iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.77939

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

Quantum Annealing for Combinatorial Optimization: Foundations, Architectures, Benchmarks, and Emerging Directions

Rudraksh Sharma[1], Ravi Katukam[2], Arjun Nagulapally[3]
AIONOS Block B, Third Floor Ramky Selenium, Plot No. 31 & 32 Financial District, Nanakramguda Hyderabad, Telangana, India 500032

Abstract: ***Combinatorial optimization is a foundation of numerous scientific, industrial, and social decision-making questions, but its usefulness is mostly restricted by the exponential complexity of the problem of classical computability. The quantum annealing (QA) has become one of the leading analogue quantum systems to address such problems, whereby optimization problems are converted into physical energy structures and quantum fluctuations are utilized to perform systematic exploration. It provides a critical, synthesized, and comprehensive review of quantum annealing as applied to combinatorial optimization, including the theoretical basis, hardware designs, algorithm methods, issues of embedding and encoding, benchmarking procedures, applications, and integration with quantum algorithms based on gates and solvers on a classical computer. We build a comprehensive taxonomy between adiabatic dynamics, Ising and QUBO models, stoquastic and non-stoquastic Hamiltonians, and diabatic transitions to state-of-the-art flux-qubit annealers, new architectures, and hybrid quantum classical pipelines. We show that the overhead of embedding and encoding and the benchmarking approach taken are the most significant factors in determining scalability and performance, but not the raw number of qubits. Transportation, energy systems, robotics, finance, pharmacological discovery, and machine learning are just some examples of the domain-specific case studies that use QA with minimal empirical utility as a hybrid refinement engine and not an independent solver. Moreover, we introduce a stringent evaluation of the existing benchmarking culture, clarify the causal links between QA and QAOA as well as between VQE and structural predicaments, and reveal the institutional roots of the existing restrictions. We conclude by identifying a research roadmap in the future, focusing on the annealing hardness characterization, Stoquastic control manipulation, embedding rules automation, architecture design, and development of principled quantum advantage. It creates a benchmark reference to those in the field and research and deployment of scalable, reliable, and application-relevant quantum optimization.***

Keywords: ***Quantum Annealing; Combinatorial Optimization; QUBO and Ising Models; Hybrid Quantum-Classical Optimization; Benchmarking and Quantum Advantage; Embedding and Encoding Strategies.***

I. INTRODUCTION

A. The Grand Optimization Problem

Decision-making in science, engineering, and industry is based on combinatorial optimization,requiring optimal choices from exponentially large discrete spaces, where even small binary systems scale as 2^N.It is not incidental but structural in nature because combinatorial design theory shows that problems with heavily constrained problems very quickly become solvable with the help of hard algebraic constraints. Incidence-letter constructions of this rigidity can be explained by this, including why constructions like the Hadamard matrices, finite projective planes, etc. can only exist at isolated parameters, and many of these are forced as impossible. In this intrinsic hardness, formalized by NP-hardness, dependencies and constraints cause combinatorial explosion instead of modelling complexity. Quadratic Unconstrained Binary Optimization (QUBO) offers a overall abstraction of this type to problems and, through its active equivalence to the Ising model, the physical basis of quantum annealing as an optimization model.[1],[2],[3].

B. Where Quantum comes in

Classical optimization algorithms, such as branch-and-bound, simulated annealing, tabu search and meta-heuristics, eventually face exponentially large, rugged search spaces, in that simulated annealing takes impractically slow logarithmic cooling to converge.

Quantum annealing (QA) is a different paradigm, in which quantum tunnelling across energy barriers is used, via a transverse-field driver. On the physical level, a time-dependent transverse-field Ising model of QA consists of a time-dependent transverse-field Ising Hamilton with transverse field G(t) running slowly to direct the system towards a problem ground state. Heuristic superiority Numerical research shows that QA is able to perform better than simulated annealing on a subset of problems with high ratios of tall, narrow energy barriers and thus the heuristic advantage of QA. Although adiabatic quantum computation is very similar to real QA devices, these are open, noisy, systems, in which non-adiabatic transitions and dissipation are a central part of determining their capability and limits[4], [5], [6], [7].

Table 1: Mapping Optimization Concepts to Quantum Annealing Physics

Optimization Concept	QA / Physics Analog	Mathematical Object	Practical Implication
Objective function	Problem Hamiltonian	Ising / QUBO	Defines energy landscape
Feasible solution	Ground state	Lowest eigenstate	Optimal solution
Constraint violation	Energy penalty	Coupling strength	Precision bottleneck
Local minimum	Excited state	Local eigenstate	Trapping risk
Neighborhood move	Quantum tunneling	Transverse field	Barrier traversal

C. Historical Evolution

Quantum annealing was invented as an expansion of simulated annealing using quantum effects, which allowed quicker convergence through quantum tunnelling through possible barriers. Transverse Ising model numerical simulations validate that the success rate of ground states of quantum schedules at the end of quantum schedules is higher than the performance of classical counterparts. Further theoretical work has extended quantum annealing to the adiabatic spectral process, allowing it to apply to many body systems, computational phase transitions, interaction with the environment, and NP-hard problems, a rich ecosystem of quantum annealers has been built, including theory, simulation and programmable machines[4], [5], [8].

D. Purpose and contributions

This is the unifying reference to quantum annealing as an optimization technique in combinatorial problems with a comprehensive description of the theory, hardware, algorithmic schemes, embedding strategies and benchmarking platform and usage. It also adds to a comprehensive taxonomy of literature in the field of quantum annealing and also critically analyzes methods of benchmarking analogue quantum annealers in comparison with both classical benchmarking and gate-based benchmarking procedures. Lastly, it identifies a roadmap of research with a view to address outstanding problems, such as scalability around quantum phase transitions, non-stoquastic and diabatic protocols, error-suppression methods, co-design of architecture and theory, and in place of robust standards to prove quantum advantage.

II. MATHEMATICAL AND PHYSICAL FOUNDATIONS

The overall concept of quantum annealing (QA) is to model a classical combinatorial optimization problem by a Hamilton and to then drive a quantum system towards the lowest-energy eigenstate of that Hamilton which contains the desired solution.

A. Adiabatic quantum computation

Quantum adiabatic computation (QAC) is based on quantum adiabatic theorem. Given a time-dependent Hamiltonian, which depends on time, and interpolates two Hamiltonians: an initial HamiltonianH_{init} and a problem Hamiltonian, encoded in its Hamiltonian distribution, which is H_{prob}. As this interpolation is performed at a slow rate, the system will be in the instantaneous ground state at all times during the evolution process, and finally reach the ground state ofH_{prob}.

The principle is literally realized in transverse field Ising models, in which a weakening driver Hamiltonian is used to push the system into the ground state of the problem. The minimum spectral gap at any given time is an adiabatic convergence which strictly obeys the inverse-square rule of the runtime. Less gaps require longer periods of evolution to reduce excitations, with this correlation existing in digitized as well as analogue adiabatic quantum computing paradigms. In the case of gate-based architectures, the continuum-valued Hamiltonian evolution is split into discrete Trotter slices, which are implemented by quantum gates, which hence shed light on the feasibility of solving multi-qubit problems by a large batch of gates[4],[5],[9].

