widehat{\mathrm{Prep}}' : |x'\rangle \mapsto |\widetilde{x}'\rangle|\mu'\rangle|p'\rangle$, we implement the next step as $\widehat{\mathrm{Prep}}_n : |\widetilde{x}'\rangle|\mu'\rangle|p'\rangle \otimes |x_n\rangle \mapsto |\widetilde{x}\rangle|\mu\rangle|p\rangle$, and set $\hat{\mathrm{P}} = \widehat{\mathrm{Prep}}_n\widehat{\mathrm{Prep}}' = \widehat{\mathrm{Prep}}_n\widehat{\mathrm{Prep}}_{n-1}\cdots\widehat{\mathrm{Prep}}_1$. The original Krovi map P is analogous (it outputs cosets $|tY_\mu\rangle$ rather than $|\widetilde{x}\rangle$) and shares most subroutines (see Secs. 6 and 8 of our first manuscript).

References

[1] K. M. R. Audenaert, J. Calsamiglia, R. Muñoz-Tapia, E. Bagan, Ll. Masanes, A. Acin, and F. Verstraete. "Discriminating States: The Quantum Chernoff Bound". In: *Physical Review Letters* 98.16 (Apr. 2007), p. 160501. DOI: 10.1103/physrevlett.98.160501.

[2] Dave Bacon, Isaac L. Chuang, and Aram W. Harrow. "Efficient Quantum Circuits for Schur and Clebsch-Gordan Transforms". In: *Physical Review Letters* 97.17 (Oct. 2006), p. 170502. DOI: 10.1103/physrevlett.97.170502. arXiv: quant-ph/0407082.

[3] Robert Beals. "Quantum computation of Fourier transforms over symmetric groups". In: *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*. STOC '97. El Paso, Texas, USA: Association for Computing Machinery, 1997, pp. 48–53. DOI: 10.1145/258533.258548. URL: https://doi.org/10.1145/258533.258548.

[4] Robin Blume-Kohout, Sarah Croke, and Daniel Gottesman. "Streaming Universal Distortion-Free Entanglement Concentration". In: *IEEE Transactions on Information Theory* 60.1 (Jan. 2014), pp. 334–350. DOI: 10.1109/tit.2013.2292135.

[5] Hao-Chung Cheng, Nilanjana Datta, Nana Liu, Theshani Nuradha, Robert Salzmann, and Mark M. Wilde. "An invitation to the sample complexity of quantum hypothesis testing". In: *npj Quantum Information, volume 11, Article number 94, June 2025* 11.1 (Mar. 26, 2024). DOI: 10.1038/s41534-025-00980-8. arXiv: 2403.17868 [quant-ph].

[6] Andrew M. Childs, Honghao Fu, Debbie Leung, Zhi Li, Maris Ozols, and Vedang Vyas. "Streaming quantum state purification". In: *Quantum* 9 (Jan. 2025), p. 1603. DOI: 10.22331/q-2025-01-21-1603.

[7] M. Christandl, G. Panova, P. Posta, and M. Walter. "Plethysm is in #BQP". In preparation. 2026.

[8] Matthias Christandl and Graeme Mitchison. "The Spectra of Density Operators and the Kronecker Coefficients of the Symmetric Group". In: *Commun. Math. Phys., Vol. 261, No. 3, pp. 789-797 (2006)* 261.3 (Sept. 2, 2004), pp. 789–797. DOI: 10.1007/s00220-005-1435-1. arXiv: quant-ph/0409016 [quant-ph].

[9] J. I. Cirac, A. K. Ekert, and C. Macchiavello. "Optimal Purification of Single Qubits". In: *Physical Review Letters* 82.21 (May 1999), pp. 4344–4347. DOI: 10.1103/physrevlett.82.4344.

[10] Filippo Girardi, Francesco Anna Mele, and Ludovico Lami. "Random purification channel made simple". In: (2025). arXiv: 2511.23451 [quant-ph]. URL: https://arxiv.org/abs/2511.23451.

[11] Filippo Girardi, Francesco Anna Mele, Haimeng Zhao, Marco Fanizza, and Ludovico Lami. *Random Stinespring superchannel: converting channel queries into dilation isometry queries*. 2025. arXiv: 2512.20599 [quant-ph]. URL: https://arxiv.org/abs/2512.20599.

[12] Dmitry Grinko, Adam Burchardt, and Maris Ozols. *Efficient quantum circuits for port-based teleportation*. 2023. arXiv: 2312.03188 [quant-ph].

[13] Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. "Sample-optimal tomography of quantum states". In: *IEEE Transactions on Information Theory* 63.9 (2017), pp. 5628–5641. DOI: 10.1109/tit.2017.2719044. arXiv: 1508.01797.

[14] Aram W. Harrow. "Applications of coherent classical communication and the Schur transform to quantum information theory". In: *Ph.D thesis, Massachusetts Institute of Technology, Cambridge, MA, 2005* (2005). arXiv: quant-ph/0512255 [quant-ph].

[15] M. Hayashi and K. Matsumoto. "Simple construction of quantum universal variable-length source coding". In: *IEEE International Symposium on Information Theory, 2003. Proceedings.* IEEE, 2003, p. 459. DOI: 10.1109/isit.2003.1228476.

[16] Masahito Hayashi. "Another quantum version of Sanov theorem". In: (July 2024). DOI: 10.48550/ARXIV.2407.18566. arXiv: 2407.18566 [quant-ph].

