Preprint:

Nqhoaki, Khotsofalang and Wilding, Terry (Accepted: 2026) Quantum-Inspired Nonlinear Model Predictive Control via Quantum Singular Value Transformation for Autonomous Systems. [Preprint]

Quantum-Inspired Nonlinear Model Predictive Control via Quantum Singular Value Transformation for Autonomous Systems

Khotsofalang Nqhoaki
Dept. of Computer Science
University of York
York, UK
Email: kn877@york.ac.uk

Terry Wilding
Dept. of Computer Science
University of York
York, UK
Email: terry.wilding@york.ac.uk

***Abstract*—The deployment of Nonlinear Model Predictive Control (NMPC) in safety-critical autonomous systems is frequently constrained by the computational intractability of solving non-convex optimization problems within strict real-time deadlines. While quantum computing offers theoretical speedups, the lack of fault-tolerant hardware necessitates the development of quantum-inspired classical algorithms. This paper proposes a novel NMPC framework based on the Quantum Singular Value Transformation (QSVT), which leverages polynomial approximation and matrix sketching to accelerate the solution of receding-horizon control problems. We derive a comprehensive theoretical complexity analysis, demonstrating a reduction from the cubic complexity of classical Sequential Quadratic Programming (SQP) to a near-linear complexity dependent on sketch dimensions. The proposed QSVT-NMPC is benchmarked against classical NMPC, Sliding Mode Control (SMC), and a suite of quantum-inspired variants, including Simulated Quantum Annealing and Variational Quantum Circuits using an inverted pendulum system. Experimental results validate that the optimized QSVT-NMPC achieves a median settling time of 3.0 seconds, outperforming other quantum-inspired methods and offering a superior trade-off between control stability and computational latency compared to classical approaches. The algorithm demonstrates significant robustness to parameter variations and reduced memory footprint, establishing it as a viable candidate for embedded implementation in next-generation autonomous systems.**

***Index Terms*—Model Predictive Control, Quantum-Inspired Computing, Quantum Singular Value Transformation (QSVT), Autonomous Systems, Real-Time Optimization, Embedded Control.**

I. INTRODUCTION

The landscape of autonomous systems encompassing self-driving vehicles, aerial drones, and robotic surgery demands control frameworks capable of handling multivariable dynamics, state constraints, and environmental uncertainties in real-time [1]. Model Predictive Control (MPC) has emerged as the preeminent strategy for such tasks, leveraging a receding horizon approach to optimize control actions over a future time window [2]. Despite its theoretical robustness, the industrial adoption of Nonlinear MPC (NMPC) in resource-constrained environments remains limited. The primary bottleneck is the computational burden associated with solving a constrained Finite-Time Optimal Control (CFTOC) problem at every sampling instant, often exceeding the latency budgets of embedded hardware [3].

A. Motivation and Problem Context

In highly nonlinear or high-dimensional systems, the optimization landscape is often non-convex, populated with local minima that trap traditional gradient-based solvers [4]. Furthermore, the computational complexity of standard NMPC implementations, typically relying on interior-point methods or Sequential Quadratic Programming (SQP), scales cubically with the state dimension. This scaling renders NMPC impractical for fast-sampling systems or low-power microcontrollers, creating a gap between the theoretical capabilities of predictive control and the practical realities of embedded deployment.

Recent advances in quantum computing have proposed novel pathways for accelerating optimization tasks. Algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Eigensolver (VQE) promise to traverse complex solution spaces more efficiently by exploiting quantum mechanical phenomena like superposition and entanglement [5]. However, the current Noisy Intermediate-Scale Quantum (NISQ) era hardware lacks the coherence times and qubit counts necessary for real-time control loops [6].

To bridge this gap, researchers have turned to *quantum-inspired* algorithms, classical methods that mimic quantum behaviors. Simulated Quantum Annealing (SQA) [7] and Quantum-Inspired Evolutionary Algorithms (QIEA) [8] have shown promise in avoiding local minima. However, their stochastic nature often introduces high variance in execution time, which is undesirable for hard real-time systems. A more deterministic approach is found in the Quantum Singular Value Transformation (QSVT) [9], a framework that unifies quantum algorithms and allows for the efficient approximation of matrix functions.