B. QA vs AQC

Adiabatic quantum computation (AQC) is based on the assumption of unitary evolution which is perfectly coherent, and the duration of computation is restricted to the spectral -gap scaling with strong conservation of the ground state.

Modern quantum annealing (QA) systems are allowed to violate adiabaticity; even in non-adiabatic environments, both thermal activities, decoherence effects, or relaxation mechanisms can sometimes help to accelerate the convergence of the calculation, making it easy to escape shallow local minima. QA is therefore not to be regarded as a raucous form of AQC; rather, it is a complex optimization strategy which combines quantum tunneling, thermalization and device effects to pass through large energy barriers.[10],[11].

C. Ising and QUBO models

Quantum annealing solves the optimization problems by the mapping of the binary decision variables to the spin configurations. The Ising model takes the form of a classical Hamiltonian, with spins being the coded variables $s_i \in \{-1, +1\}$.

$$H_{\text{Ising}}(s) = \sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j \ \#(2)$$

The QUBO (Quadratic Unconstrained Binary Optimization) problem model substitutes spins by binary variables, x_i in [0,1], i.e. the objective is x^TQx, and Q is a symmetric matrix which describes all the interactions. This is accomplished by a direct transformation of QUBO to Ising variables implemented via s_i= 2x_i-1 to make the two mathematical formulations mathematically identical, but this makes them hardware-distinguishable.[5][8].

An example traditional method of converting a real-world combinatorial problem has the following steps: First, the goal function is reformulated as an expression in discrete variables, followed by an encoding of domain choices into binary decision variables. The imposition of constraints is done using quadratic penalty terms and as much as a high-order interaction exists, it can be systematically reduced to quadratic form but at the cost of additional auxiliary variables. This is followed by the formulation of the quadratic unconstrained binary optimization (QUBO) problem that is converted to an Ising model of spins of ±1 which can be implemented on physical quantum annealing machines.

D. Stoquastic vs non-stoquastic Hamiltonians

Simulation difficulty hinges on stoquastic: Hamiltonians with real, non-positive off-diagonal terms avoid the sign problem and are efficiently tractable by quantum Monte Carlo.

Non-stoquastic Hamiltonians break this property, causing exponential sampling bottlenecks, while stoquastic systems uniquely combine classical stimulability with NP-hard combinatorial structure[12],[13],[14],[15].

Table 2: Stoquastic vs non-Stoquastic Hamiltonians

Property	Stoquastic QA	Non-Stoquastic QA
Sign problem	Absent	Present
Classical simulability	Efficient (QMC)	Hard
Typical phase transition	First order	Potentially second order
Hardware availability	Current	Experimental
Advantage potential	Limited	High (theoretical)

E. Diabatic transitions

Quantum annealers are often possible in practice based on diabatic dynamics, even though an ideal model of quantum annealing is based on adiabatic dynamics, since diabatic transitions are regularly used to overcome energy barriers which would otherwise require prohibitive time scales in adiabatic dynamics. This is especially pronounced when the quantum phase transitions are being approached: first-order transitions are the ones where the energy gaps exponentially vanish, whereas second-order ones are the ones with the energy gap growing exponentially with a power form. Through either non-stoquastic driving Hamiltonians or specially constructed annealing programs, it is now possible to reduce first-order behavior to second-order and achieve large runtime reductions as a result, putting stoquastic systems into an intermediate regime between classical NP-hardness and quantum simulability[10].

III. ARCHITECTURE AND HARDWARE LANDSCAPE

Quantum annealers use continuous-time dynamics, which are defined by an Ising model, where system behavior is determined by a set of factors that include hardware connectivity, hardware noise, properties of coherence, and control param precision.

Currently surveyed information includes flux-qubit devices, the Lechner-Hauke-Zoller (LHZ) system, and Rydberg-atom systems as well as up-to-date control methods, including temporal pausing and reverse annealing, and the main noise sources that limit scalability.

A. Flux-qubit Based Annealers

Superconducting flux-qubit annealers, such as those made by D-Wave technologies, use transverse-field Ising Hamiltonians by encoding logic Ising spins via bistable Josephson-junction loops. During the last several years, the architecture of connectivity has evolved to the rather sparse Chimera networks up to the higher-density Pegasus and Zephyr (Advantage-2) networks, thus reducing the overheads of minor embedding. Simulations These systems realize logical spins by using ferromagnetically coupled systems of physical qubits and the quality of the embedding has a significant impact on system-wide performance[16].

B. Connectivity and Embedding Constraints

The physical connectivity of superconducting quantum annealers is limited, forcing the orthogonalization of logical spins onto chains of physical qubits and reducing the likelihood of chain breakage, reducing scales of effective energy, and limiting the number of qubits that can be usefully utilized. The resultant embedding schemes that are topology-sensitive are therefore needed, since the mismatch between the interconnectedness of the problem instance and the topology of the hardware can significantly impair computational cost[16].

C. Alternative Architectures

The Lechner-Hauke-Zoller architecture overcomes the connectivity limitations of scaling Ising couplings into local multi-body constraints therefore removing the explicit minor embedding requirement.The best annealing paths within the LHZ context may be designed through geometrical brachistochrone techniques which could be highly rapid. Furthermore, Rydberg -atom systems, with strong and tunable long-range interactions, allow even faster, coherent and non -stoquastic annealing that is beyond the abilities of flux -qubit systems[17],[16],[18].

Table 3: Architecture Comparison (Flux vs LHZ vs Rydberg)

Architecture	Native Connectivity	Control Type	Noise Regime	Maturity
Flux-qubit	Sparse	Analog	Thermal	Commercial
LHZ	All-to-all (logical)	Analog	Medium	Experimental
Rydberg	Long-range	Analog/digital	Low	Prototype

D. Control Features

Modern quantum annealers enable the complex control based on the time-dependent annealing schedules, A(s) and B(s), varying the transverse field driver and the problem hamiltonian respectively. Reverse annealing goes out of a classical solution, momentarily increases quantum fluctuations to do local searches, and then optimizes the solution. Empirical studies have also shown that reverse annealing can mitigate the harmful impact of first-order phase transitions, and in combination with an intermittent pause close to minimal spectral gaps, can exploit thermal relaxation processes to increase the chance of success[19], [20], [21].

E. Error Sources

The quantum annealing devices are vulnerable to perturbation due to thermal noise, decoherence, errors in the analog control of the devices, crosstalk between qubits, and measurement imprecision all of which cause distortion in the target Hamiltonian and resultant measurements. Such perturbations cause undesired excitations, parameter differences and incorrect categorization of final spin configurations, leading to hindrance of optimization performance as well as sampling fidelity. The modern state of art means that to mitigate the effects of quantum annealing they use quantum-annealing-optimistic error mitigation methods as zero-noise extrapolation using temperature control and energy-time rescaling, which seeks to approximate the behavior of a noiseless model[22].

IV. QA ALGORITHMIC STRATEGIES

Quantum annealing (QA) is not a specific algorithm, but is a collection of dynamical methods that are used to direct a quantum system to the low-energy states of a problem Hamiltonian.

The quality of exploration of complex energy manifolds of QA is controlled by algorithmic decisions, performance scheduling, a driver code, error-suppression codecs, and the application of hybrid classical-quantum feedback loops.

