

QuantumShield AI: an AI-Powered Post-Quantum Cryptography Recommendation Platform

Bandari Venkatesh*, P. Suresh Varma

Department of Computer Science and Engineering, Adikavi Nannaya University College of Engineering, Rajamahendravaram, Andhra Pradesh, India.

ABSTRACT

The finalization of the NIST post-quantum cryptography (PQC) standards ML-KEM (FIPS 203), ML-DSA (FIPS 204), and SLH-DSA (FIPS 205), with FN-DSA (draft FIPS 206), has shifted the migration question from whether to adopt PQC to which standard best fits an organization's operational and regulatory context. Existing studies benchmark these algorithms thoroughly but stop short of a selection mechanism, leaving practitioners to interpret timing tables by hand. This paper presents QuantumShield AI, a decision-support platform that formalizes PQC algorithm selection as a multi-criteria decision-making (MCDM) problem, encodes the selection task as a Quadratic Unconstrained Binary Optimization (QUBO), and solves it with the Quantum Approximate Optimization Algorithm (QAOA) on a statevector simulator, cross-validated by an exact classical solver. Six industry profiles — banking, healthcare, IoT/edge, government, e-commerce, and defense supply criteria weightings grounded in sector requirements. The framework is delivered as a three-tier platform combining a React front end, a Spring Boot service layer with JWT-based role-controlled access, and a Python/Qiskit quantum microservice over PostgreSQL. To our knowledge this is the first framework to operationalize selection across the complete finalized NIST portfolio, converting sector requirements into an auditable, reproducible recommendation. Across all six profiles the QAOA solution recovers the exact optimum (approximation ratio = 1.0), validating the pipeline; we are explicit that at the four-algorithm scale this certifies correctness rather than quantum advantage, and the QUBO–QAOA formulation is carried as a forward-compatible path to NP-hard enterprise-portfolio selection. The contribution is a novel, deployable decision methodology and its engineered realization.

Keywords: Post-quantum cryptography, NIST FIPS 203/204/205/206, Multi-criteria decision-making, QUBO, QAOA, Decision-support systems, Cryptographic migration

INTRODUCTION

Public-key cryptography underpins nearly every secure digital interaction, and nearly all of it rests on two mathematical problems integer factorization and the discrete logarithm that Shor's algorithm solves in polynomial time on a cryptographically relevant quantum computer. The economics of the threat are already present: an adversary can record encrypted traffic today and decrypt it once suitable hardware exists, so long-lived secrets in banking, healthcare, and government carry present-day risk. This is the harvest-now-decrypt-later problem, and it is the reason cryptographic migration cannot wait for quantum hardware to arrive.

The National Institute of Standards and Technology answered the cryptographic half of this problem in

August 2024, finalizing ML-KEM for key encapsulation and ML-DSA and SLH-DSA for digital signatures, with FN-DSA following as draft FIPS 206. Standardization, however, created a second problem it did not solve: selection. The finalized algorithms differ sharply in security assumptions, key and signature sizes, and computational cost, and the empirical literature confirms that no single algorithm dominates across all criteria. The same studies conclude, in effect, that the right choice depends on the needs of a particular case, yet they provide no mechanism for making that determination.

A bank optimizing for transaction latency, a hospital bound by decade-long data-retention mandates, and a defense agency requiring maximum-strength parameters face the same four-algorithm menu but should not reach the same answer. Today they reach any

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.



answer only through manual expert interpretation of benchmark tables a process that is slow, inconsistent across organizations, and unauditable. This paper closes that gap.