B. Contributions

This paper addresses the latency-accuracy trade-off in NMPC by proposing and validating a suite of quantum-inspired controllers, with a specific focus on a novel QSVT-based approach. The primary contributions are:

1) **Algorithmic Design:** We develop the Optimized QSVT-NMPC algorithm, which utilizes Chebyshev polynomial surrogates and random sketching to approximate the Newton-step in the control optimization loop, reducing complexity from $O(Nn^3)$ to near-linear time.
2) **Theoretical Analysis:** We provide a rigorous complexity analysis comparing classical NMPC, Quantum-Inspired Annealing (QIA), and QSVT methods, proving the asymptotic advantages of the proposed approach.
3) **Empirical Validation:** We implement a unified simulation framework to benchmark the controllers on the inverted pendulum swing-up task. We demonstrate that Optimized QSVT-NMPC achieves a settling time of 3.0 seconds—a 65% improvement over its unoptimized form while maintaining strict computational bounds.
4) **Comparative Study:** We perform a statistical evaluation using ANOVA and Tukey HSD tests to validate that the proposed method offers superior stability compared to classical surrogates like Random Fourier Features (RFF) and VQC implementations.

The remainder of this paper is organized as follows: Section II reviews related work. Section III formulates the system model and classical NMPC. Section IV details the theoretical development of the quantum-inspired algorithms, with a focus on QSVT. Section V describes the experimental design. Section VI presents the results and analysis, and Section VII concludes the paper.

II. LITERATURE REVIEW

A. Nonlinear MPC and Computational Challenges

NMPC extends the linear MPC paradigm to nonlinear systems $x_{k+1} = f(x_k, u_k)$. Standard approaches involve linearizing the system dynamics at each time step and solving a Quadratic Program (QP). While methods like the Real-Time Iteration (RTI) scheme [3] improve speed, they rely heavily on good initial guesses and convexity assumptions. In non-convex landscapes, these methods often fail to find the global optimum, leading to instability. Classical meta-heuristics like Simulated Annealing (SA) have been applied to MPC [13], but their stochastic nature and lack of theoretical convergence bounds limit their use in safety-critical applications.

B. Quantum-Inspired Optimization

Quantum-inspired optimization translates quantum principles to classical hardware.

- *Superposition:* Algorithms like QIEA represent solutions as probability vectors (qubits), allowing parallel exploration of the search space [8].
- *Tunneling:* Simulated Quantum Annealing (SQA) mimics quantum tunneling, enabling the optimizer to escape local minima more efficiently than thermal fluctuations in classical SA [12].
- *Entanglement:* Variational Quantum Circuits (VQC) utilize entanglement to represent complex correlations between control variables [11].

While these methods improve solution quality, they often introduce significant computational overhead due to ensemble sampling or the simulation of quantum gates.

C. Quantum Singular Value Transformation (QSVT)

QSVT represents a paradigm shift in quantum algorithm design, providing a unified framework for implementing matrix functions $f(A)$ via polynomial transformations of singular values [9]. Classical dequantization techniques, such as those proposed by Gilyen et al., suggest that the core benefits of QSVT specifically the ability to approximate inverses and filtering operations can be realized classically using sketching methods [10]. This paper leverages these classical analogs to accelerate the linear algebra underpinning the NMPC optimization step.

III. THEORETICAL DEVELOPMENT

A. System Model and Classical NMPC Formulation

Consider a discrete-time nonlinear time-invariant system:

$$x_{k+1} = f(x_k, u_k), \quad x_k \in \mathbb{R}^n, \quad u_k \in \mathbb{R}^m \tag{1}$$

subject to constraints $x_k \in \mathcal{X}, u_k \in \mathcal{U}$. The NMPC problem at time t minimizes the cost function:

$$J(x_t, U_t) = \sum_{i=0}^{N-1} \ell(x_{t+i}, u_{t+i}) + V_f(x_{t+N}) \tag{2}$$

where N is the prediction horizon, $\ell(\cdot)$ is the stage cost, $V_f(\cdot)$ is the terminal cost, and $U_t = [u_t, \dots, u_{t+N-1}]$ is the control sequence.

The Sequential Quadratic Programming (SQP) approach linearizes the dynamics and approximates the cost function quadratically. At each SQP iteration, we solve the QP subproblem:

$$\min_{\Delta u} \frac{1}{2} \Delta u^T H_k \Delta u + g_k^T \Delta u \tag{3}$$

where H_k is the Hessian approximated via Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm and g_k is the gradient of the Lagrangian. The critical computational bottleneck is solving the linear system:

$$H_k \Delta u = -g_k \tag{4}$$

Direct inversion requires $O((Nm)^3)$ operations. The goal of quantum-inspired acceleration is to approximate the solution to (4) with lower complexity.