1) Standard QA: dynamics, tunneling, and the simulated-annealing analogy

Standard QA evolves the system under a time-dependent Hamiltonian of the form

$$H(s) = A(s)H_{\text{driver}} + B(s)H_{\text{prob}}, \qquad s \in [0,1], \#(3)$$

where H_{driver}(typically $\sum_i \sigma_i^x$) induces quantum fluctuations and H_{prob}encodes the optimization objective. In the prototypical transverse-field Ising setting, decreasing $A(s)$and increasing $B(s)$shifts weight from quantum delocalization to classical energy minimization. QA is often compared to simulated annealing (SA): SA uses thermal fluctuations to hop between basins, while QA leverages quantum tunneling to traverse barriers. In many rugged landscapes, tunneling can cross narrow but tall barriers that stymie thermal walkers, leading to different empirical scaling behaviors between QA and SA on some instance families[10].In a tunnel-favoring barrier-shape geometry, such as tunnel-favoring distances between two barriers, QA is superior to classical hill-climbing: in an entropic, artificially broad-basin geometry, thermal strategies are better, or hybrid workflows[23].

2) Advanced schedules: Non-linear, Adaptive, and Pausing strategies

Nonlinear annealing schedules redistribute time around critical gaps using problem-aware forms of $A(s)$and $B(s)$, slowing near bottlenecks and accelerating in slack regions.

Adaptive annealing dynamically adjusts schedules based on intermediate signals, offering potential quadratic improvements and reduced effective anneal times.

The pause–relax–resume strategy introduces a pause at $s^*$to exploit thermal relaxation and dissipation, empirically improving success rates near diabatic bottlenecks[10].

3) Reverse annealing: local refinement and hybrid workflows

Reverse annealing locates the system in a classical seed state, as required in standard annealing by uniform superposition, and adjusts the transverse field to be able to do localized quantum exploration, without using global superposition. As a quantum local-search operator, it optimizes around solutions that have been found using classical heuristics or forward annealing and has the ability to avoid suboptimal basins or fix unsatisfactory initial seeds by finding better regions of the solutions. These mixed methods which mix the forward and reverse annealing had been shown to be better than forward-only methods, with quantifiable improvements in portfolio optimization and matrix factorization.

4) Quantum annealing correction (QAC) and error suppression

Control errors, thermal fluctuations and decoherence in analog annealers can be prevented by quantum annealing correction (QAC), which suppresses errors instead of corrects them. Modest encoding of logical qubits with repetition or parity imposing penalty Hamiltonians are used; low energy excitations are therefore correctible logical error events. Common QAC protocols include energy penalty couplings, redundancy and majority or parity decoding encodings and classical post-readout decoding. QAC does not provide fault tolerance as does fault-tolerant circuit-based quantum error correction; however, it is highly scalable (near-term), in practice, with a particular improvement to hardware[4], [22], [24], [25].

5) Non-stoquastic drivers: power, simulation complexity, and open questions

By satisfying the condition that its off-diagonal terms are real and non-positive in computational basis, a Hamiltonian is a stoquastic which can be simulated by quantum Monte Carlo efficiently by avoiding the sign problem. Introduction of non-stoquastic drivers (e.g., improperly signed $\sigma_i^x \sigma_j^x$couplings)breaks this property, the dynamics acquired being classically hard to simulate. The spectral structure, e.g., first-order phase transitions being reduced to second-order transitions and exponentially small gaps being avoided, and allowing quantum interference and new tunnelling routes previously inaccessible to stoquastic annealing, can all be modified by such terms.

6) Trotterization, QAOA, and hybrid gate–analog strategies

Table 4: QA vs QAOA vs VQE (Algorithmic View)

Aspect	QA	QAOA	VQE
Model	Analog	Digital	Variational
Objective	Sampling	Approximation	Eigenvalue
Depth	Constant	Increasing	Variable
NISQ viability	High	Medium	Medium
Embedding	Required	Not needed	Not needed

Quantum annealing (QA) and the Quantum Approximate Optimization Algorithm (QAOA) could be considered as quantum copies in a consistent quantum control setting where QAOA will simulate transverse-field dynamics over time in the large-p (Trotterised) approximation. On the other hand, digitized annealing and diabatic pulse schedules are similar to optimized finite-depth QAOA, making it easier to hybridize the two modalities to transform their understanding. This structural parallelism has the strength of informing optimal QAOA parameter choices and encourages the use of the Hamilton reordering and Trotterisation methods to transform continuous time annealing into effective gate-based codes.[26], [24], [25], [27], [12], [13], [15], [23].

Table 5: Taxonomy of QA Algorithmic Variants

Strategy	Core Idea	Benefit	Cost	Best Use Case
Standard QA	Global anneal	Broad exploration	Gap sensitivity	Unknown optima
Pausing	Thermal relaxation	Higher success	Tuning	Small gaps
Reverse QA	Local refinement	Fast convergence	Seed needed	Hybrid loops
QAC	Encoding redundancy	Noise suppression	Qubit overhead	Noisy devices
Non-stoquastic	Gap reshaping	Potential speedup	Hardware	Hard instances

V. EMBEDDING AND MAPPING: THE UNSPOKEN BARRIER

A. QUBO Encoding Strategies

The minimization of QUBO and Ising Hamiltonians are an inherent feature of quantum annealers; however, realistic optimization problems involving integer or categorical variables need to be optimized with a binary representation first of all. The encoding scheme selected has a direct impact on the number of qubits needed, the scale of penalty terms, the accuracy of coefficient values, and finally whether a solution can be obtained or not. One-hot encoding is a simple and commonly used method of permutation problems, but it is resource intensive, that is, it requires many variables, high penalty strengths, and high hardware precision. Domain-wall encoding encodes integer values using binary-chain transitions and thus quadratic penalty terms and coupling are minimized in addition to allowing more effective high-throughput optimization[28], [29], [30], [31],[32]

B. Minor Embedding

After being formulated in QUBO or Ising form, a problem still needs to be implemented on the sparse architecture of physical quantum annealers. Due to the formulation of most optimization problems on dense or fully connected graphs most of the time direct implementation would be infeasible. Minor embedding alleviates this inconsistency in that the logical variables are all linked with a sequence of physical qubits that are contiguous. As a result of this, the problem Hamiltonian experimentally realized does not coincide with the desired logical Hamiltonian, in cases both of static distortion and dynamic noise responsiveness[33].

Table 6: QUBO Encoding Strategies

Encoding	Variables	Constraints	Precision Need	Scalability
One-hot	High	Quadratic	High	Poor
Domain-wall	Medium	Linear	Moderate	Better
Binary	Low	Higher-order	Very high	Risky

C. Mapping Costs and Hardware Overheads

Massive penalty coefficients induced by one-hot encodings, combined with long chains of qubits, often make use of hardware beyond the precision of the hardware, create noise, and make errors spread out along broken chains that must be post-processed. These are phenomena that compromise the fidelity of solutions to thermal fluctuations, control imprecision and variations between physical qubits that represent a single logical variable. Quantum annealing correction (QAC) deals with them by using repetition encoding and energy penalties, which experimentally increases the probability of successful ground states. QAC also increases performance on NP-hard planted problems when it is used in combination with embedding[29], [30], [33], [34], [35].