The novelty of this work is that it is, to our knowledge, the first system to operationalize PQC algorithm selection: where the benchmarking literature quantifies trade-offs and the deployment literature fixes a single algorithm by fiat, we convert sector requirements into an auditable, per-context recommendation over the complete finalized NIST portfolio, and deliver that capability as a running platform rather than a paper design. Concretely, the contributions are: (1) a formal MCDM model that converts sector requirements into weighted criteria over the four NIST algorithms; (2) a QUBO encoding of the selection problem, solved by QAOA and cross-validated against an exact solver, giving a solver architecture that is correct at native scale and forward-compatible with NP-hard enterprise-portfolio selection; (3) six calibrated industry profiles spanning banking, healthcare, IoT/edge, government, e-commerce, and defense; and (4) a fully engineered, security-hardened three-tier platform that delivers auditable recommendations. We are explicit about scope: at four algorithms the search space is exhaustively verifiable, so QAOA's approximation ratio of 1.0 certifies pipeline correctness rather than a performance advantage, and FN-DSA is labeled draft in every output. This scoping bounds the quantum claim; it does not bound the contribution, which is the selection methodology and its realization.

BACKGROUND

The quantum threat to classical asymmetric cryptography is asymmetric in its own way: Shor's algorithm collapses the hardness of both factorization (RSA) and discrete logarithms (ECC, Diffie–Hellman) to polynomial time, while Grover's algorithm merely halves the effective key length of symmetric ciphers, which AES-256 absorbs comfortably. The practical consequence is that key establishment and digital signatures the asymmetric primitives are where migration is urgent, and both halves of the NIST portfolio target exactly these functions.

The migration challenge is not throughput but decision-making. An organization must choose, for each function it secures, among standardized algorithms whose trade-offs pull in different directions: lattice-based key encapsulation (ML-KEM) offers compact

keys and fast operations; lattice-based signatures (ML-DSA) balance size and speed; hash-based signatures (SLH-DSA) rest on the most conservative security assumptions but produce large signatures; and FN-DSA, still in draft, offers compact signatures at the cost of implementation complexity and non-final status. No published framework converts these documented trade-offs into a per-context recommendation. That absence, rather than any deficiency in the algorithms themselves, motivates the present work.

3. LITERATURE REVIEW

3.1. Performance characterization of PQC

The most mature strand of the literature benchmarks PQC algorithms empirically. Opilka et al. [1] implemented seventeen signature variants Dilithium (levels 2/3/5), Falcon (512/1024), and twelve SPHINCS+ configurations through the liboqs library on a controlled virtual machine, measuring key generation, signing, and verification across 10 MB–1 GB files over 100 iterations with standard-error reporting. They found that Dilithium5 signs a 1 GB file with roughly 27.7% overhead relative to RSA-2048 while providing NIST security level 5, and concluded that algorithm choice depends on the use case. This study is the empirical foundation on which the present work builds, and its explicit absence of any selection mechanism is our point of departure.

3.2. Deployment-oriented PQC integration

A second strand embeds a chosen PQC algorithm into an application. Nasir et al. [2] integrated ML-KEM (Kyber768) into a real-time sealed-bid property-auction platform, reporting complete tamper detection across adversarial trials and sub-millisecond per-bid encryption; notably, the algorithm is fixed a priori and ML-DSA integration is deferred to future work. Survey treatments such as Nwaga and Idima [3] map the PQC family taxonomy onto blockchain and cloud infrastructures and articulate the harvest-now-decrypt-later risk, but predate FIPS finalization and likewise provide no decision framework. Across benchmarking, deployment, and survey work the pattern repeats: rich evidence about algorithms, no mechanism for choosing among them.

3.3. Decision methods and quantum optimization

MCDM techniques weighted-sum models, AHP, and TOPSIS are established tools for technology selection

under conflicting criteria. QUBO is the canonical encoding for binary selection problems targeted by quantum and quantum-inspired solvers, and QAOA [4] is the leading near-term variational algorithm for such encodings, with Lucas [5] cataloguing Ising formulations of many NP-hard problems. To our knowledge, no prior work composes these into a decision-support system for NIST PQC selection, nor delivers such a system as a deployable multi-tier platform.

4. RESEARCH GAP AND POSITIONING

Synthesizing the review, five gaps emerge: benchmark studies quantify trade-offs but do not decide; deployment studies decide by fiat, without justifying the choice against alternatives or contexts; no study spans the complete finalized NIST portfolio including draft FN-DSA under its FIPS identity; no study connects sector-specific regulatory reality to algorithm choice; and no study examines whether quantum optimization itself can serve the migration-planning problem. Table 1 positions QuantumShield AI against the two most closely related works.

