

Weighted Quantum Signal Processing: A Resource-Efficient Framework for Univariate Polynomial Approximation

Rohit Sarma Sarkar^{*} Rupayan Bhattacharjee^{*} Elias Combarro[†] Michele Grossi[‡]
Lirandë Pira[§] Carmen G. Almudéver[¶] Sergi Abadal^{*} and Eduard Alarcon^{*}

Abstract. Quantum Signal Processing (QSP) generates univariate polynomials but is limited by high circuit depth and parameter cost owing to its parity constraints. We introduce Weighted QSP (WQSP), where a weight function modifies the signal operator, recovering QSP as the unit-weight case. Non-trivial integer weights yield a pruned QSP with significant reduction of parameter count vis-a-vis circuit depth, achieving competitive polynomial approximation error with fewer resources. Generic weights admit error bounds and can yield exact approximation in many cases. We analyze WQSP from deterministic and quantum neural network viewpoints, showing it retains QSP-level expressivity with improved efficiency.

Keywords. Quantum Signal Processing, Weighted Quantum Signal Processing, Chebyshev Polynomials, Quantum circuit.

1 Introduction

Recent developments in theoretical quantum computing have revealed a unifying perspective in which many quantum algorithms can be interpreted as realizing spectral transformations of the form $f(H)$, where H is a Hamiltonian operator [10, 14]. Several frameworks have been developed for such transformations, including phase-estimation-based methods, linear combinations of unitaries (LCU) [2, 5], and Quantum Signal Processing (QSP)[10, 14]. QSP is particularly powerful because it realizes polynomial transformations of eigenvalues encoded in a unitary using compact single-qubit control sequences with minimal ancilla overhead. It was later generalized by Quantum Singular Value Transformation (QSVT) [10], where QSP appears as a special case. Nevertheless, QSP remains attractive due to its simplicity and its experimental demonstration on noisy trapped-ion devices [3, 12], with applications in Hamiltonian simulation and sampling [15, 6, 13, 14].

While effective, QSP has two main limitations. First, it imposes parity restrictions on realizable polynomial classes, requiring decomposition into even and odd sectors. Second, high-degree QSP requires many phase parameters growing linearly with the degree of the polynomial it approximates, which along with the parity constraint, results in two QSP-circuits with high depth and an increased overall parameter cost for approximating high-degree generic polynomial without a fixed parity. Although existence is guaranteed [14], constructing these phase parameters efficiently remains difficult, motivating numerical and iterative synthesis methods [10, 9, 4, 11, 16, 8, 7, 1].

^{*}Universitat Politècnica de Catalunya, Barcelona, Spain, E-mail: rohit.sarma.sarkar@upc.edu

[†]Universidad de Oviedo, Spain

[‡]Quantum Computing Infrastructures and Algorithms Competence Centre, CERN, Switzerland

[§]Centre for Quantum Technologies (CQT), National University of Singapore, Singapore

[¶]Universitat Politècnica de València, Spain

In this work, we introduce Weighted Quantum Signal Processing (WQSP), a structured extension of QSP in which fixed weights are incorporated into signal operators. Standard QSP is recovered in the trivial unit-weight case. For other integer weights, WQSP behaves as a pruned variant of QSP, exposing parameter redundancy and reducing circuit depth and parameter count while preserving approximation accuracy. For general weights, we derive approximation bounds quantifying compression–error trade-offs. Overall, WQSP provides a compact framework for approximation with QSP-level expressivity but reduced resource requirements.

2 Weighted Quantum Signal Processing

In this section, we provide a mathematical overview of WQSP. Let $\mathbf{w} = (w_1, \dots, w_k)^T \in \mathbb{R}^k$ and $x = \cos \theta$. We define WQSP as $U_{\Phi_{\mathbf{w}}}^{\text{WQSP}}(x) = \left(\prod_{j=1}^k R_z(\phi_j) \exp(iw_j \theta \sigma_x) \right) R_z(\phi_0)$. For $w_j = 1 \forall j$, this reduces to QSP. Restricting $\mathbf{w} \in \mathbb{N}^k$ (weights are greater than 1) yields a pruned version of QSP, where setting selected $\phi_j = 0$ merges adjacent signal operators into $U(x)^{w_j}$ and removes intermediate R_z rotations.