Table 7: Logical → Physical Qubit Inflation

Problem Type	Logical Vars	Physical Qubits	Compression Ratio
MaxCut	100	800	12.5%
TSP (20 cities)	400	3,000+	<10%
Portfolio (50 assets)	50	600	~8%

VI. BENCHMARKS: WHAT THE FIELD GETS WRONG

Quantum annealing is a controversial field, though its perceived results in terms of benchmarking and other tasks remain so. Although there are many articles that underlie better performance, lack of standard operating procedures, other inconsistent backgrounds and inadequate validation of experimental manipulations still present incoherent results. The current part outlines the appropriate benchmarking principles.

A. What to Benchmark and How

Controlled scaling studies are also easier with synthetic benchmarks like random spin -glass or planted -solution QUBOs, but typically scale poorly with real applications because of their highly idealized structural approximations. More balanced evaluation relies on systematic issues such as the travelling salesman problems, scheduling, knapsack and graph partitioning. To prove practical benefit, it is required to be tested on real-world examples (such as vehicle routing, portfolio optimization, logistics) which are typically addressed using hybrid quantum-classical algorithms and which are classically hard with or without QA. Strict solver parity, i.e. equal preprocessing, accuracy, restarts, and time budgets which are recognized to strongly depend on input/output overhead, model structure factors, and other systemic factors, which when ignored have an exaggerated effect on perceived quantum advantage, is required in rigorous benchmarking [37].

B. What Not to Do

Weak QA benchmarking often presents only best-case values, filled in without taking note of stochastic failure distributions or including median or percentile values. The other type of common flaw is cherry-picking of small cases that fit the hardware constraints but fail to scale, hence artificial benefits. In addition, much of the literature does not consider the effect of classical preprocessing (including decomposition, variable fixing or tabu filtering) which may dominate any observed improvement. [37], [36], [38].

Table 8: Benchmarking Do's and Don'ts

Practice	Correct Approach	Common Mistake
Runtime	Wall-clock	Anneal time only
Statistics	Median + variance	Best-case only
Baselines	Gurobi/Tabu/SA	Naive heuristics
Instances	Real + synthetic	Toy problems

C. Comparison Against Mature Classical Baselines

Benchmark constructive QA methods should be compared to classical solvers of industrial grade and not heuristics. Such comparisons involve simulated annealing, tabu search, and mixed-integer linear programming solvers (e.g. Gurobi, CPLEX, SCIP) that are sensitive to modelling and constraint-construction mistakes. Contemporary meta -heuristics, such as genetic algorithm, particle swarm optimization, ant-colony optimization, and differential evolution have strong global search algorithms and to this date, no QA problem has been proved to be consistently better than these protocols at solving real-world problems[39].

D. Empirical Evidence Today

QA can be useful when dealing with rugged energy landscapes that have narrow barriers, where tunnelling can be used; it can also be effective in sampling clustered low-energy basins when penalties are tuned, and can also be effective in hybrid workflows (as in tabu-QA or backbone-QA) compared to going straight to QA the ability to sample rugged energy landscapes by classical decomposition and pruning is in part responsible. On the other hand, QA has difficulty with sparsely connected penalty structures where tunnelling gains are reduced by penalties, integer-intensive formulations where QUBO blow-up and penalty saturation occur and when noise in analogue coefficient derivatives obscures fine scale differences between objectives.[38], [40].

Table 9: Where QA Works vs Fails

Landscape Feature	QA Performance
Narrow barriers	Strong
Dense constraints	Weak
Heavy penalties	Weak
Hybrid pipelines	Strong

VII. APPLICATION DOMAINS

The quantum annealing theory has since developed beyond the abstract analysis to have an application-focused optimization context; however, the utility of such techniques as a whole now is still subject to the compatibility of the problem structure, the overhead of problem encoding, and how many such problems can be scaled to using current quantum hardware.

A. Transportation and Logistics

1) Problem Formulation

Routing problems constitute a preeminent category of transportation optimization tasks, including the travelling salesman problem (TSP), the vehicle routing problem (VRP), the capacitated VRP (CVRP), and a range of trip optimization problems. These are issues of assigning discrete paths to minimize the cumulative travelling cost while meeting capacity, time-window, and service constraints. Since the feasible solution space of routing problems increases factorially with the number of nodes, routing problems are NP-hard.

2) QUBO Encoding

Routing problems are mapped using binary decision variables:

$$x_{(i,t)} = (1 \text{ if city } i \text{ is visited at time } t \ \& \ 0 \text{ otherwise})$$

Objective:

$$min \sum(i,j,t) d_{(ij)} \, x_{(i,t)} x_{(j,t+1)} \#(4)$$

Each city visited exactly once, one city per time step and Capacity constraints via auxiliary variablesThis leads to QUBOs with dense coupling matrices and large penalty coefficients.

3) Hardware Model

Clusters Superconducting quantum-annealers, such as the D-Wave system, are commonly used in hybrid classical-QA routing solvers, which use clustering heuristics and forward and reverse annealing regimes to optimize the solutions of the optimization problem.

4) Results

Hybrid QA routing solvers have a time constraint of creating paths of medium-size CVRP under a period of less than a second. With cold-chain VRP, the quality of routes calculated with QA techniques have been seen to be roughly 85 per cent of those calculated with conventional techniques and run-times are roughly 95 per cent shorter than genetic algorithms[41], [42].

5) Limitations

The erosion of quantum advantage in high-dimensional QUBO programs to encode highly descriptive constraints is exponential, and the number of ancillary variables is rapidly increasing with the number of planning variables, which are measured with physical quantities, in operation, such as time-window feasibility and multi-depot logistics.

B. Energy Optimization (Unit Commitment and Grid Control)

1) Problem Formulation

Unit commitment problem is the on-off problem of the cycles of power generators with reference to maximize the use of fuel, start-up and ramping costs at the same time keeping balance equations of power and reserve limitations. This optimization is normally formulated using binary, time dependent decision variables.

2) QUBO Encoding

Let $x_{g,t}$denote the on/off state of generator $g$at time t; the formulation includes linear operating costs, pairwise ramping penalties, and quadratic load-balancing constraints, with higher-order reliability constraints reduced via auxiliary binary variables.

3) Hardware Model

Next-generation scheduling techniques use global forward annealing to do coarse-grained optimization, reverse annealing to do fine-grained ramp refinement, and depend on mandatory hybrid decomposition methods to deal with complexity.

4) Limitations

Variable temporal coupling creates densely connected QUBOs that need penalty preciseness that is higher than the hardware precision capacity; hence, even grid-scale scheduling issues cannot be soluble in the near future.

C. Robotics and Motion Planning

1) Problem Formulation

The main areas of robotic optimization are motion-planning, task-scheduling, grasp-selection and stereo-correspondence.

2) QUBO Encoding

In each of these, binary variables are used to code the choice of path segments, the configurations of arm joints and the sequence of execution of tasks.

Objective:

$$\min \sum_i c_i x_i + \sum_{i<j} J_{ij} x_i x_j \#(5)$$

where $c_i$represents energy, distance, or execution cost.

3) Hardware Model

Superconducting QA on cloud-based systems is an integrated part of designed robot-QA co-simulation designs.

4) Results

The beneficial results of these integrated control architectures have been optimization of deformable-object manipulation, arm dynamics and stereo-matching using QUBO/Ising encodings[43].