Table 1: Positioning against the closest related work.

Aspect	Opilka et al. [1]	Nasir et al. [2]	QuantumShield AI
Goal	Benchmark signatures	Deploy one KEM	Select best algorithm per context
Portfolio	Signatures only	ML-KEM only	All four NIST standards
Decision layer	None	Fixed a priori	Formal MCDM + QUBO/QAOA
Sector context	None	Single (real estate)	Six industry profiles
Quantum used	No	No	Yes (QAOA), honestly scoped
Delivery	Python CLI	Microservice app	Three-tier platform

5. PROPOSED METHODOLOGY

QuantumShield AI assigns each stage of the decision to the paradigm that handles it most directly: a classical MCDM model reduces sector requirements to per-algorithm scores; a QUBO encoding casts selection as binary optimization; and a QAOA solver, validated by exhaustive search, produces the recommendation. The full pipeline is shown in Fig. 2, and each stage is specified below at implementation-audit level of detail.

5.1. MCDM formulation

Let $A = \{a_1, a_2, a_3, a_4\} = \{\text{ML-KEM, ML-DSA, FN-DSA, SLH-DSA}\}$ be the candidate set, evaluated over four decision criteria $C = \{\text{security, speed, size, longevity}\}$. Security is the NIST security level of the parameter set (benefit); speed is the algorithm's operations per second (benefit); size is the combined public-key and artifact ciphertext or signature byte count (a cost criterion, so smaller is better); and longevity rates conservative, long-horizon durability of the underlying assumption (benefit), on which hash-based SLH-DSA scores highest. Each raw criterion value $v_j(a)$ is min-max normalized to $[0, 1]$ across the candidate set: benefit criteria as $n_j(a) = (v_j(a) - \min_j) / (\max_j - \min_j)$ and cost criteria as $n_j(a) = (\max_j - v_j(a)) / (\max_j - \min_j)$, with $n_j(a) = 1$ where $\max_j$

$= \min_j$. Each industry profile p defines a weight vector $w(p) = (w_{\text{security}}, w_{\text{speed}}, w_{\text{size}}, w_{\text{longevity}})$ with $\sum_j w_j(p) = 1$ (validated server-side to within a 0.02 tolerance). The composite score of algorithm a under profile p is

$$S_p(a) = \sum_j w_j(p) \cdot n_j(a), j \in \{\text{security, speed, size, longevity}\}, (1)$$

where $n_j(a)$ is the normalized value of criterion j for algorithm a . This weighted-sum model yields a per-profile ranking with complete traceability from weights to recommendation, and because every criterion is normalized against the same candidate set scores are directly comparable across algorithms within a profile.

5.2. QUBO encoding

The recommendation is a single-best selection: binary variables $x \in \{0, 1\}^4$ with $x_i = 1$ iff algorithm a_i is recommended, subject to a one-hot constraint enforcing exactly one selection. The objective maximizes the total MCDM score, expressed as a minimization with a quadratic penalty:

$$H(x) = -\sum_i S_p(a_i) \cdot x_i + \lambda \cdot (\sum_i x_i - 1)^2, (2)$$

where the penalty coefficient is set adaptively to dominate the largest achievable objective gain, $\lambda = \max(1, 2 \cdot \sum_i |S_p(a_i)|)$. This construction guarantees that any solution violating the one-hot constraint incurs a

penalty larger than any score benefit, so the constrained optimum coincides with the unconstrained ground state of H . Expanding the penalty yields the quadratic form $H(x) = \sum_i Q_{ii} x_i + \sum_{i < j} Q_{ij} x_i x_j$ with $Q_{ii} = -S_p(a_i) - \lambda$ and $Q_{ij} = 2\lambda$, i.e., the QUBO matrix Q of Fig. 5. The problem is expressed as a QuadraticProgram with an explicit equality constraint ($\sum_i x_i = 1$) and solved through Qiskit's MinimumEigenOptimizer.