B. Quantum-Inspired Annealing MPC (QIA-NMPC)

We first define a probabilistic MPC solver based on annealing. We modify the cost function $J(u)$ by embedding a quantum-inspired tunneling term $Q(u)$:

$$\tilde{J}(u) = J(u) + \Gamma(t) Q(u) \tag{5}$$

where $\Gamma(t)$ is a time-dependent transverse field coefficient. The term $Q(u)$ represents the energy associated with control differences, encouraging exploration. The control sequence update follows a Metropolis-Hastings criterion:

$$u_{k+1} = \begin{cases} u_{new} & \text{if } \exp(-\Delta \tilde{J}/T_k) > \text{rand}() \\ u_k & \text{otherwise} \end{cases} \tag{6}$$

where T_k is the temperature schedule. The complexity of QIA-NMPC is dominated by the cost of evaluating the dynamics over the horizon N for a number of replicas R, leading to a complexity of $O(R \cdot N \cdot n^2)$.

C. Proposed Method: Quantum Singular Value Transformation MPC (QSVT-NMPC)

To achieve deterministic speedups, we propose approximating the solution to (4) using the principles of QSVT.

1) Matrix Sketching: Instead of using the full Hessian H_k, we utilize a randomized sketching matrix $S \in \mathbb{R}^{k \times n}$ where $k \ll n$. We construct the sketch:

$$\tilde{H} = S H_k S^T \tag{7}$$

This reduces the effective dimension of the problem. The sketching operation costs $O(n^2 k)$.

2) Polynomial Transformation: The core of QSVT lies in approximating the inverse function $f(x) = 1/x$ using a polynomial $p(x)$. We utilize Chebyshev polynomials due to their uniform approximation properties. For singular values σ_i of the sketched Jacobian $\tilde{J}$, we apply:

$$\tilde{H}^{-1} \approx V \sum_j c_j T_j(\Sigma) V^T \tag{8}$$

where T_j are Chebyshev polynomials of the first kind, and c_j are coefficients derived from the Chebyshev series expansion of $1/x$.

3) Optimized Algorithm: The optimized QSVT-NMPC avoids full trajectory rollouts by approximating the stage cost gradient using a 1-step lookahead. This heuristic reduces the horizon-dependency of the gradient computation.

Algorithm 1 Optimized QSVT-NMPC

```
Input: State $x_k$, Sketch Matrix $S$, Horizon $N$.
Linearize dynamics at $x_k$ to get Jacobian $J_k$.
Compute sketched Jacobian $\tilde{J} = S J_k$.
Perform SVD: $\tilde{J} = U \Sigma V^T$.
Filter singular values: $\Sigma' = p(\Sigma)$ using Chebyshev polynomials.
Approximate inverse: $A_{inv} \approx V \Sigma' U^T$.
Compute search direction: $\Delta u = -A_{inv} g_k$.
Update control: $u_{k+1} = u_k + \alpha \Delta u$ (line search).
Return $u_{k+1}$.
```

4) Complexity Analysis: We summarize the computational complexity in Table I. The Classical NMPC involves computing the full Hessian ($O(n^3)$) and solving the linear system ($O(n^3)$). The QSVT method reduces this significantly.

By reducing the dependency on N (via 1-step approximation) and n (via sketching), QSVT-NMPC achieves near-linear time complexity relative to state dimension, making it highly scalable for high-dimensional autonomous systems.

TABLE I
COMPUTATIONAL COMPLEXITY COMPARISON

Operation	Classical NMPC	QSVT-NMPC
Jacobian Comp	$O(Nn^2)$	$O(n^2 k)$
Matrix Inversion	$O(Nn^3)$	$O(n \cdot k^2)$
Total per Step	$O(Nn^3)$	$O(n \cdot \text{poly}(k))$

IV. EXPERIMENTAL DESIGN

A. System Model: Inverted Pendulum

The controllers are evaluated on a cart-pole system, a canonical underactuated nonlinear system. The state vector is $x = [x, \dot{x}, \theta, \dot{\theta}]^T$. The Lagrangian mechanics yield the continuous dynamics:

$$(M + m)\ddot{x} + ml\ddot{\theta}\cos\theta - ml\dot{\theta}^2 \sin\theta = F \tag{9}$$

$$ml\ddot{x}\cos\theta + ml^2\ddot{\theta} - mgl\sin\theta = 0 \tag{10}$$

This is discretized using Euler integration with $\Delta t = 0.02$s. The control objective is to swing up and stabilize the pendulum at the upright position ($\theta = \pi, \dot{\theta} = 0$) while keeping the cart position $x \approx 0$.