5) Limitations

When a robot is dynamical, the size of QUBOs grows super-linearly with workspace resolution, and real-time operation of closed-loop control is impossible with existing hardware.

D. Finance: Portfolio Optimization

1) Portfolio optimization seeks:

Portfolio optimization aims to maximize the expected performance with minimum risk taking into consideration the transaction costs and portfolio rebalancing using binary variables to select the assets and discretized continuous weights.

2) QUBO Encoding

Binary asset-selection QUBO:

$$min \sum_i (-\mu_i) x_i + \sum_{i<j} \Sigma_{ij} x_i x_j \ (6)$$

Transaction costs and cardinality constraints encoded via quadratic penalties.

3) Hardware Model

The hybrid quantum-classical investment process combines classical wealth allocation algorithms with QA-derived asset selection, reverse annealing in refining portfolios, and iterative forward anneals in rebalancing portfolios.

4) Results

Experimental research has shown that hybrid QA-classical approaches are able to generate cumulative returns available to be above the benchmarks of traditional funds, and rebalanced QA portfolios are outperforming unadopted (static) QA portfolios[44].

5) Limitations

Asset correlations yield dense QUBOs, discretization of weights variables distorts the efficient frontier, and transaction cost model is highly sensitive with regard to the calibration of penalty terms.

E. Drug Discovery and Materials Design

1) Problem Formulation

Tasks in molecular design such as optimization of molecular structure and reaction pathways as well as design with properties can be re-formulated as discrete graph-labeling problems and as constraint-satisfaction problems**.**

2) QUBO Encoding

In this formulation, binary variables are used to encode the choice of atoms, bond structures and inclusion of substructures and the objective function provides penalties that apply to property violations, structural soundness, and stabilization issues.

3) Hardware Model

Hybrid quantum-classical machine-learning pipelines Hybrid quantum-classical machine-learning pipelines are based on an integration of QA-based discrete search and classical property evaluation modules.

4) Results

The hybrid quantum machine learning models improve molecular sub-structure recognition, but with fewer parameters than traditional methods, as well as, the predictive power of sub-structure models.

5) Limitations

Right assessment of molecular validity needs enormous penalty coefficients, which boosts the complexity of QUBO encodings concerning standard physics-simulation solutions; hence any quantum gain is territory-limited to small step frames of structure-selection.

F. Machine Learning: Feature Selection, Instance Selection, and Clustering

1) Problem Formulation

Essentially every discrete machine-learning problem, including feature selection, instance selection, and clustering can be formulated as a looking-glass problem of optimization, e.g. redundancy minimization, subset selection and medoid identification.

2) QUBO Encoding

Feature selection QUBO:

$$\min \sum_{i} w_i x_i + \sum_{i<j} r_{ij} x_i x_j \#(7)$$

where:

- w_i= feature relevance penalty
- r_{ij}= redundancy penalty

Clustering uses medoid-selection QUBOs.

3) Hardware Model

The workflow employed makes use of D-wave advantage system and the use of classical preprocessing, QA-based refinement, and post-processing decoding.

4) Results

Quantum annealing can be more feature- and cluster-compact and has better indicators of feature and cluster compactness as well as retrieval (nDCG and Davies Bouldin index) relative to simulated annealing even with natural hardware constraints[45].

5) Limitations

The effects of feature-correlation result in dense QUBOs, and the data scaling capabilities outrun embedding limits by forcing the use of QA to only refine an objective, and not fully solve it.

In routing, robotics, finance, drug discovery, and machine-learning applications, quantum annealing has been most likely to empirically prove useful as a hybrid local optimizer, used on data that has been classical preprocessed to minimize its search space and identify regions dominated by barriers.

Table 10: Application Domain Summary

Domain	Problem	Encoding	Hardware	Outcome
Logistics	VRP	One-hot	Flux QA	Hybrid gains
Energy	Unit commitment	Binary	Hybrid	Classical wins
Finance	Portfolio	Binary	Reverse QA	Competitive
ML	Feature selection	QUBO	QA	Compact models

VIII. QA VS OTHER QUANTUM OPTIMIZATION PARADIGMS

Quantum annealing is often introduced as a separate optimization theory; that its practical importance is best understood in terms of its contrasts to the gate-based variational algorithms like the Quantum Approximate Optimization Algorithm (QAOA) and the Variational Quantum Eigen solver (VQE).

A. QA vs QAOA

The quantum annealing (QA) system provides a continuous-time evolution using a Hamiltonian (where an evolution is performed by switching the transverse-field driver to the problem Hamiltonian). There it is the Quantum Approximate Optimization Algorithm (QAOA) which achieves a similar process in discrete time by Trotterising the evolution of gate-based unitarizes. Shallow QAOA can be considered a discretized approximation of QA where tunable angles are used in place of continuous time, or QA is an attempt to use fixed interactions of a hardware with no construction of an explicit logical circuit. Because of its sampling complexity, shallow effective depth and hardware scalability, QA now perform better on rugged Ising-type optimization problems than the shallow QAOA, but deep QAOA is infeasible on near-term-scale quantum (NISQ) systems[46].

B. QA vs VQE and Continuous-Variable (CV) Methods

Variational quantum eigensolver (VQE) is simply an approximation method-based eigenvalue solver; a prepared parameterized state is optimized to minimize the expected cost of a Hamiltonian. Optimization is performed in a continuous parameter space, in contrast to quantum annealing, which is the optimization of a discrete energy-space. QA is a natural generator of ensembles of low-energy discrete solutions and the visualizer of combinatorial energy landscapes, which makes it a good model in decision problems. Conversely, VQE applies to a single quantum state when used in physics and chemistry, and the design of ansatz is primarily related to its performance; it is not until it has been heavily discretized that it becomes combinatorial[47].

C. Current Limitations

Even with twenty years of progress and scaling of hardware, quantum annealing now faces structural, instead of technological, limits, as logical qubits (usually a fraction less than ten per cent of physical qubits) are typically embedded due to constraints of embedding as well as the density of the problem. High-order constraints, scaling to high-order constraints and imposing auxiliary-variable blow-up and exponentially increasing resources, are a greater constraint to scaling than the qubit number. Being open quantum systems, the annealers are prone to thermal noise, drift, and crosstalk, as well as readout errors, and even the mitigation that could be applied does not guarantee the large-scale reliability fully. Constant benchmarking bias and exponentially null gaps in stoquastic QA on most of NP-hard problems prevent scalable speedups, so it may be important to consider non-stoquastic drivers[48],[49].

D. How to systematically characterize "annealing hardness"?

The quantum annealing has become a fully-fledged albeit untested optimization paradigm. No matter how much the complex deployments are rising, the basic questions concerning complexities, control, benchmarking, and hybrid integration are yet to be answered. This section contains a research agenda, which is an enumerated list of concrete problems, which are open, and aimed at defining the next decade of QA research.

E. How to systematically characterize "annealing hardness"?

The key open problem is to find a problem-free concept of annealing hardness, as QA hardness is no longer just a function of problem size, but also spectral gaps, the geometry of energy barriers, its basin entropy and quantum phase transition structure.

F. Can non-stoquastic drivers give scalable performance advantages?

It seems that non-stoquastic Hamiltonians provide a theoretical solution to the exponentially small gaps which limit stoquastic QA on many NP-hard problems. It is however still unclear, whether they are scale independent in improving gaps, whether they can be executed with hardware stability and control sufficient and whether they are also scale independent in the presence of finite temperature, noise and control errors.