5.3. QAOA solution and exact ground truth

The QUBO Hamiltonian is solved by QAOA with $p = 3$ repetition layers, using a COBYLA classical optimizer (maximum 150 iterations) and the Qiskit Aer sampler primitive. QAOA prepares the variational state

$$|\psi(\beta, \gamma)\rangle = \prod_{l=1..3} \exp(-i\beta_l H_m) \exp(-i\gamma_l H_c) |+\rangle^{\otimes 4} \quad (3)$$

with mixer $H_m = \sum_i X_i$; COBYLA tunes the six variational parameters (β, γ) to minimize the expected energy, and the optimized state yields the recommended index. For a four-algorithm instance the constrained optimum is directly computable as the highest-scoring feasible selection, and this exact argmax serves as ground truth. The approximation ratio is reported as

$$r = S_p(a_{\text{chosen}}) / S_p(a_{\text{exact}}), \quad (4)$$

measuring whether QAOA recovered the ground-truth optimum. The service is designed to degrade gracefully: if the quantum backend is unavailable or a QAOA run fails, it falls back to the exact solver and reports the solver path explicitly. In the experiments reported here, QAOA executed successfully and returned the ground-truth optimum on every profile (Section 8). We emphasize that $r = 1.0$ at this scale is a correctness result QAOA recovering an optimum that is independently and trivially computable and not evidence of quantum advantage. The value of the formulation is that it provides a validated pipeline and the exact-solver reference that any future scaling to NP-hard portfolio instances, where ground truth is no longer computable by enumeration, will require.

6. SYSTEM ARCHITECTURE




Figure 1: Three-tier architecture of QuantumShield AI, comprising the React presentation tier, the Spring Boot application tier, and the Python/Qiskit quantum microservice over PostgreSQL.

The platform comprises three services. The presentation tier is a React/TypeScript single-page application (Vite, port 5173) providing profile selection, custom weighting, results visualization, and administrative screens. The application tier is a Java 21 Spring Boot service (port 8080) implementing business logic, PostgreSQL 16 persistence, and security: stateless JWT authentication with a 30-minute token lifetime and role-based access control across Admin, Analyst, and Viewer roles, with self-registration restricted to the non-administrative roles. The quantum tier is a Python 3.12 FastAPI microservice (port 8000, entry-point `app.main:app`) exposing the MCDM scorer, QUBO builder, QAOA solver, and exact validator over REST, executing on the Qiskit Aer statevector simulator. Python is pinned to 3.12 for Qiskit/NumPy wheel compatibility.




Figure 2: End-to-end recommendation workflow, from authenticated intake through MCDM scoring, QUBO construction, QAOA execution, and exact validation to a persisted, ranked recommendation.

Table 2: NIST algorithm parameters used as decision-matrix inputs. Starred columns are relative engineering proxies pending measured benchmarks; all other values are from the standardized parameter sets.

Algorithm	NIST standard	Sec. level	Public key (B)	Artifact (B)	Ops/s*	Long.*
ML-KEM	FIPS 203	3	1184	1088	9500	7
ML-DSA	FIPS 204	3	1952	3309	4200	7
FN-DSA	FIPS 206 (draft)	1	897	666	1800	6
SLH-DSA	FIPS 205	1	32	17088	60	10

The six industry profiles of Table 3 constitute the experimental conditions: each supplies a weight vector over the four criteria (validated server-side to sum to 1.0 within a 0.02 tolerance) and is executed through the full MCDM → QUBO → QAOA pipeline. The quantum solver is configured with QAOA depth $p = 3$, a COBYLA optimizer capped at 150 iterations, and

A request flows from authentication and profile intake, through MCDM scoring and QUBO construction, into QAOA execution and exact validation, and finally to a persisted, ranked recommendation carrying the FN-DSA draft-status disclaimer. Isolating the quantum solver behind a REST boundary lets the decision engine evolve for instance toward hardware backends without disturbing the security or presentation tiers.