Similar to QSP, WQSP admits the block form $U_{\Phi_{\mathbf{w}}}^{\text{WQSP}}(x) = \begin{bmatrix} P(x) & iQ(x)\sqrt{1-x^2} \\ iQ(x)^* & P(x)^* \end{bmatrix}$, where $P(x) = \sum_{r \in \mathcal{R}_{\mathbf{w}}} c_r e^{i\phi_0} T_r(x)$, $Q(x) = \sum_{r \in \mathcal{R}_{\mathbf{w}}} d_r \sin \theta e^{i\phi_0} U_{r-1}(x)$, $T_r(x) = \cos(r\theta)$, $U_{r-1} = \frac{\sin(r\theta)}{\sin \theta}$, and $\mathcal{R}_{\mathbf{w}} = \{\sum_{p=1}^k |\eta_p w_p| \mid \eta_p \in \{\pm 1\}\}$, with c_r linear in $\Phi_{\mathbf{w}}$. For integer r , $P(x)$ reduces to a Chebyshev expansion, showing that suitable weights can realize polynomials beyond standard QSP structure.

We give deterministic and learning-based constructions. In the deterministic case, the error is bounded by $\|\mathbf{c}\|$ and vanishes as $\|\mathbf{c}\| \rightarrow 0$. In particular, the error becomes

$$\underbrace{D^2}_{\text{Error}} \leq 2\pi O(\|\mathbf{c}\|_2^2 + C(\rho)^2).$$

where, $C(\rho) := \sup_{|x| \leq \rho} |r(x)|$ where $r(x) := \arccos(x) - \frac{\pi}{2} + x$, and $\rho := \|W\mathbf{c}\|_{\infty}$ for $W = (\sqrt{2}H)^{\otimes k-1}P$ such that P is a permutation and H is the Hadamard gate.

For a learning problem, we solve the following:

$$\min_{\Phi} \|P(x) - \text{Re}(\langle 0 | U^{\text{WQSP}}_{\Phi}(x) | 0 \rangle)\|_{L_2} = \min_{\Phi} \sqrt{\int_{-1}^1 \left(P(x) - \text{Re}(\langle 0 | U^{\text{WQSP}}_{\Phi}(x) | 0 \rangle) \right)^2 dx}.$$

The learning approach optimizes $\Phi_{\mathbf{w}}$ for WQSP, revealing parameter redundancy in QSP. For degree d polynomials, standard QSP uses $\Theta(d)$ parameters and depth, while WQSP with $w_j = 2^{j-1}$ (or $w_j = j$) achieves comparable accuracy using $k = \Theta(\log d)$ (or $k = \Theta(\sqrt{d})$) parameters.

3 Results

We illustrate the key properties of the WQSP framework through polynomial approximation. We show that QSP can be pruned into shallower circuits, achieving linear to exponential reductions in parameters without significant performance loss.




Figure 2: WQSP fitting of a real 7th degree target polynomial (fixed odd parity) with different weight partitions whose coefficients are randomly generated. In 2a and 2b the weight partitions are respectively $(3, 2, 1, 1)$ and $(1, 1, 1, 1, 1, 1, 1)$ (standard-QSP) whose weight vector lengths are 5 and 8 respectively. We observe that the QSP case produces an error of a similar order compared to 2a which provides appreciable parameter reduction. In 2a, we see that the weight vector length is equal to the coefficient vector length of our degree 7 target polynomial (4 coefficients for fixed odd parity) and thus such a setup captures the degrees of freedom necessary to approximate the target polynomial.




Figure 1: WQSP fitting error across all partitions for a real degree-14 (1a) and degree-10 (1b) polynomial. Polynomial coefficients were generated by uniformly sampling from $[-1/8, 1/8]$ and $[-1/6, 1/6]$ respectively (for polynomial normalization), and approximated using WQSP. Phase parameters were optimized using the Levenberg–Marquardt method with 20 random restarts, and the solution achieving the lowest approximation error was reported.