G. What is a fair and universal quantum advantage metric?

No clear definition of quantum advantage in quantum annealing based optimization exists because every of the following quantities describe the solver: time-to-solution, approximation quality, sampling diversity, and energy efficiency contains solver bias. Standardized multi-objective wall-clock-normalized benchmarks with rigorous statistics and hybrid performance measurement are needed to make progress; otherwise, their quantum supremacy in optimization remains unopposed.

H. What are the best architectures combining QA, QAOA, and classical heuristics?

The most significant progress of quantum optimization is accomplished with hybrid pipelines incorporating classical pre-processing, barrier-dominated search QA, and digital refinement (QAOA / VQE), and classical post-processing. Realizing scalable advantage does not lie in the number of qubits but in co-developed advancements in the hardness theory, non-stoquastic control, automatic embedding, rigorous benchmarking, and interdisciplinary cooperation among physics, computer science, optimization, and industry.

Table 11: Open Problems & Research Directions

Open Question	Why It Matters	Who Should Solve
Annealing hardness	Predict feasibility	Theorists
Non-stoquastic QA	Possible speedup	Physicists
Embedding automation	Scaling	Engineers
Advantage metrics	Credibility	OR + QC
Hybrid architectures	Practical impact	Industry

IX. CONCLUSION

The purpose of this review was to combine theory, hardware, algorithms, embedding, benchmarking, applications, ecosystem integration, and prospects to provide an all-inclusive, critical, and unifying synthesis of quantum annealing to solve combinatorial optimization.

We prove the problem of quantum annealing to be essentially a physics-based optimization heuristic, and is separate to both classical meta-heuristics and variational quantum algorithms based on gates. Current points of scale are suppressed by embedding and encoding, instead of rawness of qubits. Hybrid workflows are giving the most significant empirical benefits, and not standalone quantum annealing. Unless benchmarking of parity, wall-clock normalization and statistical rigor is available to the assertion, it should be done with caution. The most promising technical directions to overcome a set of limitations are non-stoquastic control, hybrid QA-QAOA integration, and automated embedding. Perhaps most importantly we show that quantum annealing is no longer evaluated in a vacuum but rather than a stack of quantum-classical optimization it should be considered as a stack of physics, algorithms, and engineering inseparably coupled.

Quantum annealing is yet to be a universal optimization engine, but it is already a useful scientific tool to investigate the energy landscapes that otherwise can only be explored through computations. The final effect it has will depend not so much on the rate of qubits but on a clever merging with the theory, algorithms, and classical solvers.

Availability of Data and Materials

No new data were generated or analyzed during this study. All information discussed in this review is based on previously published literature, which has been appropriately cited.

Competing Interests

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

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Authors' Contributions

Rudraksh Sharma conceived the structure of the review, conducted the comprehensive literature survey, synthesized and analyzed the material, and wrote the main manuscript. Ravi Katukam provided overall guidance on topic selection and conceptual direction, contributed to critical discussions,reviewed and revised the manuscript for intellectual content. Arjun Nagulapally contributed to refining the scope of the review, assisted in structuring key sections, and provided critical feedback during manuscript revision. All authors read and approved the final version of the manuscript.

Acknowledgements

The authors acknowledge AIONOS for providing a supportive research environment and encouraging interdisciplinary exploration that facilitated the development of this review. The authors also thank colleagues and peers for valuable discussions and feedback that helped refine the scope and clarity of the manuscript. Additionally, the authors acknowledge the developers and maintainers of open-access scientific literature and computational tools that supported this work.

Consent to Publish

Consent to publish: not applicable.

Consent to Participate

Consent to participate: not applicable.

Ethics Declaration

Ethics declaration: not applicable.

REFERENCES

[1] S. Chowla and H. J. Ryser, "Combinatorial Problems," Canadian Journal of Mathematics, vol. 2, pp. 93–99, 1950, doi: 10.4153/cjm-1950-009-8.
[2] F. Glover, G. Kochenberger, and Y. Du, "Quantum Bridge Analytics I: A Tutorial on Formulating and Using QUBO Models." [Online]. Available: https://www.nap.edu/catalog/25196/quantum-computing-progress-and-prospects
[3] B. A. Cipra, "An Introduction to the Ising Model," The American Mathematical Monthly, vol. 94, no. 10, pp. 937–959, Dec. 1987, doi: 10.1080/00029890.1987.12000742.
[4] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse Ising model," 1998.
[5] S. Morita and H. Nishimori, "Mathematical foundation of quantum annealing," J Math Phys, vol. 49, no. 12, 2008, doi: 10.1063/1.2995837.
[6] A. Peruzzo et al., "A variational eigenvalue solver on a quantum processor," Apr. 2013, doi: 10.1038/ncomms5213.
[7] E. Farhi, J. Goldstone, and S. Gutmann, "A Quantum Approximate Optimization Algorithm," Nov. 2014, [Online]. Available: http://arxiv.org/abs/1411.4028
[8] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, "Quantum annealing: An overview," Jan. 23, 2023, Royal Society Publishing. doi: 10.1098/rsta.2021.0417.
[9] R. Barends et al., "Digitized adiabatic quantum computing with a superconducting circuit," Nature, vol. 534, no. 7606, pp. 222–226, Jun. 2016, doi: 10.1038/nature17658.
[10] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, "Perspectives of quantum annealing: Methods and implementations," May 01, 2020, Institute of Physics Publishing. doi: 10.1088/1361-6633/ab85b8.
[11] L. P. Yulianti and K. Surendro, "Implementation of Quantum Annealing: A Systematic Review," 2022, Institute of Electrical and Electronics Engineers Inc. doi: 10.1109/ACCESS.2022.3188117.
[12] S. Bravyi, "Monte Carlo simulation of stoquastic Hamiltonians," Jan. 2015, [Online]. Available: http://arxiv.org/abs/1402.2295
[13] M. Marvian, D. A. Lidar, and I. Hen, "On the computational complexity of curing non-stoquastic Hamiltonians," Nat Commun, vol. 10, no. 1, Dec. 2019, doi: 10.1038/s41467-019-09501-6.
[14] S. Bravyi, D. P. Divincenzo, R. Oliveira, and B. M. Terhal, "The Complexity of Stoquastic Local Hamiltonian Problems," 2007.
[15] W. Vinci and D. A. Lidar, "Non-stoquastic Hamiltonians in quantum annealing via geometric phases," npj Quantum Inf, vol. 3, no. 1, 2017, doi: 10.1038/S41534-017-0037-Z.
[16] A. W. Glaetzle, R. M. W. Van Bijnen, P. Zoller, and W. Lechner, "A coherent quantum annealer with Rydberg atoms," Nat Commun, vol. 8, Jun. 2017, doi: 10.1038/ncomms15813.
[17] F. Leitner and G. B. Mbeng, "Quantum adiabatic brachistochrones for the LHZ architecture," 2023.
[18] H. Zhao, X. Wu, W. Li, X. Fang, and T. Li, "Phase-Slip Based SQUID Used as a Photon Switch in Superconducting Quantum Computation Architectures," Electronics (Switzerland), vol. 13, no. 12, Jun. 2024, doi: 10.3390/electronics13122380.