7. EXPERIMENTAL SETUP

The decision engine is exercised on the four finalized or near-finalized NIST algorithms, parameterized as in Table 2. Security level and the public-key and artifact (ciphertext or signature) byte sizes are taken directly from the standardized parameter sets (ML-KEM-768, ML-DSA-65, SLH-DSA-128f, and draft FN-DSA-512). The operations-per-second and long-term-validity figures are relative engineering proxies, included to exercise the speed and longevity criteria and flagged in the system for replacement with measured liboqs benchmarks; the numeric scores in Section 8 should therefore be read as demonstrating the mechanism rather than as certified performance measurements.

the Qiskit Aer sampler primitive (qiskit 1.2.4, qiskit-aer 0.15.1); the exact argmax over the feasible set is computed as ground truth, and the service falls back to the exact solver automatically on any quantum-backend failure. All results in Section 8 are produced by this pipeline on the data of Tables 2 and 3.

Table 3: Industry profiles and their criterion weight vectors (security / speed / size / longevity), with the algorithm selected by the deployed engine.

Profile	Sector	Weights (sec / speed / size / long.)	Selected
Banking	Finance	0.40 / 0.25 / 0.15 / 0.20	ML-KEM
E-Commerce	Retail	0.25 / 0.45 / 0.20 / 0.10	ML-KEM
IoT/Edge	Embedded	0.15 / 0.40 / 0.40 / 0.05	ML-KEM
Healthcare	Health	0.35 / 0.10 / 0.10 / 0.45	ML-KEM
Government Archive	Public Sector	0.10 / 0.05 / 0.10 / 0.75	SLH-DSA
Defense	Defense	0.15 / 0.05 / 0.05 / 0.75	SLH-DSA

RESULTS AND DISCUSSION

8.1. Per-profile MCDM scores




Figure 3: Composite MCDM scores $S_p(a)$ of the four NIST algorithms under the six industry profiles, computed by the deployed engine on the inputs of Tables 2 and 3.

Fig. 3 reports the composite scores $S_p(a)$ produced by the deployed engine for all four algorithms under the six profiles. The weight vectors reshape the ranking as intended. In the profiles that emphasize security and speed banking, e-commerce, and IoT/edge ML-KEM leads by a wide margin (0.843, 0.916, and 0.944 respectively), reflecting its high throughput and compact key-and-ciphertext footprint. In the two profiles that place 0.75 weight on long-horizon durability government archive and defense the ranking inverts and hash-based SLH-DSA wins (0.750 in both), consistent with its conservative security assumption. Healthcare is the closest and most interesting contest:

under the weighted-sum model ML-KEM leads (0.658) ahead of ML-DSA (0.583) and SLH-DSA (0.450). FN-DSA is last or near-last in every profile, penalized by its draft status' low longevity rating and its lower security level despite the smallest signatures. The leading algorithm therefore changes with context precisely the dependence the prior literature identified but did not operationalize.

8.2. Method robustness: weighted-sum versus TOPSIS and weight sensitivity

To test whether a recommendation is an artifact of the weighted-sum model (WSM), we re-ran all six profiles through TOPSIS on the same normalized decision matrix and weights. The two methods agree on the top choice in five of the six profiles (Table 4). The single divergence is healthcare: WSM ranks ML-KEM first, whereas TOPSIS which rewards proximity to the per-criterion ideal and distance from the anti-ideal elevates SLH-DSA on the strength of its maximal longevity and conservative security posture. This is a genuine property of the healthcare weighting rather than a defect: the two algorithms encode different risk postures, and the disagreement usefully flags healthcare as the profile most warranting expert review. We further probed stability by perturbing each weight by ± 0.10 and renormalizing, giving eight perturbations per profile; no perturbation changed the top recommendation in any profile, indicating the selections are robust to modest weight mis-specification.

Table 4: Top recommendation under the weighted-sum model versus TOPSIS, on identical weights and normalized inputs.