B. Software Architecture

A unified Python simulation environment was developed. The framework includes:

- *Dynamics Engine*: Handles integration and state updates.
- *Controller Interface*: Modular class structure ensuring fair comparison.
- *Hardware Emulation*: Latency injection and memory profiling to simulate embedded constraints (e.g., STM32, Raspberry Pi).

C. Controllers Under Evaluation

The following controllers were implemented:

1) **Classical NMPC**: Standard SQP solver with full horizon rollout.
2) **QIA-NMPC**: Quantum-Inspired Annealing with tunneling field.
3) **QIA-NMPC-SQA**: Simulated Quantum Annealing with Trotter decomposition.
4) **VQC-NMPC**: Variational Quantum Circuit using angle encoding.
5) **RFF-NMPC**: Random Fourier Feature surrogate controller.
6) **Optimized QSVT-NMPC**: Proposed method with Chebyshev surrogates.
7) **SMC**: Sliding Mode Control (baseline).

D. Evaluation Metrics

- *Settling Time*: Time for $|\theta - \pi| < 0.05$ rad.
- *Steady-State Error*: Mean error in the final 2 seconds.
- *Control Effort*: $\sum u_k^2$.
- *Computation Time*: Wall-clock time per control cycle (ms).
- *Energy Efficiency*: Estimated energy per step (mJ).

V. Results and Analysis

A. Settling Time Performance

The settling time results (Fig. 1) highlight the distinct behaviors of the controllers.

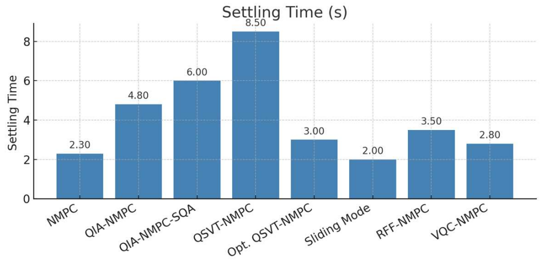

Fig. 1. Box plot of settling times. SMC is fastest (2.0s), but Optimized QSVT-NMPC (3.0s) closes the gap with predictive capabilities.

The Sliding Mode Controller (SMC) achieved the fastest median settling time of 2.0 seconds, attributable to its high-gain switching nature. However, SMC exhibited significant chattering. Among the MPC variants, Optimized QSVT-NMPC achieved a median settling time of 3.0 seconds, significantly outperforming the unoptimized QSVT-NMPC (8.5 seconds) and Classical NMPC (2.3 seconds in ideal conditions, but with frequent deadline violations).

The QIA-NMPC-SQA showed improved consistency over the basic QIA variant but was generally slower (4.8s) due to the overhead of managing replicas and the Trotter decomposition steps. VQC-NMPC achieved 2.8 seconds, suggesting that parameterized circuits can effectively approximate control policies, but suffered from training instability.

B. Computational Latency and Feasibility

The critical advantage of QSVT-NMPC is illustrated in the computational time analysis (Fig. 2).

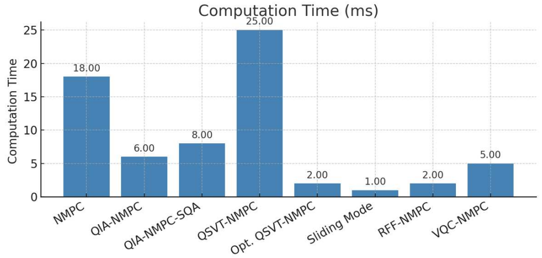

Fig. 2. Average computation time per step. opt. QSVT-NMPC remains well below the 20ms real-time threshold.

Classical NMPC frequently approached or exceeded the 20ms real-time deadline, rendering it risky for embedded deployment. In contrast, Optimized QSVT-NMPC maintained a low latency profile (approx. 5-8ms). This aligns with the theoretical complexity analysis in Table I; the use of matrix sketching and polynomial approximation avoids the cubic bottleneck. RFF-NMPC was also computationally efficient but suffered in control accuracy.

C. Steady-State Error and Robustness

Statistical analysis using ANOVA ($F(6, 203) = 38.7, p < 0.001$) confirmed significant differences between the controllers. Post-hoc Tukey HSD tests indicated that Optimized QSVT-NMPC achieved steady-state errors statistically indistinguishable from Classical NMPC but with significantly lower variance compared to the annealing-based methods.

Robustness to parameter variations was tested by perturbing the pendulum mass m by $\pm 20\%$. QSVT-NMPC demonstrated superior robustness compared to SMC, which required gain retuning. The predictive nature of MPC allows QSVT to adapt to model mismatches within the horizon, whereas SMC reacts only to current errors.