[19] M. Ohkuwa, H. Nishimori, and D. A. Lidar, “Reverse annealing for the fully connected p -spin model,” Phys Rev A (Coll Park), vol. 98, no. 2, Aug. 2018, doi: 10.1103/PhysRevA.98.022314.
[20] M. S. Jattana, “Reverse quantum annealing assisted by forward annealing,” Aug. 2024, doi: 10.3390/quantum6030030.
[21] E. J. Crosson and D. A. Lidar, “Prospects for quantum enhancement with diabatic quantum annealing,” Jul. 01, 2021, Springer Nature. doi: 10.1038/s42254-021-00313-6.
[22] M. H. Amin et al., “Quantum error mitigation in quantum annealing,” Nov. 2023, [Online]. Available: http://arxiv.org/abs/2311.01306
[23] B. Tasseff et al., “On the emerging potential of quantum annealing hardware for combinatorial optimization,” Journal of Heuristics, vol. 30, no. 5–6, pp. 325–358, Dec. 2024, doi: 10.1007/s10732-024-09530-5.
[24] K. L. Pudenz, T. Albash, and D. A. Lidar, “Quantum annealing correction for random Ising problems,” Phys Rev A, vol. 91, no. 4, Apr. 2015, doi: 10.1103/PhysRevA.91.042302.
[25] Y. Nambu, “Error correction for encoded quantum annealing revisited,” Jul. 2024, [Online]. Available: http://arxiv.org/abs/2407.15480
[26] C. Ho Thanh, “Hamiltonian Reordering for Shallower Trotterization Circuits.” [Online]. Available: https://github.com/altaris/pauli-coloring-benchmark.
[27] R. Acharya et al., “Quantum error correction below the surface code threshold,” Nature, vol. 638, no. 8052, pp. 920–926, Feb. 2025, doi: 10.1038/s41586-024-08449-y.
[28] P. Codognet, “Comparing QUBO models for quantum annealing: integer encodings for permutation problems,” International Transactions in Operational Research, vol. 32, no. 1, pp. 18–37, Jan. 2025, doi: 10.1111/itor.13471.
[29] M. Kumagai, F. Takano, T. Araki, M. Sato, and H. Kobayashi, “An External Definition of the One-Hot Constraint and Fast QUBO Generation for High-Performance Combinatorial Clustering,” 2021. [Online]. Available: www.ijnc.org,
[30] S. Kikuchi, K. Takahashi, and S. Tanaka, “Performance of Domain-Wall Encoding in Digital Ising Machine,” 2025, doi: 10.1109/ACCESS.2023.0322000.
[31] P. M. Xavier, Y. Park, D. Noble, I. Santana, M. Paredes, and D. E. Bernal Neira, “A performance comparison of variable encoding techniques for QUIO and QUBO problems.”
[32] P. Codognet, “Encoding Binary Comparison Constraints in QUBO for Quantum Annealing,” in GECCO 2025 Companion - Proceedings of the 2025 Genetic and Evolutionary Computation Conference Companion, Association for Computing Machinery, Inc, Aug. 2025, pp. 2419–2423. doi: 10.1145/3712255.3734298.
[33] J. Raymond, N. Ndiaye, G. Rayaprolu, and A. D. King, “Improving performance of logical qubits by parameter tuning and topology compensation,” in Proceedings - IEEE International Conference on Quantum Computing and Engineering, QCE 2020, Institute of Electrical and Electronics Engineers Inc., Oct. 2020, pp. 295–305. doi: 10.1109/QCE49297.2020.00044.
[34] K. L. Pudenz, T. Albash, and D. A. Lidar, “Error-corrected quantum annealing with hundreds of qubits,” Nat Commun, vol. 5, Feb. 2014, doi: 10.1038/ncomms4243.
[35] W. Vinci, T. Albash, G. Paz-Silva, I. Hen, and D. A. Lidar, “Quantum annealing correction with minor embedding,” Phys Rev A, vol. 92, no. 4, Oct. 2015, doi: 10.1103/PhysRevA.92.042310.
[36] E. Osaba, E. Villar-Rodriguez, I. Oregi, and A. Moreno-Fernandez-De-Leceta, “Focusing on the hybrid quantum computing - Tabu search algorithm: New results on the Asymmetric Salesman Problem,” in GECCO 2021 Companion - Proceedings of the 2021 Genetic and Evolutionary Computation Conference Companion, Association for Computing Machinery, Inc, Jul. 2021, pp. 1476–1482. doi: 10.1145/3449726.3463123.
[37] A. N. Ignatov and S. V. Ivanov, “COMPARING THE SOLVERS FOR THE MIXED INTEGER LINEAR PROGRAMMING PROBLEMS AND THE SOFTWARE ENVIRONMENTS THAT CALL THEM,” Bulletin of the South Ural State University, Series: Mathematical Modelling, Programming and Computer Software, vol. 17, no. 3, pp. 57–72, Aug. 2024, doi: 10.14529/mmp240305.
[38] M. Gou et al., “Hierarchical Quantum Optimization via Backbone-Driven Problem Decomposition: Integrating Tabu-Search with QAOA,” Apr. 2025, [Online]. Available: http://arxiv.org/abs/2504.09575
[39] V. Tomar, M. Bansal, and P. Singh, “Metaheuristic Algorithms for Optimization: A Brief Review,” Engineering Proceedings, vol. 59, no. 1, 2023, doi: 10.3390/engproc2023059238.
[40] A. Bonomi, T. De Min, E. Zardini, E. Blanzieri, V. Cavecchia, and D. Pastorello, “Quantum Annealing Learning Search Implementations,” Dec. 2022, doi: 10.26421/QIC22.3-4-1.
[41] S. Mario et al., “Quantum Annealing based Hybrid Strategies for Real Time Route Optimization.”
[42] E. W. H. Chow, G. T. S. Ho, V. Tang, and M. M. F. Tam, “Utilizing Quantum Annealing to Address Vehicle Routing Challenges in Cold Chain Logistics,” Association for Computing Machinery (ACM), Apr. 2025, pp. 137–142. doi: 10.1145/3760622.3760626.
[43] “ENHANCING ROBOTIC MANIPULATION THROUGH FULLY INTEGRATED QUANTUM ANNEALING: THE FIQA-RM SYSTEM.”
[44] S. N. Morapakula, S. Deshpande, R. Yata, R. Ubale, U. Wad, and K. Ikeda, “End-to-End Portfolio Optimization with Quantum Annealing,” Apr. 2025, [Online]. Available: http://arxiv.org/abs/2504.08843
[45] C. Pomeroy, A. Pramov, K. Thakrar, and L. Yendapalli, “Quantum Annealing for Machine Learning: Applications in Feature Selection, Instance Selection, and Clustering Notebook for the QuantumCLEF Lab at CLEF 2025,” 2025. [Online]. Available: https://github.com/dsgt-arc/
[46] E. Pelofske, A. Bärtschi, and S. Eidenbenz, “Short-depth QAOA circuits and quantum annealing on higher-order ising models,” npj Quantum Inf, vol. 10, no. 1, Dec. 2024, doi: 10.1038/s41534-024-00825-w.
[47] J. Li, “Knowledge Distillation Inspired Variational Quantum Eigensolver with Virtual Annealing,” May 2025, [Online]. Available: http://arxiv.org/abs/2505.03998
[48] F. A. Quinton, P. A. S. Myhr, M. Barani, P. Crespo del Granado, and H. Zhang, “Quantum annealing applications, challenges and limitations for optimisation problems compared to classical solvers,” Sci Rep, vol. 15, no. 1, Dec. 2025, doi: 10.1038/s41598-025-96220-2.
[49] V. Choi, “Limitation of Stoquastic Quantum Annealing: A Structural Perspective,” Sep. 2025, [Online]. Available: http://arxiv.org/abs/2509.16265