[17] Masahito Hayashi. "Asymptotics of Quantum Relative Entropy From Representation Theoretical Viewpoint". In: *J.Phys.A34:3413-3419,2001* 34.16 (Apr. 24, 1997), pp. 3413–3419. DOI: 10.1088/0305-4470/34/16/309. arXiv: quant-ph/9704040 [quant-ph].

[18] Masahito Hayashi. "Optimal sequence of quantum measurements in the sense of Stein s lemma in quantum hypothesis testing". In: *Journal of Physics A: Mathematical and General* 35.50 (Dec. 2002), pp. 10759–10773. DOI: 10.1088/0305-4470/35/50/307.

[19] Masahito Hayashi and Keiji Matsumoto. "Quantum universal variable-length source coding". In: *Physical Review A* 66.2 (Aug. 2002), p. 022311. DOI: 10.1103/physreva.66.022311.

[20] Yanglin Hu, Enrique Cervero-Martín, Elias Theil, Laura Mančinska, and Marco Tomamichel. *Sample Optimal and Memory Efficient Quantum State Tomography*. 2024. arXiv: 2410.16220 [quant-ph]. URL: https://arxiv.org/abs/2410.16220.

[21] Richard Jozsa, Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. "Universal Quantum Information Compression". In: *Physical Review Letters* 81.8 (Aug. 1998), pp. 1714–1717. DOI: 10.1103/physrevlett.81.1714.

[22] Yasuhito Kawano and Hiroshi Sekigawa. "Quantum Fourier transform over symmetric groups". In: *Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation*. ISSAC '13. Boston, Maine, USA: Association for Computing Machinery, 2013, pp. 227–234.

[23] Yasuhito Kawano and Hiroshi Sekigawa. "Quantum Fourier transform over symmetric groups—improved result". In: *Journal of Symbolic Computation* 75 (2016), pp. 219–243. DOI: 10.1016/j.jsc.2015.11.016.

[24] M. Keyl and R. F. Werner. "Estimating the spectrum of a density operator". In: *Physical Review A* 64.5 (Oct. 2001), p. 052311. DOI: 10.1103/physreva.64.052311. arXiv: quant-ph/0102027.

[25] M. Keyl and R.F. Werner. "The Rate of Optimal Purification Procedures". In: *Annales Henri Poincaré* 2.1 (Feb. 2001), pp. 1–26. DOI: 10.1007/pl00001027.

[26] Hari Krovi. "An efficient high dimensional quantum Schur transform". In: *Quantum* 3 (Feb. 2019), p. 122. DOI: 10.22331/q-2019-02-14-122.

[27] Zhaoyi Li, Honghao Fu, Takuya Isogawa, and Isaac Chuang. "Optimal Quantum Purity Amplification". In: (Sept. 2024). DOI: 10.48550/ARXIV.2409.18167. arXiv: 2409.18167 [quant-ph].

[28] Keiji Matsumoto and Masahito Hayashi. "Universal distortion-free entanglement concentration". In: *Physical Review A* 75.6 (June 2007), p. 062338. DOI: 10.1103/physreva.75.062338.

[29] Antonio Anna Mele and Lennart Bittel. *Optimal learning of quantum channels in diamond distance*. 2026. arXiv: 2512.10214 [quant-ph]. URL: https://arxiv.org/abs/2512.10214.

[30] Christopher Moore, Daniel Rockmore, and Alexander Russell. "Generic Quantum Fourier Transforms". In: *ACM Trans. Algorithms* 2.4 (2006), pp. 707–723.

[31] J Nötzel. "Hypothesis testing on invariant subspaces of the symmetric group: part I. Quantum Sanov's theorem and arbitrarily varying sources". In: *Journal of Physics A: Mathematical and Theoretical* 47.23 (May 2014), p. 235303. DOI: 10.1088/1751-8113/47/23/235303.

[32] Ryan O'Donnell and John Wright. "Efficient quantum tomography". In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. STOC'16. New York, NY, USA: Association for Computing Machinery, 2016, pp. 899–912. DOI: 10.1145/2897518.2897544. arXiv: 1508.01907.

[33] Ryan O'Donnell and John Wright. "Efficient quantum tomography II". In: *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*. STOC'17. New York, NY, USA: Association for Computing Machinery, 2017, pp. 962–974. DOI: 10.1145/3055399.3055454. arXiv: 1612.00034.

[34] Angelos Pelecanos, Jack Spilecki, Ewin Tang, and John Wright. *Mixed state tomography reduces to pure state tomography*. 2025. arXiv: 2511.15806 [quant-ph]. URL: https://arxiv.org/abs/2511.15806.

[35] Benjamin Schumacher. "Quantum coding". In: *Physical Review A* 51.4 (Apr. 1995), pp. 2738–2747. DOI: 10.1103/physreva.51.2738.

[36] Ewin Tang, John Wright, and Mark Zhandry. "Conjugate queries can help". In: (2025). arXiv: 2510.07622 [quant-ph]. URL: https://arxiv.org/abs/2510.07622.

[37] Yuxiang Yang, Giulio Chiribella, and Masahito Hayashi. "Optimal Compression for Identically Prepared Qubit States". In: *Physical Review Letters* 117.9 (Aug. 2016), p. 090502. DOI: 10.1103/physrevlett.117.090502.

[38] Satoshi Yoshida, Ryotaro Niwa, and Mio Murao. *Random dilation superchannel*. 2025. arXiv: 2512.21260 [quant-ph]. URL: https://arxiv.org/abs/2512.21260.