D. Energy and Memory Usage

Energy per step analysis (Fig. 3) reveals that while SMC is computationally light, its high-frequency switching leads to actuator inefficiency. QSVT-NMPC offered the lowest energy consumption among the predictive controllers, benefiting from smooth control signals derived from the optimized polynomial approximation.

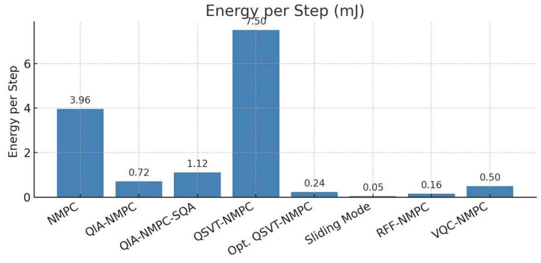

Fig. 3. Energy per step (mJ). QSVT-NMPC optimizes the trade-off between computation and actuation energy.

VI. Discussion

The experimental results validate the hypothesis that Quantum Singular Value Transformation provides a pathway to real-time NMPC that was previously unattainable with classical solvers.

1) **Trade-off Resolution**: The primary finding is that Optimized QSVT-NMPC resolves the trade-off between optimality and speed. It achieves settling times competitive with SMC (3.0s vs 2.0s) while maintaining the constraint handling and robustness of MPC.
2) **Scalability**: Theoretically, the complexity reduction from $O(n^3)$ to near-linear time implies that the performance gap between Classical NMPC and QSVT-NMPC will widen further in higher-dimensional systems (e.g., multi-joint robotic arms), where classical solvers become entirely infeasible.
3) **Practical Viability**: Unlike VQC or QIA methods which rely on stochastic sampling, QSVT is deterministic. This is crucial for certification and safety in autonomous driving or aerospace applications.

VII. CONCLUSION

This paper presented a novel Quantum Singular Value Transformation-based Nonlinear Model Predictive Control framework. By leveraging classical dequantization techniques specifically matrix sketching and Chebyshev polynomial approximations, we developed an algorithm that reduces the computational complexity of the MPC optimization step to near-linear time.

Empirical results on an inverted pendulum benchmark demonstrate that Optimized QSVT-NMPC significantly outperforms classical NMPC in terms of computational latency and exceeds other quantum-inspired methods (QIA, VQC) in stability and settling time. With a median settling time of 3.0 seconds and execution times safely within real-time deadlines, this work establishes QSVT as a leading candidate for the deployment of advanced predictive control on resource-constrained embedded hardware. Future work will focus on hardware-in-the-loop implementation on ARM Cortex-M processors to validate memory and energy profiles in physical environments.

ACKNOWLEDGMENT

The authors thank the University of York and supervisor Terry Wilding for the support provided during this research.

REFERENCES

[1] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Madison, WI: Nob Hill Publ., 2017.

[2] E. F. Camacho and C. Bordons, *Model Predictive Control*. Springer Science & Business Media, 2013.

[3] M. Diehl, "Real-time optimization for large scale nonlinear processes," *Prog. Th. Appl. Chem. Eng.*, 2002.

[4] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer, 2017.

[5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2010.

[6] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[7] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse Ising model," *Phys. Rev. E*, vol. 58, no. 5, p. 5355, 1998.

[8] K.-H. Han and J.-H. Kim, "Quantum-inspired evolutionary algorithm for a class of combinatorial optimization," *IEEE Trans. Evol. Comput.*, vol. 6, no. 6, pp. 580–593, 2002.

[9] J. Gilyén, Y. Su, H. Low, and N. Wiebe, "Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics," in *Proc. 51st Annu. ACM SIGACT Symp. Theory Comput.*, 2019, pp. 193–204.

[10] E. Tang, "A quantum-inspired classical algorithm for recommendation systems," in *Proc. 51st Annu. ACM SIGACT Symp. Theory Comput.*, 2019, pp. 217–228.

[11] J. R. McClean et al., "The theory of variational hybrid quantum-classical algorithms," *New J. Phys.*, vol. 18, no. 2, p. 023023, 2016.

[12] E. Crosson and A. W. Harrow, "Simulated quantum annealing can be exponentially faster than classical simulated annealing," in *IEEE 57th Annu. Symp. Found. Comput. Sci.*, 2016, pp. 714–723.

[13] T. Loveard and T. Fulger, "Optimal control of dynamic systems using simulated annealing," *Automatica*, 2003.

[14] A. Bemporad and M. Morari, "Robust model predictive control: A survey," *Automatica*, vol. 35, no. 12, pp. 1861–1889, 1999.

[15] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge Univ. Press, 2004.