Profile	WSM pick	WSM score	TOPSIS pick	Agree
Banking	ML-KEM	0.843	ML-KEM	Yes
E-Commerce	ML-KEM	0.916	ML-KEM	Yes
IoT/Edge	ML-KEM	0.944	ML-KEM	Yes
Healthcare	ML-KEM	0.658	SLH-DSA	No
Government Archive	SLH-DSA	0.750	SLH-DSA	Yes
Defense	SLH-DSA	0.750	SLH-DSA	Yes

8.3. Quantum solver versus exact baseline




Figure 4: Median solve time per profile, QAOA (Aer statevector, $p = 3$) versus exact enumeration of the 2^4 feasible states, on a log scale. Both solvers return the identical selection (approximation ratio = 1.0).

Table 5 and Fig. 4 report the quantum solver against the exact baseline. On every profile QAOA (Aer

statevector, $p = 3$, COBYLA at 150 iterations) returned exactly the ground-truth optimum — an approximation ratio of 1.0. The cost of that agreement is decisive: QAOA required a median of about 315 ms per instance, while exhaustive enumeration of the $2^4 = 16$ feasible states returned the same answer in roughly 6.5 μ s, a factor of approximately 49,000 in favor of the classical baseline. We present this as the paper's most direct honesty check: at four algorithms the quantum route is not merely without advantage, it is about 5×10^4 times slower for an identical result. QAOA's role at this scale is solely to demonstrate that the MCDM $\rightarrow$ QUBO $\rightarrow$ QAOA pipeline is wired correctly against a certifiable ground truth; its practical relevance is deferred to the portfolio regime of Section 8.5.

Table 5: Per-profile approximation ratio and median solve time, QAOA versus exact enumeration (steady state, Qiskit Aer statevector; warm-up runs discarded).

Profile	Ratio r	QAOA (ms)	Exact (μ s)	Exact faster by
Banking	1.00	304.0	6.85	$\sim 44,000\times$
E-Commerce	1.00	315.5	6.65	$\sim 47,000\times$
IoT/Edge	1.00	324.8	6.26	$\sim 52,000\times$
Healthcare	1.00	314.1	6.36	$\sim 49,000\times$
Government Archive	1.00	337.8	6.21	$\sim 54,000\times$
Defense	1.00	304.8	7.00	$\sim 44,000\times$

8.4. QUBO structure




Figure 5: QUBO matrix Q for the banking profile, built with the adaptive penalty $\lambda = \max(1, 2 \cdot \sum |S_p(a_i)|) = 3.83$. The most negative diagonal entry marks the selected algorithm.

Fig. 5 shows the QUBO matrix Q for the banking profile, built with the adaptive penalty $\lambda = \max(1,$

$2 \cdot \sum |S_p(a_i)|) = 3.83$. Diagonal entries $Q_{ii} = -S_p(a_i) - \lambda$ carry the negated scores offset by the penalty, so the most negative diagonal (ML-KEM, -4.67) marks the profile winner; the uniform off-diagonals $Q_{ij} = 2\lambda = 7.65$ penalize selecting more than one algorithm. Because λ dominates the score spread, the feasible optimum is trivially certifiable at this size which is exactly why the interesting regime is the larger portfolio problem, where Q grows quadratically in the number of system-to-algorithm assignments and exhaustive certification fails.

8.5. Discussion and threats to validity

The central result must be read precisely: an approximation ratio of 1.0 across all six profiles is agreement between QAOA and exhaustive search over sixteen states a correctness certificate, not quantum advantage, which cannot exist at this size and, as Section 8.3 quantifies, comes at roughly a 49,000 $\times$ time penalty. Quantum relevance is prospective: when the decision expands to enterprise portfolios (many systems $\times$ many algorithms $\times$ interoperability, budget,

and migration-order constraints), the selection problem becomes NP-hard and QAOA-class heuristics become a plausible tool. Three limitations bound the present findings. First, execution uses a statevector simulator rather than quantum hardware, so NISQ noise behavior is untested. Second, and most important for reproducibility, the speed and longevity criteria are relative engineering proxies rather than measured quantities; security level and byte sizes come from the standardized parameter sets, but the operations-per-second and long-term-validity figures must be replaced with measured liboqs benchmarks before the numeric scores are treated as anything beyond a demonstration of the mechanism. Third, the profile weights are expert-calibrated rather than empirically derived, so each recommendation inherits that calibration which is why the TOPSIS cross-check and the weight-sensitivity analysis of Section 8.2 are reported alongside the headline recommendations.

CONCLUSION

QuantumShield AI establishes a clear contribution: to our knowledge it is the first system to turn PQC algorithm selection from manual expert judgment into an auditable, reproducible computation over the complete finalized NIST portfolio. Where prior work concluded only that algorithm choice depends on the needs of a particular case, QuantumShield AI supplies the mechanism by which that dependence is computed, and delivers it as a security-hardened, three-tier platform rather than a paper design. The QUBO–QAOA formulation is validated end to end against an exact solver, yielding a solver architecture that is correct at native scale and forward-compatible with the NP-hard enterprise-portfolio regime in which quantum-class heuristics become relevant. We are precise about scope: what this paper establishes is methodological and architectural novelty, not demonstrated quantum advantage, and we claim only the former. Future work extends the framework to enterprise-portfolio selection, where ground truth can no longer be obtained by enumeration and the choice of solver begins to carry real weight.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

ACKNOWLEDGEMENTS

The authors thank the Department of Computer Science and Engineering, Adikavi Nannaya University College of Engineering, Rajamahendravaram, for the computational facilities and academic support that made this work possible.

REFERENCES

1. Opiłka F, Niemiec M, Gagliardi M, Kourtis MA. Performance analysis of post-quantum cryptography algorithms for digital signature. *Appl Sci* 2024;14 (12):4994. <https://doi.org/10.3390/app14124994>.
2. Nasir E, et al. QTrustBid: post-quantum cryptography and AI-driven recommendations for real-time property auction platform. *Spectr Eng Sci* 2026;4 (5):2077–93.
3. Nwaga P, Idima S. Post-quantum cryptographic algorithms for secure communication in decentralized blockchain and cloud infrastructure. *Int J Comput Appl Technol Res* 2022;11 (4):155–70.
4. Farhi E, Goldstone J, Gutmann S. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*; 2014.
5. Lucas A. Ising formulations of many NP problems. *Front Phys* 2014;2:5. <https://doi.org/10.3389/fphy.2014.00005>.
6. National Institute of Standards and Technology. Module-lattice-based key-encapsulation mechanism standard. FIPS 203. Washington (DC): U.S. Department of Commerce; 2024. <https://doi.org/10.6028/NIST.FIPS.203>.
7. National Institute of Standards and Technology. Module-lattice-based digital signature standard. FIPS 204. Washington (DC): U.S. Department of Commerce; 2024. <https://doi.org/10.6028/NIST.FIPS.204>.
8. National Institute of Standards and Technology. Stateless hash-based digital signature standard. FIPS 205. Washington (DC): U.S. Department of Commerce; 2024. <https://doi.org/10.6028/NIST.FIPS.205>.
9. Shor PW. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J Comput* 1997;26 (5):1484–509.

10. Preskill J. Quantum computing in the NISQ era and beyond. *Quantum* 2018;2:79.
11. Qiskit Contributors. Qiskit: an open-source framework for quantum computing [software]; 2024. <https://www.qiskit.org> (accessed 31 July 2026).
12. Triantaphyllou E. Multi-criteria decision making methods: a comparative study. Boston (MA): Springer; 2000.

HOW TO CITE: Bandari Venkatesh, P. Suresh Varma, QuantumShield AI: an AI-Powered Post-Quantum Cryptography Recommendation Platform, *Int. J. Sci. R. Tech.*, 2026, 3 (10), 626-634. <https://doi.org/10.5281/zenodo.23261604>