Conclusion. We introduce WQSP, a weighted extension of QSP that reduces circuit depth and parameter count from linear to exponential scales, and enables efficient polynomial approximation through N -weighted constructions with different weight partitions of the target polynomial degree. It achieves QSP-level expressivity with fewer resources and admits high-quality learning-based implementations for univariate polynomial approximation.

References

- [1] Michel Alexis, Lin Lin, Gevorg Mnatsakanyan, Christoph Thiele, and Jiasu Wang. Infinite quantum signal processing for arbitrary Szegő functions. *Communications on Pure and Applied Mathematics*, 79(1):123–174, 2026.
- [2] Dominic W. Berry, Andrew M. Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. In *Proceedings of the 56th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 792–809. IEEE, 2015.
- [3] J.-T. Bu, Lei Zhang, Zhan Yu, Jing-Bo Wang, W.-Q. Ding, W.-F. Yuan, B. Wang, H.-J. Du, W.-J. Chen, L. Chen, J.-W. Zhang, J.-C. Li, F. Zhou, Xin Wang, and M. Feng. Exploring the experimental limit of deep quantum signal processing using a trapped-ion simulator. *Phys. Rev. Appl.*, 23:034073, Mar 2025.
- [4] Rui Chao, Dawei Ding, András Gilyén, Cupjin Huang, and Mario Szegedy. Finding Angles for Quantum Signal Processing with Machine Precision. *arXiv preprint arXiv:2003.02831*, 2020.
- [5] Andrew M. Childs, Robin Kothari, and Rolando D. Somma. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6):1920–1950, 2017.
- [6] Alexander M. Dalzell, Sam McArdle, Mario Andrea Berta, Przemyslaw Bienias, Chi-Fang Chen, András Gilyén, Connor T. Hann, Michael J. Kastoryano, Emil T. Khabiboulline, Aleksander Kubica, Grant Salton, Samson Wang, and Fernando G. S. L. Brandão. *Quantum Algorithms: A Survey of Applications and End-to-End Complexities*. Cambridge University Press, Cambridge, 2025.
- [7] Yulong Dong, Lin Lin, Hongkang Ni, and Jiasu Wang. Infinite quantum signal processing. *Quantum*, 8:1558, 2024.
- [8] Yulong Dong, Lin Lin, Hongkang Ni, and Jiasu Wang. Robust iterative method for symmetric quantum signal processing in all parameter regimes. *SIAM Journal on Scientific Computing*, 46(5):A2951–A2971, 2024.
- [9] Yulong Dong, Xiang Meng, K. Birgitta Whaley, and Lin Lin. Efficient phase-factor evaluation in quantum signal processing. *Phys. Rev. A*, 103:042419, Apr 2021.
- [10] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, page 193–204, New York, NY, USA, 2019. Association for Computing Machinery.
- [11] Jeongwan Haah. Product Decomposition of Periodic Functions in Quantum Signal Processing. *Quantum*, 3:190, October 2019.
- [12] Yuta Kikuchi, Conor Mc Keever, Luuk Coopmans, Michael Lubasch, and Marcello Benedetti. Realization of quantum signal processing on a noisy quantum computer. *npj Quantum Information*, 9(1):93, Sep 2023.
- [13] Lorenzo Laneve. Quantum signal processing, phase extraction, and proportional sampling. *arXiv preprint arXiv:2303.11077*, 2024.

- [14] Guang Hao Low and Isaac L. Chuang. Optimal hamiltonian simulation by quantum signal processing. *Phys. Rev. Lett.*, 118:010501, Jan 2017.
- [15] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. Grand unification of quantum algorithms. *PRX Quantum*, 2:040203, Dec 2021.
- [16] Lexing Ying. Stable factorization for phase factors of quantum signal processing. *Quantum*, 6:842, October 2022.