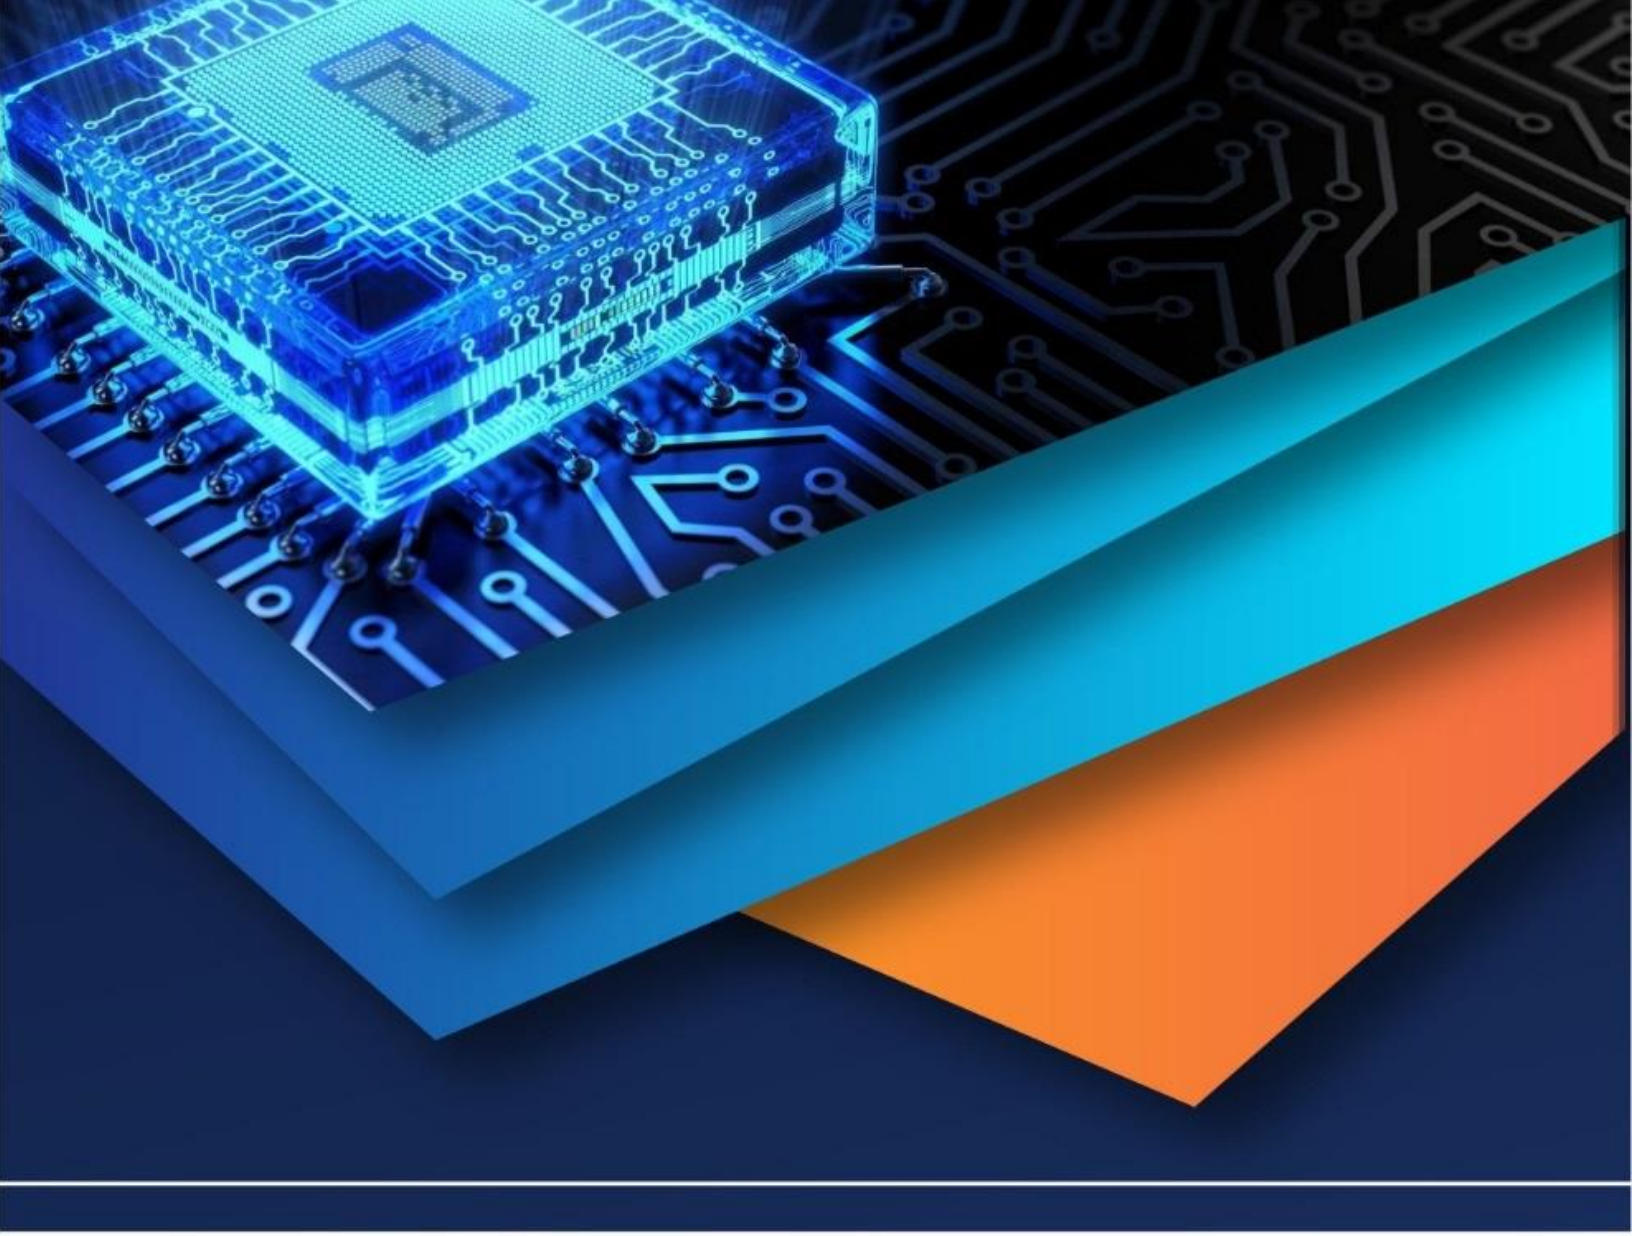

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY